Time, Hope, and Independence: An Argument for More Structure in Decision Theory

Goreti Faria

A thesis submitted to the Department of Philosophy, Logic and Scientific Method of the London School of Economics and Political Science for the degree of Doctor of Philosophy, September 2016.

Declaration

I certify that the thesis I have presented for examination for the MPhil/PhD degree of the London School of Economics and Political Science is solely my own work. The copyright of this thesis rests with the author. Quotation from it is permitted, provided that full acknowledgement is made. This thesis may not be reproduced without my prior written consent. I warrant that this authorisation does not, to the best of my belief, infringe the rights of any third party.

I declare that my thesis consists of 34996 words.

Abstract

My thesis explores alternatives to the orthodox model of decision theory. I criticise it by focusing on motivations that, for different reasons, I believe should not be labelled as a consequence. I start by describing what I call the orthodox theory of choice. I describe both the theories proposed by Von Neumann and Morgenstern (1944), and Savage (1954). This is followed by a discussion of the individuation strategy, and, in particular, John Broome's proposal of an individuation by justifiers (Broome (1991)). I then focus on the Allais's problem, and on how individuation can be employed to solve the problematic situation in which it puts orthodox decision theory. I argue against the practice of using this strategy as a general solution to cases such as the Allais's problem. I then extend the concerns raised by the Allais's problem to motivations that do not concern probabilities. I focus on dynamic decision making, and on a property of certain actions that arises from the passage of time: hope. This is interesting because it is no longer the probabilities that affect the decision maker, but an element that is not even represented in the standard model, namely time. I then describe the Kreps and Porteus's model (Kreps and Porteus (1978)) for a preference for the timing of the resolution of uncertainty, where a preference for either early or late resolution of uncertainty is modelled explicitly. I show how making time an explicit part of the model allows one to model utility as depending on something other than consequences, while not violating dynamic consistency. I use the above case to claim

that in some contexts - for example, if the decision maker is deciding for herself, and time is passing - it is preferable to model concerns that do not quite fit the label of a consequence explicitly, because the benefits of doing so surpass the costs. This seems to indicate that decision theory should be moving towards a pluralist approach, where different models are used depending on the decision context.

Acknowledgements

I would like to thank the Portuguese Science and Technology Foundation (FCT) for the Ph.D. grant SFRH/BD/78906/2011, the Department of Philosophy, Logic and Scientific Method of the London School of Economics and Political Science, and my supervisors Richard Bradley and Katie Steele.

CONTENTS

LIST OF FIGURES

"I find many men in our dangerous age who seem to be in love with misery and death, and who grow angry when hopes are suggested to them. They think hope is irrational and that, in sitting down to lazy despair, they are merely facing facts. I cannot agree with these men. To preserve hope in our world makes calls upon our intelligence and our energy. In those who despair it is frequently the energy that is lacking."

Bertrand Russell (1950, p. 700)

1

INTRODUCTION

1.1 INTRODUCTION

This thesis is about hope, or at least that's what I tell people when they ask me what I am writing my PhD on. In truth, hope is just one example used to argue against the consequentialist nature of orthodox decision theory. It is however an important example that I believe does a good job at conveying one of the main points of this thesis: not everything that motivates decision makers should be classified as a consequence. José Saramago, who praises hope and dreams about a better world throughout his entire work, writes that "[h]ope is like salt, there's no nourishment in it but it gives the bread its taste" (Saramago, 2004, p. 38). That is what I want to say, I do not believe that hope's place is in a consequence (the "nourishment"), but I believe hope imprints something non-consequential on an alternative (the "taste") and it should be represented as such.

Normative decision theory – which is the theory under discussion in this thesis – aims at answering the question of how people **should** make decisions. Orthodox decision theory answers this question with the prescription that agents ought to prefer, among a set of alternatives, that one with the highest expected value. The expected value of an alternative, as the name implies, is the sum of its different possible values weighted by the probability of each value obtaining. How are the different possible values represented? In orthodox decision theory these values are represented by the utility function, which arguably represents agents' desires regarding the **consequences** of the different possible alternatives.[1] Therefore, apart from probabilities, which have a weighting-only role, the only thing that ought to play a role in decision making is the agents' desires with respect to the set of consequences associated with the different alternatives. As a result, everything that motivates decision makers - everything they care about and everything they desire and hope for - must be classified as a consequence. I disagree with this approach.

However, disagreeing with theories that classify everything as a consequence is a difficult stance to hold and to argue for, as it always seems to be possible to explain away any clashes with the theory by maintaining that the agent in question faces a different problem than the one modelled, and that by correctly re-describing the

1 That utility represents desires is not at all uncontroversial, but that discussion is not within the scope of this thesis.

consequences the theory can remain unchallenged. This is one of the goals of this thesis: to argue that it is indeed problematic to classify certain things as consequences, and to justify this claim with more than an angel-like intuition. The other goal is to offer alternatives to orthodox decision theory.

These two goals are intertwined. As I said above, arguing against the practice of *consequentialising* everything is difficult, and I believe that no argument can fully show it to be mistaken. The best one can hope for, then, is to criticise it and offer an alternative theory that is, all things considered, more attractive. That is precisely the strategy I follow in this thesis: I begin by discussing the problems with the strategy of describing everything that matters to decision makers as consequences, and I then propose an alternative model that captures considerations concerning time as something other than a consequence. In particular, this model allows preferences for the moment in time at which uncertainty about a decision is resolved to be modelled explicitly. This provides a way to explicitly model hope, as one way to think about hope is as a preference for leaving uncertainty unresolved for longer. What is more, I believe this model has the properties one should want it to have.

The significance of this project lies in the fact that orthodox decision theory is broadly used, and it influences decision making and economics in the real world. Furthermore, I think it also influences how people think and reason about their own decision making and that of others (why else would my economist mother keep telling me that 'there are no free lunches'?). It seems to me that making considerations such as hope part of an immense set of consequences disguises them as something less special and makes them less salient. It is important to recognise hope for what it is, and make it as salient as possible. This is what the quote motivating this thesis by Bertrand Russell conveys: it is important to be hopeful if we want things to change. This thesis is my effort to make people more hopeful.[2]

1.2 SUMMARY OF CHAPTERS

This thesis is composed of the following chapters.

Chapter 2: Orthodox Decision Theory

In chapter 2 I briefly describe what I above called orthodox decision theory. Arguably, the two most well-known axiomatic approaches are von Neumann and Morgenstern's (Von Neumann and Morgenstern (1944), henceforth vNM) and Savage's (Savage (1954)). The main difference between these two approaches regards the way probabilities are featured and thus interpreted. The interpretation of probabilities however is not crucial for my thesis, and therefore the difference in the type of probabilities will not be particularly relevant. I use them in turn in the following chapters, depending on the issue being discussed. I describe both these theories informally, with an

2 Noam Chomsky urges us to do the same when he says that "[i]f you assume that there is no hope, you guarantee that there will be no hope." It seems then that the most salient hope is, the more hope there will be too.

emphasis on the issues that are going to be relevant for this thesis. Additionally, I show how the two features of orthodox decision theory under discussion throughout the thesis are shared by both theories. These are: that on the desire side the only primitive is the consequences, and that the utility function can ultimately only depend on consequences.

Chapter 3: Against Individuation

In chapter 3 I argue against the individuation strategy. Focusing on Savage's decision theory, I describe the Allais Problem, which is an example of a decision situation that casts doubt on one of the main axioms of orthodox theory: the Sure-Thing Principle (and the independence axiom, in vNM's framework). It does so because most people's intuitions are in agreement with a pattern of choices that is inconsistent with the axiom. I will explain this case in detail in chapter 3, but the general idea is that it is a preference for a certain alternative (over a risky/uncertain one) that is causing such a pattern of choice. Since there is nothing irrational about an agent preferring an option that delivers a consequence with certainty (especially in the realm of Humean decision theory - a decision theory that does not assess desires, but rather the means used to satisfy them), orthodox theory might be in trouble.

The individuation strategy is the orthodox way to respond to the Allais problem.[3] This strategy consists in individuating consequences more finely to include the mo-

3 Another alternative would be to hold that the Allais preferences are irrational, but I am here dismissing this possibility because I am concerned with ways in which we can accommodate these preferences. Additionally, I do think - as most people do - that the Allais preferences are intuitively very reasonable, and that it would be inconsistent to exclude such preferences in the context of Humean decision theory.

tivation causing the problem. In the Allais case this is done by making certainty a part of all of the consequences associated with the certain alternative. In chapter 3 I describe and evaluate Broome's version of the individuation strategy, which I believe to be the most defensible version of the strategy. I make a comparative evaluation of Broome's proposal. When one is faced with cases such as the Allais problem, there are mainly two routes that can be taken to avoid the paradoxical outcome: keep the orthodox model of decision theory, and individuate outcomes more finely, or have a finer model of decision making, and coarser individuation of outcomes. I believe that it is not the case that the latter is always preferred to the former, but that this is the case for certain motivations and contexts. In particular, in chapter 3 I argue against the individuation strategy when global properties of actions are concerned, which are properties of actions that only arise due to the conjunction of all the outcomes obtained in the different states of the world (such as the properties of certainty and hope).

I am not looking for an argument that will completely dethrone individuation. Even in contexts where I argue against the use of individuation, my argument is - in my view - not enough to get the orthodox model to be abandoned in those situations. I need to additionally propose alternatives and try to show that their advantages surpass their disadvantages in comparison to an orthodox individuated theory. Decision making is a very complex process, and there will not be one perfect normative theory, but rather one that is slightly better than another in a given context. Therefore, the next two chapters are concerned with exploring an alternative to the orthodox model when hope-type motivations are influencing the decision makers, and showing how this model's advantages surpass its disadvantages with regards to the orthodox model and in context where time and hope matter.

Chapter 4: Kreps and Porteus: An Example

In chapter 4 I describe Kreps and Porteus's (Kreps and Porteus (1978), henceforth KP) model of decision making. Their model offers an axiomatic approach to dynamic decision problems that allows to explicitly model preferences with regards to when uncertainty is resolved. For example, if one has a preference for alternatives that allow hope to be kept for longer, then one would have a preference for late resolution of uncertainty.

I explain the model in natural language and, as much as possible, in very simple terms. I also try to make the motivation behind it clearer. I then discuss how it differs from the orthodox model: it ceases having utility depend only on consequences by adding a function that depends on time, but it also adds time as an explicit part of the model. In orthodox decision theory preferences over alternatives are assumed to be fully determined by decision makers' desires on expected consequences. What is more, the model only includes alternatives, probabilities, and the utility and probability functions. KP add a further element to these: time. I explain how time is added as an explicit part of the model, and how utility can then be made to explicitly depend on time.

Chapter 5: Dynamic Consistency and Independence

In chapter 5, I discuss the disadvantages of the KP model, and argue that they are not as problematic as could initially be thought and that the use of this framework to model motivations such as hope is overall preferable to the orthodox individuated model. In order to show that such is the case, I answer the following questions: which of the orthodox axioms are violated by KP's model?, and how problematic is such a violation?

To answer the first question I start by explaining the orthodox conditions of rationality imposed on dynamic decision making. I discuss the independence axiom in the context of dynamic decisions, and how independence follows by implication from the remaining conditions imposed on dynamic choice. One of the main worries with models of dynamic decision making that deviate from the standard axioms is that they might violate one condition in specific: dynamic consistency (what I choose tomorrow is what I would choose to choose tomorrow if I was forced to make tomorrow's choice today). Therefore, in order to preserve the viability of KP's model as an alternative to individuation, an important goal of this chapter is to show that KP do not, in fact, violate dynamic consistency. This is what I show in chapter 5. I also stress the fact that adding structure to the model in the form of time being added as a primitive is what allows for one to have utility depend on something other than consequences while not violating compound independence nor dynamic consistency. I claim that the condition that is indeed violated is a reduction axiom. I discuss the importance of this violation and argue that, in certain contexts, it makes sense to give up this condition.

Chapter 6: Conclusion

I conclude by suggesting that no model should be used all the time across all decision contexts. For example, in the contexts that KP are trying to address, it seems reasonable to use a class of models that does not subscribe to a reduction axiom. However, if the context is a very simple static one, then perhaps the simplicity of the orthodox model is desirable. This suggests the use of a pluralist decision theory that uses different models according to the context the decision is taking place.

2

ORTHODOX DECISION THEORY

2.1 INTRODUCTION

I take orthodox decision theory to include all decision theories that subscribe to *maximisation of expected utility*. When interpreted from a normative standpoint - which is the interpretation being discussed throughout this thesis - this means that in order for an agent's preferences to be rational, she ought to prefer - among a set of alternatives - that with the highest expected value. The expectation is a reflection of probabilities, which are either exogenously given, or a reflection of agents' beliefs, depending on the precise model used. The value is a reflection of agents' desires, and represented by a utility function. It is with the latter side of the theory that is the focus of this thesis.

Decision theory encompasses many different decision contexts and situations. One important qualification is whether a decision maker is facing decisions under certainty, risk, or uncertainty (as divided, for example, by Luce and Raiffa (1957). Decision making under certainty describes situations where each possible alternative leads to one and only specific consequence. When on the other hand the state of the world is not known, there is another element that becomes relevant: the probability. Decision making under risk describes situations where each alternative leads to a set of possible consequences, and each consequence within the set has a known probability of occurring. Decision makers are given and therefore know what these probabilities are. Finally, decision making under uncertainty describes a situation where alternatives lead to a set of consequences, but the probabilities of each of these consequences obtaining are not known to the decision maker.

Arguably, the two most well-known versions of orthodox decision theory are von Neumann and Morgenstern's (Von Neumann and Morgenstern (1944)) and Savage's (Savage (1954)). The former applies to decision making under risk, and the latter to decision making under uncertainty. The main difference between these two approaches regards the way probabilities feature and how they are thus interpreted. vNM present an axiomatic theory where the probabilities are objective (i.e., independent of the decision maker), and Savage a theory with subjective probabilities (i.e., internal to the decision maker) that are endogenously provided by the model. The role and interpretation of probabilities is not crucial for this thesis, and therefore this difference will not be particularly relevant. However, due to this difference, the two theories make use of slightly different frameworks, and use different terminologies and notations.

I will use both theories throughout this thesis, according to the one I consider most appropriate for the reader's understanding of the point that is being made. Therefore, in this chapter I describe both these theories informally, with an emphasis on the issues that are going to be relevant throughout this thesis. Additionally, I draw a connection between the main axioms from both theories that are going to be under discussion in the following chapters. The general framework used when *dynamic decisions* are concerned is also going to be relevant for this thesis, but I will leave its definition and presentation for chapter 4, where dynamic decision making is first discussed.

This chapter is organised as follows. In the next section I introduce decision under certainty. In section 3 I discuss vNM, and in the following section Savage. In section 5 I compare vNM's and Savage's frameworks with an eye on the issues discussed in this thesis. I then conclude.

2.2 DECISIONS UNDER CERTAINTY: BERNOULLI UTILITY

When decisions are made under certainty then the utility function is the only element needed for one to make a decision. A decision is made over a set of alternatives $X = x, y, z, ...$ Each agent has a preference relation over the different possible alternatives, which is represented by $\succcurlyeq$, and stands for a weak preference (this means that if x is weakly preferred to y then the agent either prefers x or is indifferent between the two alternatives). If the agent respects a number of axioms on this preference relation, then her preferences can be represented by an ordinal utility function. This utility function is called a Bernoulli utility function, named after Daniel Bernoulli, who in 1738 argued that "the value of an item must not be based on its price, but rather on the utility it yields" (Bernoulli (1738)). This function represents preferences ordinally: it can be shown that it is unique up to monotonic transformations, which are changes that preserve the order of preferences. For example, if one prefers x to y then all the utility functions that represent one's preferences will keep x above y in one's preference ranking.

The Axioms and Representation Theorem

Which axioms does the relation need to respect in order for the utility function to be reached? The first two axioms are:

- **Completeness**: If $x, y \in X$, either $x \succcurlyeq y$, or $y \succcurlyeq x$, or both.
- **Transitivity**: For $x, y, z \in X$, if $x \succcurlyeq y$ and $y \succcurlyeq z$, then $x \succcurlyeq z$.

According to orthodox theory, for a preference relation to be rational it is enough that it satisfies these two axioms, and it is in this way that rationality is defined. In fact, without any utility connotation, Cantor (1895) proved that if a binary preference relation $\succcurlyeq$ on a countable set $X = x, y, z...$ satisfies completeness and transitivity, then there exists a function f such that $x \succcurlyeq y$ if and only if $f(x) \geq f(y)$.

However, if we allow for the set of consequences to be uncountably infinite, then not any rational preference relation can be represented by a utility function (only the reverse holds).[4] In order for it to be represented by a utility function one technical axiom needs to be fulfilled:

- **Continuity**: For any $y \in X$, the sets $\{x \in X : x \succcurlyeq y\}$ and $\{x \in X : y \succcurlyeq x\}$ are closed.

Given these three axioms, Debreu (1954) generalised Cantor's result, and accompanied it by a change in the interpretation of the function on the left-hand side so that it becomes a utility function.

Theorem Debreu. Suppose that $X = \mathbb{R}_0^n$ and that preferences are complete, transitive, and continuous, then there exists a continuous function $U : X \to \mathbb{R}_0^n$ such that, for any $x, y \in X$,

$$U(x) \geq U(y) \text{ if and only if } x \succcurlyeq y$$

4 The lexicographic ordering on $\mathbb{R}^2$ defined by $(x_1, x_2) \succcurlyeq (y_1, y_2)$ iff either $x_1 > y_1$ or ($x_1 = y_1$ and $x_2 \geq y_2$) is complete and transitive but admits no utility representation if we allow for the set of alternatives to be uncountable. A proof can be found in Mas-Colell et al. (2006).

2.3 DECISIONS UNDER RISK: VNM, 1944

When decisions are made under risk, it is still the case that decision makers are choosing one of several available alternatives. However, these alternatives are now named lotteries, which are probability distributions. This is due to the fact that, as mentioned above, now there is a set of possible consequences associated with each alternative, and decision makers are choosing one of possible probability distributions over this set of consequences. For example, suppose one can choose between playing a game that offers £6 if a fair dice comes out 2 and £1 otherwise, or a game that offers £2 if a fair coin comes out heads and £1 if it comes out tails. In this latter case one is choosing among the following lotteries (where in $(x, y; z, w)$, x and z stand for consequences, and $y \in [0, 1]$ and $w \in [0, 1]$ the probability of obtaining the outcomes x and z respectively):

$$A = (6, \tfrac{1}{6}; 1, \tfrac{5}{6}) \text{ and } B = (2, \tfrac{1}{2}; 1, \tfrac{1}{2})$$

A lottery assigns a probability for all the possible outcomes of choosing that lottery. For example, suppose I choose lottery A. This lottery is formally represented by the vector $(\frac{1}{6}, \frac{5}{6})$. More generally: there is a set of alternatives associated with each decision problem under risk. Each alternative is a different probability distribution over the outcomes in this set (some distributions might very well attribute a probability of 0 to some of the outcomes in the set obtaining), and a decision maker must choose one of these probability distributions (represented by vectors such as the one above). Mostly everything that holds for the certainty case also holds for the risk case, provided alternatives are now represented as lotteries, and the set of axioms is slightly modified to accommodate probabilities. If the agent's preferences obey this revised set of axioms, then these preferences can be represented by a more informative utility function, which will now be called expected utility.

The idea is as follows: in the case of decisions under certainty the agent has a Bernoulli utility function, which assesses the desirability of an alternative by applying that utility function to that alternative's associated consequence. In the case of decisions under risk the agent has that same Bernoulli utility function, which once again evaluates consequences. However, since alternatives no longer lead to consequences with certainty, in order to assess an alternative the Bernoulli utility is now multiplied by the probabilities of the different consequences associated with an alternative occurring. This results in the vNM utility function on lotteries: $EU = \sum_{i=1}^{N} p_i u_i$, and vNM requires one to choose the lottery that maximises this quantity.

Axioms and Representation Theorem

Something needs to be said about the new set of axioms governing preferences over lotteries. Completeness, and transitivity are maintained. The Archimedean Axiom is added, and it implies continuity of the lottery space.[5] A second axiom is introduced, which is the central axiom in vNM and it almost became a synonym of

5 Archimedean Axiom: For all lotteries $L, L', L'' \in \mathcal{L}$ (a set of such lotteries), if $L \geq L' \geq L''$ then there exists $\alpha, \beta \in (0,1)$ such that $\alpha L + (1-\alpha)L'' \geq L' \geq \beta L + (1-\beta)L''$.

the theory itself: the independence axiom. Let L, L', and L'' be lotteries, and $\mathcal{L}$ a set of such lotteries, from which the agent is choosing.

(4b) Independence: Preferences satisfy independence if for any $L, L', L'' \in \mathcal{L}$ and $\alpha \in [0, 1]$:

$$L \succcurlyeq L' \text{ iff } (\alpha L, (1-\alpha)L'') \succcurlyeq (\alpha L', (1-\alpha)L'')$$

Adding this axiom to the other three guarantees that the expected utility function is linear in the probabilities, which informally means that expected utility is the result of Bernoulli utilities weighted linearly by the probabilities. This will become clearer in the last section of this chapter, when I discuss the shape of the expected utility functional and the separability of preferences.

There is a further axiom (or assumption - it is not quite clear how it is meant to be perceived) that is necessary on the way to obtaining a representation theorem. This is the reduction of compound lotteries, which says that any compound lottery (a lottery over lotteries) can be reduced to a simple one (a lottery over outcomes) by operating with the probabilities according to ordinary probability calculus. If we allow $\mathcal{L}$ above to include both simple and compound lotteries, then this axiom is already encoded in the independence axiom. I will came back to the reduction assumption on chapter 5.

Given these axioms, vNM derived a representation theorem, which guarantees the existence of an expected utility function:

Theorem vNM. If completeness, transitivity, continuity, and independence hold, then there exists a vector of utilities $(u_1, ..., u_N)$ such that for any two lotteries $L = (p_1, ..., p_N)$ and $L' = (p'_1, ..., p'_N)$:

$$L \succcurlyeq L' \text{ iff } \sum_{i=1}^{N} p_i u_i \geq \sum_{i=1}^{N} p'_i u_i$$

The independence axiom guarantees that utility is invariant up to positive affine transformations, and therefore that utility is cardinal.

2.4 decisions under uncertainty: savage, 1954

Savage's framework includes acts (or actions), states of the world, and consequences. States of the world S are mutually exclusive and jointly exhaustive descriptions of different features of the world that are relevant for the decision in question and regarding which there is uncertainty (and it is assumed that any partition of states is acceptable). The probability function has as its domain precisely these states of the world. Elements of S are specific states of the world and are represented by $s, s', ...$ Events are subsets of the set S, and I henceforth use E to represent an event. Consequences are a set F that includes all possible consequences of different actions under all possible states of the world. Elements of F are represented as $x, y, ...$ Finally, acts (Savage's alternatives) are functions $f, g, ...$ from the set of states of the world S to the set of consequences F. Therefore, $f(s)$ also represents a specific consequence, just like x and y. The set of actions A is the set that contains these functions. An agent will have as primitive a preference relation $\succcurlyeq$ on the set of acts, and a set of axioms will be imposed on this preference relation such that, if it satisfies these axioms, then a utility and a probability function can be derived.

The theory applies in situations of choice under uncertainty, and it adds to the utility function a probability function, which represents the agents' beliefs much in the same way that the utility function represents agents' desires.[6] When both of these functions are given, subjective expected utility (SEU) of actions can be obtained: $\int_S u(f(s))dP$, where P is a probability measure on the set of events and u the utility function. Then Savage's theory claims that one ought to choose the action with the highest subjective expected value.[7]

Axioms and Representation Theorem

What conditions are required in order for preferences to be numerically represented by a SEU function? The first axiom simply states that the preference relation $\succcurlyeq$ is complete and transitive. We are familiar with these basic axioms from the previous section, and they apply to Savage's acts similarly to how they apply to vNM's lotteries. There are several other axioms, both regarding the utility and the probability function.

6 It is not at all uncontroversial that these functions represent beliefs and desires respectively. There are many different interpretations regarding the relationship between these functions and mental states, and some authors fail to acknowledge that the relationship exists. Nonetheless, there is a tendency to assume that these functions do represent beliefs and desires, and I will do the same in this thesis.

7 Savage's theory belongs to the set of SEU theories. Similar theories were developed by others - for example, Ramsey (1926), Jeffrey (1965), Anscombe and Aumann (1963).

However, I will mainly focus on those on the utility side, and I will try not to go into unnecessary details.

With an eye on getting to a utility function on consequences, Savage defines a preference relation over consequences from a preference relation over constant acts. Constant acts are those that have the same consequence under all states of the world. The definition is the following:

Definition (Savage's D2): $y \succcurlyeq y'$ if and only if $f \succcurlyeq f'$ when $f(s) = y$ and $f'(s) = y'$ for every $s \in S$.

Then there is one axiom - the sure thing principle, or STP - that is used to ensure that if two acts differ only in their consequences in some state then the one with better consequences in that state is preferred, and there is another axiom - state independence, or SI - that this ranking is state-independent.

The STP states that the evaluation of acts should not depend on the events in which they agree, but only in those in which they don't.

Savage's STP: Let E and $\neg E$ be mutually exclusive and exhaustive events, let a_1, a_2, a_3, and a_4 be four available actions, and let the consequences for these actions be as follows:

	E	$\neg E$
a_1	c_1	c_3
a_2	c_2	c_3
a_3	c_1	c_4
a_4	c_2	c_4

Figure 2.1: Consequence Matrix to illustrate Savage's STP

Then $a_1 \succcurlyeq a_2$ iff $a_3 \succcurlyeq a_4$.

In other words, this axiom guarantees that if two acts differ only in their consequences in some state then the one with better consequences in that state is preferred. This axiom also relies on the assumptions that actions do not affect the probabilities of the states of the world. Without this assumption a_1 could change the probabilities in E in a different way than a_3, and therefore the conclusion that $a_3 \succcurlyeq a_4$ would no longer follow.

SI means that the preference relation is based on the consequences, and the state under which these consequences obtain has no influence on the preference ranking. Before the axiom a further definition is needed though (**D3**): the definition of a null event. It says that an event $E \subseteq S$ is called null if for all f and g from A, $f \sim g$ given E. Finally, the axiom.

Savage's SI: If E is not null and if $f(s) = x$ and $g(s) = y$ for all $s \in E$, then $f \succcurlyeq g$ given E iff $x \succcurlyeq y$.

Informally, state independence says that only the consequences play a role in preferences, not the state of the world in which they occur.

Savage then derives a representation theorem using all his axioms $P1 - P7$, of which I described only two here:

Theorem Savage. a preference $\succcurlyeq$ satisfies $P1 - P7$ if and only if there is a finitely additive probability measure P on the set of events and a function $u : C \rightarrow R$ such that for every pair of acts f and g

$f \succcurlyeq g$ iff $\int_S u(f(s))dP \geq \int_S u(g(s))dP$

And P is unique and u is unique up to positive affine transformations.

This result depends on an additional assumption that Savage does not explicitly mention - Broome (1991) calls this assumption $P0$, or the rectangular field assumption (RFA). I will discuss this assumption in the following chapter, and therefore I will leave it unattended for now.

2.5 SAVAGE, VNM, AND THIS THESIS

Throughout this thesis I use both vNM's and Savage's framework. This is because the two features of orthodox decision theory under discussion throughout this thesis are quite general, and therefore are shared by both theories. They are: that on the desire side the only primitive is the consequences, and the elements that the utility function can depend on (which is both a consequence of the axioms and of the first feature stated). I will here briefly explain how these issues are addressed in the chapters that follow, and how they apply to both Savage and vNM.[8]

First of all, it is important to note that it is indeed quite hard to compare the technical details of these two theories because they are defined in different frameworks: one framework uses objective probabilities and does not model states of the world, and the other uses subjective probabilities and states are modelled explicitly. Savage's theory includes additional axioms that are related precisely to the fact that it uses subjective probabilities and to the existence of states of the world in the model. However, one can try to understand in which ways they differ in light of what is going to be discussed throughout this thesis.

In terms of structure, both Savage's and vNM's decision theory have similar ones: they include the same basic functions - probabilities and utilities -, alternatives, and consequences. Savage's theory has more structure: the primitives are not only consequences (as in vNM), but also states of the world, there is a distinction between acts, states of the world, and consequences, and finally the probability function represents agents' beliefs (conversely, vNM take probabilities as something that is exogenously given). However, this is not what I am concerned with: as I said before, in this thesis I mostly ignore the belief side of decision making and focus on the desire side (and the

8 Throughout this thesis I will sometimes use *alternatives* as a general term that encompasses both vNM's lotteries and Savage's acts. I will additionally use *outcomes*, *prizes*, and *consequences* interchangeably.

added elements in Savage's theory come from the fact that probability is subjective). In chapters 4 and 5 I am concerned with the possibility of something other than consequences and probabilities to influence the expected value of alternatives, and for that to be the case additional primitives would need to be added explicitly to the existing ones in both vNM's and Savage's models.

In terms of the utility function over alternatives, and what is allowed to influence it and in what way, both of these theories subscribe to the separability of preferences, which is one of the defining characteristics of orthodox decision theory. I will start by briefly explaining what this means.

Recall that the expected utility of an action (in Savage) or a lottery (in vNM) is roughly given by $\sum_{i=1}^{n} U(c_i) \cdot p_i = U(c_1) \cdot p_1 + ... + U(c_n) \cdot p_n$, where c_i are the different outcomes, and p_i are the probabilities of each of these outcomes obtaining. This means that the utility U of each consequence is not influenced by the probability attached to the state of the world in which that consequence occurs, and that each consequence/probability pair is independent of all the others. This is known in the literature as *preference separability*. Stated more formally, preference separability is equivalent to the following two things: (1) the contribution of each utility/probability pair to the overall sum that constitutes expected utility is independent of the other utility/probability pairs, and (2) concerns the relationship between the utility and the probabilities within each pair. It says that for each utility/probability pair, the utility is independent of the probability.[9] This is what makes it the case that only what is classified as a consequence can influence the agent's Bernoulli utility, and that the probabilities can only play the role of weighting the utilities in expected utility.

Let us now turn to the finer details where we see how the theories differ. The expected utility function has the shape described in the previous paragraph *if and only if* the decision maker's preferences obey the axioms described above. Here is where the theories differ.

In Savage, there are two axioms that play this role: the sure-thing principle (STP), and state independence (SI). Savage starts by defining a preference relation over consequences, and the SI axiom guarantees that this ranking is state independent and the STP guarantees that if two acts differ only in their consequences in some event then the one with better consequences in that event is preferred. With regards to separability, SI guarantees that the state of the world in which a consequence occurs does not affect how that consequence is assessed (component 2 above) and the STP guarantees the separability of probability-utility pairs (component 1). Recall that the STP states that the evaluation of acts should not depend on the events in which they agree, but only in those in which they do not. In other words, this axiom guarantees that if two acts differ only in their consequences in some event then the one with better consequences in that event is preferred. This is equivalent to component (1) of separability

9 Machina (1989) calls (1) replacement separability (RS) and (2) mixture separability (MS). He discusses these concepts in the context of vNM and lotteries only, and therefore one should be careful in extrapolating these definitions to discuss the two frameworks simultaneously. However, for reasons of tractability, I will sometimes refer to these two concepts using Machina's terminology.

because by demanding that the decision maker does not pay attention to events where consequences are the same one is effectively saying that what happens in other events does not influence our assessment of the current state.

In vNM, independence guarantees component (1) of separability quite straightforwardly. As described above, independence says that if a lottery is preferred to another then a combination of that lottery with some other must also be preferred to the same combination applied to the lottery that it was preferred to. Therefore, one pair cannot affect the other, otherwise this would not necessarily be the case. As for component (2), it is not as straightforward, but it is also guaranteed by independence. One of the lemma's used to prove vNM's representation theorem is the following:

Lemma: Suppose that there is a preference relation $\succ$ on a set of lotteries $L = X, Y, Z, ...$ that satisfies vNM's axioms. Then, the following holds: If $X \succ Y$ and $0 \leq a < b \leq 1$ then $bX + (1-b)Y \succ aX + (1-a)Y$.

This lemma says that if a lottery is preferred to another then so is the compound lottery that gives the agent the preferred lottery with a higher probability. This seems obvious, but it is only true if utilities are not affected by probabilities, i.e., if component (2) of separability is preserved. This lemma can be proved using only vNM's independence axiom.[10] Therefore, independence guarantees the two components of separability. This seems to already indicate that there is at least a sense in which the STP is weaker than independence.

10 This is part of the proof of vNM's representation theorem

This is interesting because the independence axiom is often used as synonym for vNM's theory, and in the same way the STP as synonym for Savage's theory. Simultaneously, the two principles are often referred to (even if implicitly) in the literature as identical. However, given the above, they differ in some respects. Indeed, it is possible to concoct a decision situation where independence places restrictions on what the agent can do but the STP remains silent. Consider the following decision matrices:

	S1 ($\frac{1}{2}$)	S2 ($\frac{1}{2}$)
f	a	b
g	c	a

	S1 ($\frac{1}{2}$)	S2 ($\frac{1}{2}$)
f'	a^*	b
g'	c	a^*

Figure 2.2: Savage's STP and vNM's Independence Axiom

Representing these actions using vNM results in the following lotteries: $f = (a, 1/2; b, 1/2)$, $g = (c, 1/2; a, 1/2)$, $f' = (a^*, 1/2; b, 1/2)$, and $g' = (c, 1/2; a^*, 1/2)$. Now suppose that b is preferred to c. By independence we can derive the following biconditional: $b \succcurlyeq c \leftrightarrow b \cdot \frac{1}{2} + a \cdot (1 - \frac{1}{2}) \succcurlyeq c \cdot \frac{1}{2} + a \cdot (1 - \frac{1}{2})$. This is true for any outcome written instead of a. Therefore, it follows that $f \succcurlyeq g \leftrightarrow f' \succcurlyeq g'$.

However, the STP cannot say anything regarding this situation. Recall that the STP is a conditional that has as the antecedent that two acts differ only in their consequence in some event. This is not true in this case, and therefore, given that the antecedent is false, the consequence does not need to be true. Savage's theory taken as a whole would say the same as vNM's, but the STP alone does not say anything in this case.

In this thesis my main concern is not the subtle differences between these theories, but the fact that they both only allow for expected utility to depend on outcomes (and linearly on probabilities - i.e., probabilities are treated as weights). Both of them have this feature, even if achieved by a combination of different axioms. What is more, I focus on a violation of these theories brought into evidence by the Allais problem, which violates component (1) of separability and therefore both Savage's STP and vNM's independence. It is important to keep in mind that the two principles are not identical, but what is going to be said in the following chapter applies to both axioms.

In chapters 4 and 5 two-stage lotteries play an important role, and there is another distinction that is important to make. So far the discussion has been restricted to one-stage lotteries, and reduction of compound lotteries to simple ones plays no role in this analysis. In such a setting there might be a sense in which the STP is actually stronger than independence, but I will say more about this at the end of chapter 5 when discussing the applicability of what was said to the STP.

2.6 conclusion

In this chapter I prepared the reader for what is to be discussed in the next chapters. I described the orthodox theory of choice by describing two of the most used frameworks: Savage's and vNM's framework. I then established a relation between the two and made clear how my criticism in this thesis applies rather to something that both frameworks share.

3

AGAINST INDIVIDUATION

3.1 INTRODUCTION

In this chapter I discuss a strategy used in decision theory that came to be known as *the individuation strategy of outcomes* (or simply the individuation strategy, or individuation, and sometimes reindividuation). In particular, I discuss the use of this strategy as a solution to the so-called Allais Problem, which is an objection to orthodox decision theory.

Orthodox decision theory represents desires via a unique utility function on consequences.[11] In the Allais case the same consequences are assessed differently in slightly different choice contexts, which contradicts the existence of a unique utility function that maps consequences into the real numbers, and that is linear in the probabilities.

11 Unique up to positive affine transformations.

Therefore, this case presents a problem for the theory. Different possible ways out have been proposed. The one I discuss in this chapter is the individuation strategy, which consists in individuating the consequences more finely to make them distinct from each other, so that the domain of the utility function is actually different in the two choice contexts and therefore it is no longer the case that the same consequence is assessed differently.

Individuation is usually charged of stripping decision theory of any normative force. If every time a pattern of preferences violates the theory one redefined the consequence set so as to prevent such a violation, then it does not seem that the theory is making any prescriptions on decision making behaviour. Given this line of criticism, different authors have proposed restrictions to individuation, to avoid that decision theory collapses into this overly permissive framework. One such author is Broome (1991).

In this chapter I describe and evaluate Broome's proposal. Cases like the Allais problem show that there are intuitive violations of one of the main axioms of orthodox decision theory, and if individuation is shown to be problematic then there is a greater case for considering alternative models to the orthodoxy. I make a comparative evaluation of the strategy. When one is faced with cases such as the Allais problem, there are mainly two routes that can be taken to avoid the paradoxical outcome: keep the orthodox model of decision theory, and individuate outcomes more finely, or have a finer model of decision making (such as the one discussed in the next chapter), and coarser individuation of outcomes. The point I want to make is not that the latter is always preferred to the former, but that this is the case when the properties of the actions influencing the decision maker are of a certain kind, and additionally there is a plausible modelling alternative (the existence of which will be discussed in the following chapters). In particular, in this chapter I argue against the individuation strategy when global properties of actions are concerned, which are properties of actions that only arise due to the conjunction of all the outcomes obtained in the different states of the world (but that the individuation strategy nonetheless attributes them to particular outcomes).

However, even when global properties are concerned, the alternative to individuation might not be good enough. There are advantages and disadvantages to taking any one of these two routes. In this chapter I discuss the disadvantages associated with the individuation strategy, and in the following chapters I argue against some of the disadvantages usually attributed to finer models of decision making. I eventually make a case for a pluralist decision theory, where the model used depends on the specific decision context and situation. In this chapter I use Savage's decision theory because I believe that explicitly modelling the states of the world makes it easier to understand what is going on in the Allais case. This chapter is organised as follows. In the next section I describe the Allais problem, and the motivation behind the pattern of preferences exhibited by decision makers. I then describe how the individuation strategy can be used as a way out of the Allais problem. In section 5 I argue against the individuation strategy. I then conclude.

3.2 ALLAIS'S PROBLEM AND A PREFERENCE FOR CERTAINTY

Savage's Sure-Thing Principle (STP), formally defined in the previous chapter, is one of the most important axioms in Savage's theory, and it is needed for the general separable form of the utility function to be derived, as explained in the previous chapter (in particular, the STP implies what I referred to in the previous chapter as component (1) of preference separability). In Savage's framework, states of the world represent all things about the world that can affect the outcome of a given action. They are mutually exclusive and jointly exhaustive. It is in terms of these states that the STP is defined. The STP requires that a preference between acts depends solely on the consequences in states in which the payoffs of the two acts being compared are distinct. This implies that the valuation of the consequences of an act in one event is independent of the consequences of that same act in the complementary event.

The intuition behind the STP is that what happens in a given state of the world can be evaluated separately from what happens in the other states of the world. However, is the evaluation of an outcome in a given state of the world really independent of what happens in other states of the world? Is this a reasonable assumption to make?

The literature in decision theory offers plenty of examples that seem to contradict this separability between states of the world. Many authors argue that in fact outcomes in different states of the world are often complementary of each other, and given these complementarities, the STP is no longer a reasonable axiom. The most well-known example of this is the Allais problem. Allais (1953) produced an example where people tend to choose in a way that goes against the spirit of the STP. In his example, agents have to make two choices: one between *a* and *b* (decision 1), and one between *c* and *d* (decision 2).

	A (0.89)	B (0.1)	C (0.01)
a	1*m*	1*m*	1*m*
b	1*m*	5*m*	*Nothing*
c	*Nothing*	1*m*	1*m*
d	*Nothing*	5*m*	*Nothing*

Figure 3.1: Consequence Matrix for the Allais's Problem

Allais's intuition, as well as Savage's and many others that were faced with this problem, is that decision makers will prefer *a* to *b* and *d* to *c*. This pattern of choices can be explained by t he fact that action *a* gives the agent one million pounds for sure, and this seems to be important enough for the agent to dismiss the possibility of getting five million pounds. However, there is no certainty in action *c*, and therefore the possibility of five million pounds by choosing *d* gains relevance. This is a violation of the STP because agents are not ignoring the event in which the consequences are the same for both actions, as the principle says they should. If they were, then the consequences under state of the world A could be dismissed, and action a would become indistinguishable from c, and action b indistinguishable from d.

This inconsistency cannot be solved by appealing to the particular shape of the expected utility function. If one initially prefers a to b this must mean, according to the expected utility hypothesis, that the expected utility of a is bigger than the expected utility of b. Given how expected utility is constructed, this means that $u(1m) > u(1m) \cdot 0.89 + u(5m) \cdot 0.1 + u(0) \cdot 0.01$. This is not problematic in itself, but once we put this together with the preference between c and d, a contradiction arises. If d is preferred to c then again it follows that the expected utility of d is higher than the expected utility of c, and therefore that the following also holds: $u(0) \cdot 0.9 + u(5m) \cdot 0.1 > u(0) \cdot 0.89 + u(1m) \cdot 0.11$. If these two preferences are held together, then the following biconditional must be true: $u(1m) \cdot 0.11 > +u(5m) \cdot 0.1 + u(0) \cdot 0.01 \iff u(1m) \cdot 0.11 < u(5m) \cdot 0.1 + u(0) \cdot 0.01$. However, this biconditional is clearly a contradiction (it says that $x > y \iff x < y$ and these two things can never have the same truth value), no matter which utility function is used to represent agents' desires.

Why do people exhibit this pattern of preferences in the Allais problem? It seems to be the case that people have a preference for a property of action a that action c does not have, which is the property of being risk free. This property of action a cuts across the different states of the world, and therefore does not respect the separability of preferences. In fact, a simpler example can demonstrate the same thing. This case is useful because it is indeed very simple, but for it to work one needs to make an assumption regarding the decision maker's utility function, while in the Allais case this was not even necessary (no function could elicit the pattern of preferences exhibited by the decision maker). However, now that we know that agents can exhibit a preference for the global property of certainty, it is reasonable to make use of the following simpler example.

Example 2: Movies and Music

Suppose Anna is trying to decide between two actions: buy a ticket to an indoor film festival, or buy one to an outdoor music festival. She is uncertain about the weather though, and although Anna will always be warm at the film festival, she will be cold at the music festival if the weather is bad. Anna prefers being warm to being cold. Additionally, all things equal, Anna prefers music to films. Below is Anna's decision matrix.

	Good Weather	Bad Weather
(a) Music Festival	Music While Warm	Music While Cold
(b) Film Festival	Movie While Warm	Movie While Warm

Figure 3.2: Example 1: Movies and Music

If Anna was an expected utility maximiser, then she would buy the ticket for the music festival if and only if the probability of good weather was above a certain threshold value that would make the expected utility of the music festival larger than that of the film festival. Now, for the sake of discussion, let us assume that the probability of good weather p is indeed larger than the threshold value. If this is the case then the expected utility of choosing the music festival is higher than the expected utility of choosing the film festival and therefore, if she was a utility maximiser, Anna would choose to buy tickets for the music festival. However, suppose that Anna has a preference for the global property of certainty, just like the Allais choosers. She dislikes risk and therefore going to the film festival attracts her because it is risk free. Suppose this preference is strong enough to overtake the above considerations. Again, what increases the desirability of going to the film festival is the fact that it results in the same outcome in all the states of the world.

Both of these examples demonstrate the existence of interactions or complementarities between the states of the world. In other words, there are properties of the actions available to the decision maker that cut across the different states of the world. The decision maker is exhibiting a preference towards this property, and not towards a specific outcome. I have been calling these properties global properties of actions, following Buchak (2013). Global properties are, according to Buchak, "features [of gambles] that do not supervene on any particular outcome" (p. 28). She discusses these properties in the context of valuing certainty, just as per the examples above. However, I want to use the term global properties more widely: instead of qualifying these properties as properties of gambles, I would like to qualify them as properties of actions more generally. A gamble is a probability distribution over outcomes, and therefore it cannot have any property - for example, it cannot have properties that are related to the passage of time, such as the property of allowing hope to be kept for longer. One could very well have a preference for the action that resolves uncertainty later rather than sooner (and allows one to mantain hope for longer), even though everything else about the actions is the same.

If one wants to explicitly represent these complementarities, or the sensitivity of preference to global properties of acts, then this effectively means giving up on the STP. If the way a given outcome is valued depends on outcomes associated with other states of the world, then clearly the STP is not being respected. One can however preserve the STP and instead individuate the outcomes more finely. In what follows I describe the individuation strategy and argue against its use in the context of the examples here described.

3.3 THE INDIVIDUATION STRATEGY

One possible way to deal with the apparent violations of the STP is to claim that the outcomes are not well represented, and that the complementarities observed between states of the world should somehow be included in the outcomes. This is usually referred to as a refinement of the individuation of outcomes, and the general strategy of doing so is called an individuation strategy. In this section I describe the individuation strategy in the context of violations of the STP.

Starting with Allais's problem, one might say that the agent is not facing those outcomes represented in the consequence matrix represented in 3.1, but that she is rather facing a matrix of the following kind:

	A (0.89)	B (0.1)	C (0.01)
a	$1m \wedge y$	$1m \wedge y$	$1m \wedge y$
b	$1m$	$5m$	*Nothing*
c	*Nothing*	$1m$	$1m$
d	*Nothing*	$5m$	*Nothing*

Figure 3.3: Revised consequence matrix for the Allais's problem

The difference here with respect to the original matrix is that the consequences associated with action a were redescribed to include an additional amount y. This was done with the justification that the outcomes associated with an action which is risk free are different from those which are not, and that this should be specified in the outcome description. Then it is assumed that the agents with Allais preferences get utility from this property and therefore the utility of the redescribed outcomes is higher than the utility of the original ones.

This change solves the two problems above. First, the pattern of preferences consisting on preferring a to b in decision 1 and d to c in decision 2 will no longer result in a contradiction. Recall that in the original problem these preferences resulted in the following contradiction: $u(1m) \cdot 0.11 > +u(5m) \cdot 0.1 + u(0) \cdot 0.01 \iff u(1m) \cdot 0.11 < u(5m) \cdot 0.1 + u(0) \cdot 0.01$. Now the expected utility of a is given by $u(1m + y)$ and a contradiction can no longer be derived as this term will not appear in the expected value of any of the other options. Second, the STP no longer dictates that column A should be ignored when choosing between *a* and *b*, because the consequences for these two actions under *A* are no longer the same. As a result, the choice between *a* and *b* does not become identical to that between *c* and *d*.

With regards to the second example, the same strategy can be applied. Assuming again that the agent gets something out of the risk free action *b*, one can add this to the original matrix:

	Good Weather	Bad Weather
(a) Music Festival	Music While Warm	Music While Cold
(b) Film Festival	Movie While Warm $\wedge$ Certainty	Movie While Warm $\wedge$ Certainty

Figure 3.4: Revised consequence matrix for the Music/Film decision

A common concern with the individuation strategy is that it makes decision theory, to the extent that it relies on the STP, vacuous. If every time a pattern of preferences violates the STP, one redefines the consequences so as to prevent the STP from being violated, then any preference exhibited by a decision maker can be made consistent with the STP. As a result, the STP does not constrain behaviour in any way. Samuelson (1952) for example has said that "[i]f every time you find an axiom falsified, I tell you to go to a space of still higher dimensions, you can legitimately regard my theories as irrefutable and meaningless" (pp. 676-77).

To address this concern, different authors that support the individuation strategy have suggested the placement of different restrictions on the strategy. For example, Allais (1953) defends that consequences should only be differentiated if they differ in terms of money prizes, and if this is the case then the Allais problem will certainly pose a problem to the orthodox theory. However, this seems to be too restrictive. Machina (1981) argues in favour of distinguishing outcomes in terms of physically observed aspects. If there are no physical differences then there is no room for making two outcomes different. Sen (1993) argues that everything in the real world (except in the mind) can be used for differentiating outcomes. These are two similar proposals but again, I believe these restrictions to be too strong. There are two other well-known proposals: Pettit's and Broome's, which I believe to be more promising. I will briefly describe Pettit's proposal, but I focus on Broome's. This is because I believe the former is no longer an individuation strategy, but closer to what I will discuss in the next chapter, which is the strategy of adding structure to the model of decision making. I still describe it here as support for what is to come in the following chapter, and I then move on to Broome's proposal.

Pettit (1991) suggests that one individuates by properties, and defends the claim that outcomes should be differentiated only if there is a difference in the properties that the agent is concerned about . Properties are for Pettit the fundamental unit of desires, and each action has several different properties. On both cases described above, decision makers had a preference for the risk-free option. One of the options had the property of being risk-free, and that property was desired by the decision maker, and therefore it influenced her choice. Pettit argues for the thesis that decision theory neglects the assumption of desiderative structure, and that such an assumption is a reasonable one to make. The assumption of desiderative structure, in Pettit's words, is "that, not only do we desire prospects, we also desire properties, and that we always desire prospects for the properties we think they have." Additionally, he says that "to desire a prospect is to opt for it, or to form the intention of opting for it, among the set of available alternatives; to desire a property is to value it, being disposed, if other things are equal, to desire any prospect that displays the property."

In fact, Pettit claims that by ignoring the assumption of desiderative structure, decision theory is an incomplete account of agents' decision making. This is so because it will not be able to reflect what matters to the agents' decisions. Additionally, since agents are systematically influenced by properties of actions, decision theory is systematically incomplete.

Finally, Broome (1991), which is the proposal I will mainly focus on, argues for an individuation by justifiers. I spell out what this means in the next section.

3.3.1 *Broome's Proposal*

Broome requires two things so that individuation can take place: the first is that the outcomes in question must be different in some sense and the second is that it is rational for the agent to have a preference between these two different outcomes.

Principle of individuation by justifiers (Broome (1991)): outcomes should be distinguished as different if and only if they differ in a way that makes it rational to have a preference between them.

A justifier for difference between two outcomes exists if there is no rational requirement of indifference, and it is therefore permissible to find them different. In other words, a decision maker needs to have a rational reason to distinguish them from one another. The idea is that if it is rational to feel differently about the two outcomes, then individuation is permitted.

Broome addresses the issue of individuation as a solution both to violations of transitivity and violations of the STP. Here I am particularly interested in violations of the STP and therefore I focus on the latter. However, Broome's suggestion is the same in both cases, and the difference is just on how the issue is introduced rather than any actually significant difference. Broome recognises a couple of problems with the individuation strategy, and I will discuss part of them in this section in order to make his proposal clearer to the reader.

Complementarities Between States of the World

The individuation strategy disguises the complementarity between states of the world as something else entirely. In the examples above the reason why one of the actions is preferable to the other is that there is something about that action that cuts across different states of the world. What makes the certain action attractive in the

Allais problem is precisely that it delivers a certain outcome, and this comes from the fact that it produces the same outcome under all states of the world. Therefore, for an agent to realise this, she needs to be looking across states of the world. By individuating outcomes, and preserving the STP, one is saying that actually it is something about each outcome individually that is good, but that is untrue, as one outcome only could not be endowed with the certainty benefit.

Broome recognises this problem when he says the following:

> Allais's fundamental objection to the sure-thing principle is that there may be interactions or complementarities between states of nature. The sure-thing principle requires outcomes to be assessed individually, one state at a time. But if there are interactions between states, they will not show up in such a state-by-state assessment. How did my argument overcome this objection?
>
> Broome (1991), p. 110

It is clear how the nature of the motivation in the examples above goes against the spirit of the STP, and therefore it is suspicious that one can reconcile the two. Broome replies to this objection by appealing to what he calls the dispersion of value amongst states of nature. By this he means that he assumes that the effects of the interaction between states of the world are reflected in the outcomes in the individual states of nature. Broome claims that we are talking about feelings, and that these feelings can be experienced in each state of the world individually, and therefore it is acceptable to make them properties of each individual outcome, instead of properties of the action itself.[12] These feelings are indeed influenced by beliefs on what could have happened in other states, but that is the extent of the interaction between states.

Broome admits that this solution is not without its problems. In particular, he is worried about those cases in which the attractiveness of an action is not due to feelings experienced by the decision maker, but rather by something else. He gives as an example a case discussed by Diamond (1967), where the attractiveness of an action is that it is fair, as opposed to unfair. Diamond initially proposed this example as an objection to utilitarianism as a social choice framework, not in an individual setting. Nonetheless, it is a relevant example for individual choice too. There are two individuals A and B, facing a choice between two alternatives α and β, and two possible and equally probable states of nature θ_1 and θ_2. In action α the utility of A is always 1, and the utility of B is always 0, independently of which state of the world occurs. In action β, on the other hand, the utility values switch under θ_2 (i.e., A gets 0 utility and B gets 1). Again what happens is that in terms of expected utility one should be indifferent between these two alternatives because the utility of one person living is the same in both alternatives.[13] However, as Diamond says and most people

12 By feelings I believe Broome is simply referring to experiential states (such as pain, and anxiety) that can themselves be desirable or undesirable.

13 Assuming the utility function depends only on the number of persons alive.

would agree, alternative β seems preferable because it gives a fair chance to both A and B. Can this concern about fairness be individuated into the outcomes?[14]

Broome believes this example presents two objections to the individuation strategy. The first is that fairness is a property of the process of choice rather than a property of the results of the process. The second objection is that the wrongness of choosing action α is independent of the individual decision maker's feelings, and the unfairness associated with α only exists because of what would have happened in the other state of the world.

Broome replies to the first objection by saying that an outcome is meant to include everything that happens in a given state, both before the state is revealed and after. He replies to the second objection by claiming that he still thinks unfairness should count as a property of an outcome, a property with modal elements. He still thinks this is alright, because he believes that a property of an outcome can have modal elements. He gives the example of a ship that never burns and decays under the sea, but that is nonetheless inflammable.[15]

It seems to me that Broome is in all these three replies saying the same thing: namely that everything should be inserted in the outcomes and that if something matters to the decision maker, then by definition this something must be contained in the outcomes. I disagree with this view, and I argue against it in the following section.

14 The question here is one of whether state-by-state analysis of social states satisfying STP is appropriate. This discussion has been significantly enhanced by work by Fleurbaey (e.g. Fleurbaey (2008)).

15 Although Broome recognises later that Diamond's fairness should not be individuated into outcomes (Broome (2006), pp. 37–39).

3.4 problems with the individuation strategy

As mentioned above, when one is facing situations like the one brought to light by global properties of actions, the two main possible routes are to individuate the consequences more finely and keep the orthodox model of choice, or rather keep the consequences as they are and make the model of decision making finer. My goal in this thesis is not to claim that the latter is always preferred to the former, but rather that sometimes it is. In order to show this, I first argue against the individuation strategy in the case of global properties. It is in part a comparative judgement, because I don't think one can say that individuation is wrong *per se*, but rather that - in some contexts - modelling in a finer fashion is preferable to individuating in one. I will - in the following chapters - argue in favour of the use of finer models of decision making in some contexts, and end up making a case for pluralist models of decision making. For now, however, I present three arguments against the use of the individuation strategy.

3.4.1 *A Technical Concern*

The first argument against the individuation strategy is a technical one. It arises from an existing tension between individuating outcomes and the representation theorem of orthodox decision theory.

One of the assumptions needed for Savage's representation theorem is the constant acts assumption. I mentioned constant acts in the previous chapter, but here I mention them in the form of an assumption, and I discuss their importance to the representation theorem and how individuation puts them in jeopardy. The constant acts assumption is the following (stated informally):

Constant Acts Assumption (CAA): For every consequence in the consequence set there exists a constant act, such that it delivers that same consequence under all states of the world.

Fishburn (1970), Pratt (1974), Seidenfeld and Schervish (1983), and Shafer (1986) argued that the above assumption is needed to obtain Savage's representation theorem. However, this assumption and its use in the representation theorem is problematic in light of the individuation strategy.

Every constant act has - by definition - the same consequence under all states of the world. Additionally, as I will further explain below, it is assume that all these acts have a place in agents' preference rankings. Now suppose that the reason why one prefers the certain act in the Allais case is because the uncertain one is associated with anxiety, and therefore the individuation strategy would alter the consequences associated with the uncertain act by adding 'anxiety' to them. Then, given the CAA, there must exist an act that delivers this altered consequence no matter what the state of the world is. However, the feeling of anxiety comes from the fact that the initial scenario allowed for the possibility of ending up with nothing. How could one experience such anxiety without the uncertainty? In Broome's words, this act "seems causally impossible, and that may make it doubtful that it will have a place in your preferences" (**?**[p.116]broome91).

Gaifman and Liu (2015) argue that the assumption is not needed if one restricts analysis to finitistic problems, and that all that is needed is the existence of two different constant acts. However, even then, infinity is needed for many of Savage's important results. In fact, Buchak (2013) discusses a similar problem. She says that the impossibility of having the same consequence across different situations will limit consistency constraints.

What is more, the CAA is implied by an even stronger assumption that is equally needed for Savage's representation theorem (**?**[p.18]Savage54): the rectangular field assumption.

Rectangular Field Assumption (RFA): If one arbitrarily assigns a consequence from the consequence set to each state of the world, then all the arbitrary acts defined in this way must have a place in the preference ranking.[16]

The RFA - if not as striking in the context of the Allais problem - is even more problematic to individuation than the CAA. This is because the problem discussed above is no longer restricted to causally impossible acts. This problem is also discussed by Broome (**?**[p.116]broome91): he uses Diamond's example of fairness (that he introduces a few pages back) to say that the RFA implies that the consequence that includes 'being treated unfairly' can be placed as a consequence of an act where the person in question had a fair chance.

A second problem that is related to the Allais problem arises when once considers the RFA. If the RFA is indeed assumed, then it follows that agents must have a preference regarding the original Allais acts - but if this is the case, then we are back to our initial problem, even if individuation takes place.

As mentioned throughout, Broome discusses this issue. He agrees that RFA is threatened by the use of individuation, and therefore when trying to save expected utility from a violation of independence, one loses it from another side. However, he dismisses it as a non serious problem because, since RFA cannot be derived from the expected utility formula, even if it is false, expected utility can still be true. Addition-

16 It is easy to see how this implies the CAA - the CAA is simply a special case of this, where the same consequence is assigned to the different states of the world. Additionally, the RFA justifies the claim above that there must be a preference ranking between the constant acts.

ally, one can very likely get the same result without the RFA, even though this seems to be only an intuition, and no proof is provided. However, as I said at the beginning of this chapters, my criticisms were not meant to dethrone the individuation strategy, but rather to cast doubt on its widespread use.

3.4.2 *What Makes a Good Model*

This is the second argument against individuation, and it is a comparative argument regarding what a good model should look like. In what follows I present two reasons in favour of having a finer model and coarser individuation rather than finer individuation and a coarser model.

Reason 1: The Meaning of Utility

Expected utility theory associates a probability function with beliefs, and a utility function with desires. Even though there is disagreement as to what the exact relationship is between the utility function and the agent's desires (some argue that there is a correspondence, others think that utility allows one to represent preferences exhibited by agents, even though it says nothing about underlying mental states), everyone agrees that utility is tracking a preference. Everyone agrees that this is the utility function's job, and not the job of consequences.

However, when using the individuation strategy to solve the Allais problem it seems to me that the difference between the agent who values certainty and the one who does not is pushed into the consequences rather than the utility function. This is because when we look at the Allais problem we realise that - for those consequences and that pattern of choice - no function could solve the problem. As a result, we conclude that we must be representing the consequences wrong. In fact, it would in principle be possible for two agents to have the same utility function, and the only difference between them would be in the way they interpreted what initially seemed like the same consequence: suppose the domain of both of their functions is the power set of the set of consequences, but one of the agents when facing the Allais problem sees the singleton containing only $1m$ with certainty and the singleton containing only $1m$ with no certainty, while the other has a less fine-grained perception of the Allais consequences and doesn't make a distinction between $1m$ with certainty and without certainty, and therefore for both consequences she sees the set that contains both $1m$ with certainty and $1m$ with no certainty. In this case, both agents could still have the same utility function, but different perceptions of the Allais problem that picked up different elements from the domain as the consequences under examination.

Given that this is about personal evaluation judgements, it seems that it should be the role of the utility function to capture these judgements, and not the perception of consequences.

Reason 2: Loss of information, problems of explanation and prediction.

Another reason for wanting to keep the connection between motivation and action explicit is that by incorporating the property of the action influencing the agent in the set of consequences one is making the set of consequences almost imperceptible to the modeller, and therefore information that could otherwise be learnt is lost, and so is the causal connection between the agent and her actions. As a result, the possibility of explanation and future prediction becomes dimmer. If one is at all worried about explanation and prediction, one should be worried about individuation. Buchak also offers support for this reason for modelling motivations explicitly for what they are. She claims that "individuation breaks the connection between what an agent chooses in one decision scenario and what he will or should choose in other, very similar scenarios: this robs decision theory of some of its interpretive and prescriptive power (Buchak (2013), p. 146).

In relation to empirical decision models, Steele (2014) claims that "to be useful, a model must be true enough, and it must yield insights or predictions that justify the trouble of modelling in the first place. Adequate explanation and prediction of human behaviour often requires understanding the beliefs and values that motivate action." In descriptive terms, if one knows what is behind an agent's action, this is not only an explanation but the basis for prediction. I would like to apply this reasoning to the case of normative models too, as I believe explanation and prediction to be desirable even for normative models, as I will go on to explain.

Suppose that one could deal with the Allais problem in the following two ways: (1) individuate the consequences of the certain alternative more finely and include 'certain option' as part of the consequences, or (2) include a subjective function that assesses how much certainty is desired. Both ways would say that it is normatively allowed to prefer the certain option *a* over the uncertain option *b*, but prefer the uncertain option *d* over the also uncertain option *c*. However, (1) would do a worse job at normatively advising the decision maker. When the decision maker is trying to decide what to do, if she is using (1), then she needs to come up with a representation of the consequences she is facing, and then use utilities and probabilities to maximise expected utility in the orthodox way. Nothing is said on how consequences are to be represented - in fact, this is not meant to be a subjective process at all (if it was, then we would already be in the explicit modelling realm).

3.4.3 *The Normative Importance of Hope as Hope*

This is mostly a definitional concern, but one that I believe can ultimately affect behaviour. Is it correct to call everything a consequence? I believe that some things are clearly not consequences and part of what defines them is precisely not being a consequence. There are certain properties that, from the moment they are seen as consequences, something fundamental about them changes. I am going to use the example of hope to illustrate the point I am trying to make here.

Hope

The reason why I use hope as an example now and throughout this thesis is because I think hope is a desire that illustrates well the point I want to make here. Before explaining why, I give a brief example to make more precise what I mean by hope, and to show how it is (or else it can be) a global property of an action. Suppose one needs to choose between alternatives x and y and, excluding hope, they have the same expected value. However, x gets resolved later - for example, the agent only learns the terrible outcome of a referendum or presidential election one year after it happens -, and throughout that time the decision maker gets to keep the hope that everything was going to turn out positively. Keeping the hope alive might have value for some decision makers and also make them have a more positive outlook on life for being more rather than less hopeful. This seems to indicate some pessimism with regards to the final outcome, and it is true that the less sure one is about the favourable outcome occurring, the less one wants uncertainty to be resolved. However, one must still believe that there is a chance that things turn out for the best – otherwise there would be no place for hope. What is more, the space for having hope that x will be resolved favourably might make one a more hopeful person more generally.

While I strongly believe in the arguments above, my main argument, and probably the reason of existence of this thesis, is that making hope a consequence changes its fundamental characteristic. As soon as hope is made a consequence, it becomes the kind of thing that can be computed numerically and added to other very different concerns. We are again human beings that care only about consequences calculating the advantages and disadvantages of their actions. While this must be true to a certain extent - yes, we do decide based on what we believe to be best, all things considered - framing it that way makes us perceive ourselves and others in a darker shade.[17]

3.5 CONCLUSION

I put forward several different concerns with the use of the individuation strategy to save orthodox decision theory from cases such as the Allais problem. In what follows I will generalise the problem presented by Allais to dynamic decisions, and explore alternative ways to deal with such cases.

17 This issue might be familiar if we consider consequentialism as a moral theory. There are numerous objections in the literature to the practice of consequentialising all moral concerns to advance moral consequentialism.

4

KREPS AND PORTEUS: AN EXAMPLE

4.1 INTRODUCTION

In the previous chapter I discussed the use of the individuation strategy to deal with intuitively appealing deviations from what the orthodox theory of choice prescribes agents to do. These intuitive deviations occur mainly when agents care about aspects of decisions that are not traditionally (or, at least, in a colloquial sense) conceived as consequences, but rather as part of the process of decision making. For example, when an agent gets utility out of gambling, it is understandable that she will prefer a compound lottery to a simple one. However, according to the reduction principle present in orthodox theories, the agent should be indifferent between the two when they have the same expected value. One can force these concerns into the consequences via individuation, but I have been arguing against (the widespread) use of this strategy.

An alternative way of solving this problem is to make changes to the overall value of alternatives, and make it dependent on the risk associated with each action, or, in other words, make it dependent on the probabilities in a non-linear fashion. This strategy results in a deviation from expected utility theory, as the independence axiom - which guarantees linearity of the expected utility on probabilities - needs to be relaxed.

A utility of gambling however is a special case. For suppose there is something else (not related to the probabilities) regarding the process of decision making that the agent cares about. In the previous chapter I mentioned the importance of hope as a global property in the context of dynamic decision making, and how unreasonable it would be to try to individuate it into a consequence. This is a concern that is related to the passage of time, and not with the probabilities. In this chapter, I present a framework that aims at modelling time explicitly rather than via individuation.

However, this implies a different departure from orthodoxy. In the orthodox model of decision making the only elements that are explicitly modelled are the probabilities (and states of the world and a probability function, if we consider Savage's model), the consequences, and the utility function. In order to make the value of alternatives depend on time in an explicit fashion (and not through the individuation of consequences), one first needs to add more structure to the model (so that time is modelled explicitly). Consider the following example.[18]

18 In this chapter - for reasons of simplicity and consistency with the model being discussed - I will use vNM as my example of orthodox decision theory.

Example: The Scary Car Park

Anna is on holiday, and parked her car in a car park, but forgot to check the price. The next day she has to decide whether to go **now** to the park and check the price or whether to do so **later**. For simplicity, I assume that the price won't vary between now and later (for example, there's a unique lost ticket fee). Going now might give her peace of mind if the price is reasonable, or she might be upset if the price is really high. Going later will leave her in a state where she is still worried, but at least not sad for knowing that it is indeed expensive.

Anna's story seems quite plausible. She wants to keep enjoying her holiday, and therefore she wants to postpone the resolution of uncertainty in case the news she will get is bad.[19] This kind of concern involves time: the timing of the decision making process is relevant for the decision maker. This can also be represented in the orthodox model. However, since in the orthodox model a modelling choice is made not to represent time explicitly, but only via the consequence set, one cannot make the overall value of acts **explicitly** depend on time-related motivations. This is not something that follows from the axioms in the first instance, it is a choice, one that relies on the assumption that the time when a decision is made should have no impact on decision making.[20]

19 It is assumed throughout that the information about the price of the ticket is irrelevant to any further decisions that Anna has to make - it is not valuable information in this sense.

20 One can of course appeal to individuation, but it should be granted that there is an assumption regarding neutrality of time.

As a result, orthodox theory will make no explicit distinction between Anna's two possibilities. In fact, there is not even a decision to be made: finding out **now** or finding out **later** - since it won't change any consequences - is one and the same thing. Figure 4.1 illustrates this puzzling result. L stands for choosing to go to the park later, and N for going now. NPY stand for not paying yet, and it represents the "payoff" after one period, and then pay little and pay a lot are the payoffs after the second period.

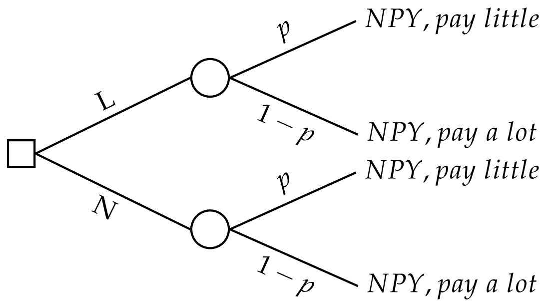

Figure 4.1: vNM Decision Tree

To model the situation described in the example using the orthodox model one would need to individuate the consequences. In this particular case, in order to reflect Anna's preference for late resolution of uncertainty, one should add the consequence of 'deferring knowledge' to the two consequences on the upper part of the tree, the ones associated with finding out later.

In order to avoid individuation in this case one would in the first instance need to add structure to the model in such a way so as to make different time periods distinguishable from one another. Afterwards, one could make the overall value of acts depend on different time periods in different ways. This is precisely what Kreps and Porteus (1978) do in the model I discuss in this chapter.

I find the analysis of this model of great interest. First, it shows us how it is possible to make the value of actions explicitly depend on elements other than the ones explicitly considered by the orthodox model of choice. Second, as I will discuss in the following chapter, this model makes it possible for the value of an action not to depend only on the consequences (and linearly on probabilities) while not violating the independence axiom. This makes a good argument in favour of changing the model rather than applying the individuation strategy - given certain conditions, independence need not be violated.

Finally, this model paved the way for modelling the distinction between risk aversion and the elasticity of intertemporal consumption (Epstein and Zin (1989); Weil (1989)), which is a valuable distinction for many macroeconomic models.[21] Therefore, this model makes it clear how advantageous it can be - in terms of explanation and prediction - to add structure to a model and represent preferences for different properties of actions in an explicit fashion. This discussion is beyond the scope of this thesis, but I refer to it for motivation.

This chapter is organised as follows. The next section describes the motivation behind Kreps and Porteus's model. Section 4.3 describes the general idea of Kreps and

21 The additional flexibility of KP's model has proven useful in applications to macroeconomic models of asset pricing (Epstein and Zin (1989, 1991)), precautionary savings (Weil (1993)), and business cycles (Tallarini (2000)) (see Backus et al. (2004) for a survey of these and related papers).

Porteus's model, including a brief description of the notation to be used throughout the chapter, and an example to illustrate how the model works. Section 4.4 discusses an increase in structure as an alternative to individuation, and finally section 4.5 concludes.

4.2 THE IMPORTANCE OF THE TIMING OF THE RESOLUTION OF UNCERTAINTY

In real life it is usually the case that one has to make a decision that is related to past and future decisions in several different ways. It is very rarely the case that agents have to make decisions that are somehow isolated from all else going on in their lives. This is not necessarily problematic for the orthodox theory, but the more realistic decision situations are, the harder it is for everything to be considered in a model that classifies all things regarding which agents have desires as consequences. When dynamic decisions and time are considered, these worries become even more relevant. For example, some people might like travelling because of the anticipatory feelings of preparing the trip, but not so much because of the trip itself. My parents enjoy knowing that I am going home, and they are happier (or so they tell me) because of knowing this. Sometimes we might not want to lie to a friend - even though we think it would help her - only because we do not want to be the kind of person who lies. And so on.

I am concerned with situations such as the car park case described above, and as I indicated then I will propose that we model such situations using Kreps and Porteus's model. The motivation behind the creation of this model is however different from mine. The particular case that motivated the appearance of Kreps and Porteus's model was the following: suppose one has to make a decision about the present while being uncertain about some factor that would only be resolved in the future. Several authors discussed this issue, for example, Mossin (1969), Dreze and Modigliani (1972), and Spence and Zeckhauser (1972). The most common examples of such a situation concern an agent who has to make a decision about present and future consumption while being uncertain about how much income she will have in the future. The examples used for illustrative purposes are highly simplified. There are two periods, the agent has an initial level of wealth w, she might have a level of present plus future income of y, and there is an uncertain amount x that may or may not be received in the future (for example, there is a lottery to be resolved in the second period that might give the agent more or less money).

This seems quite a possible real life scenario. Suppose for example Anna's pay day is three days away and she has £50 left in her account. She has to figure out how much to spend tomorrow and the day after tomorrow. She is going out for dinner those two nights, and in the first night she is sitting in the restaurant and her most preferred option - disregarding budget constraints for a blissful moment - is the juicy steak. However, this option costs the £50 Anna has available at the moment, and she will be left with no money for next day's dinner. However, before going to dinner, Anna bought lottery $L = (£50, .5; £0, .5)$, but Anna will only find out the result of the lottery the next day, after she has to make a decision regarding tonight's dinner.

If it is inconceivable for Anna to miss the next day's dinner and assuming there is no other way for her to get the money for it, it may very well be the case that, even though she is not particularly risk averse, she ends up valuing this future lottery much less than its expected value, simply because the fact that it only gets solved later rather than sooner makes it unusable for the situation that Anna wanted to use it for. Suppose tomorrow's dinner has a fixed cost of £25. Then, if Anna really thinks it is inconceivable to miss the dinner tomorrow, she will not want to risk it, and she will not choose the steak today. Rather, she will spend only £25, and leave the remaining £25 for tomorrow. Therefore, for the particular situation Anna is in, tomorrow's lottery ends up having very little value. On the other hand, if the lottery was resolved before Anna went out to the first dinner, there would be a chance she would win and be able to actually use the money won to buy the steak.

The difference between Anna's two options is in the timing of resolution of uncertainty. As we said it would be inconceivable for Anna to miss tomorrow's dinner, she will not go for the steak unless she strongly believes (or, taken to the extreme of inconceivability, if she knows) that she will get an extra £50 from the lottery. This is a very likely motivation that can justify a preference for the timing of the resolution of uncertainty.[22] Anna's preferences could also be modelled by simply describing the consequences as, for example, 'steak' and 'no steak,' which wouldn't require any appeal to individuation on the basis of psychological machinery. This is an important difference between the planning cases that motivated the development of KP's model and the kinds of cases I am concerned with.

In contrast to the car park example, in this case Anna's preferences for early resolution of uncertainty are not caused by a feeling provoked by the uncertainty in itself, but rather by an instrumental reason (in the steak example, certainty might allow Anna to have steak). My focus is on cases where the reason for wanting uncertainty to resolve earlier or later is a non-instrumental one, and my suggestion is that we apply KP's model to these particular cases (as an alternative to the individuation strategy, which I argued to be problematic in the non-instrumental cases). This should be straightforward because, even though the motivations for the two cases are different, in both cases two things are perceived as different just because at the point in time in which they are resolved.[23]

22 In this example, Anna has a preference for early resolution of uncertainty. One could easily alter this example so that Anna's preference would be for late resolution: Anna knows she cannot have the steak if she does not win the lottery, but she might accept to go for the steak on the possibility of the 50 pounds being won. In other words, if Anna wanted to be irresponsible, not knowing the result of the lottery would give her the chance of spending more money today. However, let us focus on Anna's responsible side for the purpose of the present example.

23 One can argue here that by classifying all these behaviours as a preference for early/late resolution of uncertainty, one is falling prey to the same problem I attributed to orthodox theory. One might even say that individuation would allow one to add all of these behaviours exactly for what they are. This is true if we are thinking about the adding process, but afterwards individuation leaves us with nothing but a bloated consequence, while KP leaves us with a preference for early or late resolution of uncertainty, and only then a bloated justification.

Kreps and Porteus's model aims precisely at representing such situations. A future lottery is first assessed by its expected value, but then this value can be weighted up or down depending on when the uncertainty associated with the lottery is resolved. This is why it is said to be a recursive model. Suppose there are two stages, the value of a lottery that is resolved in the second stage can be weighted up or down in the present stage. It will be weighted down if the agent has a preference for early resolution of uncertainty, which is the case with the planning cases that motivated the model. However, it can also be weighted up, as in the example mentioned in the introduction where Anna gets utility from not knowing the price of the car park. Then the value of the options at the first stage is constructed via the value carried on from the second stage.

In the next section I explain their model more formally. I then give a very general example to explain the mechanics of the model.

4.3 KP MODEL

Kreps and Porteus (1978) offer an axiomatic approach to dynamic decision problems that allows one to explicitly model preferences that depend on the different time periods that occur throughout the decision. In particular, the aim is to model preferences for the timing of the resolution of uncertainty. They start by adding more structure to the orthodox model (time is modelled explicitly), and they then make the value of actions dependent on this added structure.

One of the examples used in the paper to illustrate this type of motivating behaviour is the following: a coin is being flipped either at $t = 0$ or at $t = 1$. In both cases, the result of the coin flip will determine the prize received in the same way and the decision maker will receive this prize in $t = 1$. In the orthodox model, there is no way to model this distinction in the lottery process explicitly (only via further individuation of the consequences), but with the present model it is possible to model such a distinction, and the decision maker can be explicitly modelled as having a preference between flipping the coin at $t = 0$ or $t = 1$. In order to achieve this, KP first add more structure to the way decision situations are modelled, and second they allow for different times to have a different impact on utility.

In order to better explain how the model works, I will very carefully go through the example above. My aim is to make the model's explanation and the machinery applied to it more suitable for a philosophical audience. Before doing so, I will briefly introduce the model and the notation used by KP.

4.3.1 *General Framework*

In this subsection I will very briefly define the notation used in KP's model, from the set of payoffs up to the utility functions.

Payoffs

The set of possible payoffs is indexed to the time period in which one learns which payoff will be obtained: for each time period $t = 0, 1, ..., T$, there is a set Z_t of possible payoffs with generic element $z_t \in \mathbb{R}$.

Payoff Histories

Set Y_t is the set of payoff histories up to (and not including) time t. A generic element of the set Y_t is a vector of the payoffs obtained up to $t-1$.

The set Y_t is constructed as follows:

$$\begin{cases} Y_1 = Z_0 \\ \text{For } t = 2, ..., T+1, \text{ then } Y_t = Y_{t-1} \times Z_{t-1} \end{cases}$$

Actions

D_T is the set of actions available to the decision maker throughout the entire decision problem. It is defined as the set of Borel probability measures on the set of payoffs Z_T. This simply means that each element of the set of actions is a probability distribution over a number of possible payoffs (in this case, a lottery, which might of course be generate).

States

X_t is the set of nonempty closed subsets of D_T, and its elements are represented as x_t. It is defined recursively on D_T. Informally this means nothing more than the following: since we are talking about a dynamic setting, a decision will not only be associated with a probability distribution on the payoffs for that period, but it will also be associated with the place where the agent will be taken via that decision. Each element of X_t determines which decisions from D_T will be available to the decision maker in the next period. The reason why X_t is said to be defined recursively is that it is the decisions that can be made in the next period that define X_t.

Utility Functions

There are two utility functions in KP's model:

Payoff Utility: $U(z_0, ..., z_T) : Z_T \to \mathbb{R}$

Expected Moderated Utility: $u(z_0, \gamma) : \gamma \to \gamma^x$, $\gamma_t = EU_{t+1}, x \in \mathbb{R}_{>0}$

The first utility function - that I called general utility - is the same utility we are used to dealing with in the orthodox model, and it is applied statically. It is simply a mapping from the payoffs (in this case obtained in different periods) to the reals. The difference with regards to orthodoxy is that this function does not fully characterise the decision maker's behaviour, a second function is needed. The second function - that I called expected moderated utility (EMU) - is responsible for representing agents' sensitivity to time periods. It depends on z_0 because this is the first term of the recursion, but in subsequent periods it depends on γ, where γ is the result of the maximisation of expected utility of the following period.

4.3.2 *Example: Flipping Coins at Different Times*

This example is used by KP (p. 193) when trying to explain their model in terms of temporal lotteries. I here use it without that intent, but rather as a first - very simple - example of how their model is meant to work and how it differs from the standard model. I simply go through the example while trying to explain how KP's utility functions work in these cases and why they allow one to have preferences for the timing of the resolution of uncertainty. This example is particularly interesting because it is clear that under vNM the agent would be indifferent between the two options available to her, at least as originally described; indeed, and as it happened with the examples in the introduction, there would be no way to distinguish them under a vNM lottery description of the problem.

In this example there are two periods of time, and therefore $T = 1$: in the first period $t = 0$ and in the second period $t = 1$. In $t = 0$ the payoff is certain and it is 5 monetary units, and in the second period the payoff depends on the outcome of a flip of a fair coin - it can be either 0 (for example, if the coin comes up heads) or 10 (if the coin comes up tails). Therefore, using KP's notation, $Z_1 = (5, 0)$ or $Z_1 = (5, 10)$, where the first element of the pair stands for z_0 and the second stands for z_1. In this example, there is only one choice for the agent to make - $D_T = \{d_0(a), d_0(b)\}$, and the choice is between $d_0(a)$ and $d_0(b)$ - which occurs in period $t = 0$, and in practical terms it is a choice between the coin being flipped at $t = 0$ (if she chooses $d_0(b)$) or $t = 1$ (if she chooses $d_0(a)$). This coin flip only influences the prizes relative to the second period (as the payoff is certain in the first period), and therefore the result being known later rather than sooner does not change the payoff structure. The example can be summarised using the following decision tree[24]:

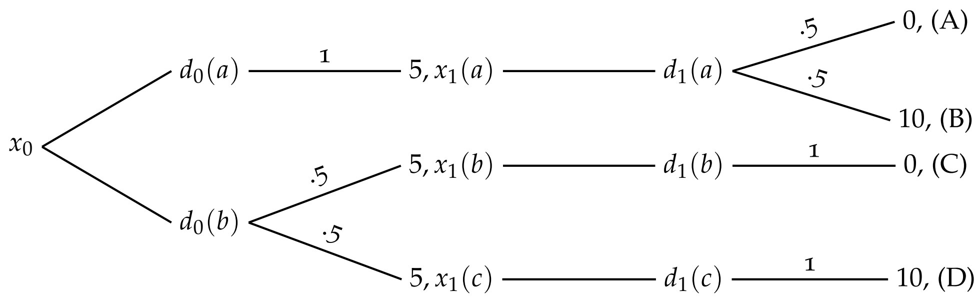

Figure 4.2: Flipping Coins at Different Times

The decision maker needs to decide between the two actions available in period $t = 0$: between $d_0(a)$ and $d_0(b)$. That is the decision that needs to be made. In $d_0(a)$, the coin is flipped at $t = 1$ while in $d_0(b)$ the coin is flipped at $t = 0$, and that is the

24 Capital letters $(A), (B), (C), (D)$ are represented in order to name the branches so that below the reader can better associate the computations made with the different branches.

only difference between the two options. In order to make the example more tractable, I will specify the utilities, and I will use the same functional forms as KP do.

The utility function of payoffs (general utility) is given by

$U(z_0, z_1) = (z_0 + z_1)^{1/2}$

and utility's sensitivity to time periods (EMU) is given by

$u_0(z_0, \gamma) = \gamma^2$, where γ is the next period's ($t = 1$) expected utility, computed using the general utility function. As explained before, z_0 is only included in the domain to guarantee its completeness with regards to the first period.

These two functions represent the decision maker's choice behaviour. The first function simply states how the payoffs in each period contribute to utility, and the second defines the interaction between time periods and utilities. This is because γ stands for the expected utility of the action with the highest expected utility at each period. Let me explain this step by step.

The problem is solved by backwards induction, and therefore we should start by looking at the end of the tree.[25] For each terminal node one uses the utility function on the payoffs $U(z_0, z_1) = (z_0 + z_1)^{1/2}$ to compute the utility that one would get were one to end up in that terminal node (the utility of the vector of consumption as a whole). For the upper most branch (A) the payoffs are 5 in the first period and 0 in the second

25 It follows from KP's temporal consistency axiom and the subsequent recursive functional form of KP's model that the model is solvable by backwards induction. Additionally, backwards induction is used in KP's paper to solve dynamic choice problems according to their model Kreps and Porteus (1978, pp. 192-193).

period. Therefore, the utility is given by $U(5,0) = (5+0)^{1/2} \approx 2.236$, and the same computation is done for the remaining terminal nodes as follows.

Branch (B): $U(5,10) \approx 3.873$

Branch (C): $U(5,0) \approx 2.236$

Branch (D): $U(5,10) \approx 3.873$

Focusing again on the upper part of the tree, one then computes the expected utility of $d_1(a)$, which is simply .5 multiplied by the utility of the payoffs of the upper most branch (A) plus .5 multiplied by the utility of the payoffs of the branch just below (B). This results in 3.054 and this value is our γ. According to the decision maker's behaviour represented by u_t, this γ should be squared ($3.054^2 = 9.330$) and carried over to the proceeding (here, first) period.[26] Now the same expected utility computation should be applied at $d_0(a)$, where the probability of each branch occurring is multiplied by the current value associated with that branch (i.e., the squared γ). However, in this case, given that there is only one branch that obtains with probability 1, this value becomes just the squared γ, i.e., the 9.330 obtained above. This is then the expected value of $d_0(a)$, which to avoid confusion I will simply refer to as the EMU of $d_0(a)$.

The same computation is done for the lower part of the tree to compute the EMU of $d_0(b)$. Starting with branch (C), the utility of this combination of payoffs is given by $U(5,0) = (5+0)^{1/2} \approx 2.236$. Since at this stage this is the only possible prize,

26 For decision situations with more than two periods, this process would be repeated for the remaining periods, the value carried over being the one obtained from the periods ahead.

this value is multiplied by 1 to obtain the expected utility of $d_1(b)$. This value (γ) is then transported to the additional resting point, $x_1(b)$, where the function u is applied: $u_0(z_0, \gamma) = \gamma^2 = 2.236^2 = 5$. The same is done for branch (D), but here γ is 3.873 and therefore the value at $x_1(c)$ is given by $3.873^2 = 15$. Now to compute the EMU of $d_0(b)$ one simply multiplies these two values for the probabilities of them obtaining, which is .5 for each. Therefore, the EMU of $d_0(b)$ is given by $.5 \cdot 5 + .5 \cdot 15 = 10$.

The decision maker should then choose the action with the highest EMU. For this particular example, $d_0(b)$ would have a higher value ($10 > 9.330$), which means that this agent has a preference for early resolution of uncertainty. In fact, it is the shape of the EMU that is going to determine the agent's attitude with regards to the timing of the resolution of uncertainty. What is the role of γ in the results obtained? First notice that, given the utility function assigned to the decision maker in this example: $U(z_0, z_1) = (z_0 + z_1)^{(1/2)}$, the payoffs of the two periods are weighted equally according to this utility function. Additionally, the two possible events (heads or tails) are equiprobable. Therefore, it is possible to reason in terms of averages, which makes this example particularly useful in studying the role of γ and its exponent.

The difference between the EMU of the two options is that one is the square of the average (upper branch) while the other is the average of the squares (lower branch). Which of these values will be higher? This can be answered generally, for the answer will hold across different values for the two prizes in question, and one can easily obtain the range of exponents of γ for which an agent exhibits a preference for early or late resolution of uncertainty, and for which value of this exponent will an agent be neutral (as per the orthodox model) to the timing of the resolution of uncertainty. For any x and y, let us compare the square of an average (1) to the average of the squares (2):

$$\left[\frac{x+y}{2}\right]^2 = \frac{(x+y)^2}{4} = \frac{x^2 + 2 \cdot x \cdot y + y^2}{4} \tag{1}$$

$$\frac{x^2 + y^2}{2} = \frac{2 \cdot x^2 + 2 \cdot y^2}{4} \tag{2}$$

In order to assess which of these expressions will yield a higher value, we can simply do $(2) - (1)$ to get the following:

$$\frac{(x^2 + y^2) - 2 \cdot x \cdot y}{4} \tag{3}$$

Solving for the inequality to 0 the following result obtains:

$$\begin{cases} x^2 + y^2 > 2 \cdot x \cdot y \text{ for } x < y \text{ and } x > y \\ 2 \cdot x \cdot y = x^2 + y^2 \text{ for } x = y \end{cases}$$

Therefore, whenever there is a difference in prizes (and if there is not, there is no uncertainty, so this is not a case that concerns us), squaring after the average will result in a smaller value than squaring before the average is computed. This is the same as saying that the agent prefers the lottery in which the uncertainty resolves earlier. This is true more generally for any $x > 1$ in γ^x. The converse holds for $x < 1$, for note that $[\frac{x+y}{2}]^{1/2} = [\frac{2}{x+y}]^2$. The orthodox neutrality obtains for the special case when $x = 1$.

In fact, going back to the original example, suppose the agent had a preference for late resolution of uncertainty, and her EMU was given by $u_0(z_0, \gamma) = \gamma^{\frac{1}{2}}$. Then $EMU(d_0(a)) \approx 1.748$, and $EMU(d_0(b)) \approx 1.732$, and she would prefer option $d_0(a)$, the option where the uncertainty resolves later. Now suppose the agent is indifferent between early and late uncertainty resolution, and that her EMU was given by $u_0(z_0, \gamma) = \gamma^1$. Then $EMU(d_0(a)) = EMU(d_0(b)) \approx 3.054$ and she would be indifferent between options $d_0(a)$ and $d_0(b)$.

In summary, for the simple generic shape γ^x an agent will have a preference for early resolution of uncertainty if $x > 1$ and a preference for late resolution if $x < 1$ and she will be neutral with regards to this timing if $x = 1$. This latter case would reduce KP model to the orthodox vNM model.

4.4 ENRICHING THE STRUCTURE OF THE MODEL

In this section I reflect on how the KP model adds structure to the orthodox model of choice. They do so in a different way than non-expected utility models because their model aims at making the overall value of alternatives depend explicitly on time; however, unlike probabilities, this is not one of the elements represented in the orthodox model.

In orthodox decision theory the number of parameters subjectively determined by the decision maker are the utility function (in both vNM and in Savage) and the probability function that assigns probabilities to the states of the world (in Savage only). Preferences over alternatives are assumed to be fully determined by decision makers' desires on consequences (in the form of the utility function), and the probabilities of these consequences obtaining (either in the form of a subjective probability function, or in the form of exogenous probabilities). The overall value of an alternative is given by the probability weighted average of the utilities of the different possible consequences associated with that alternative.

There are several examples in the literature of departures from orthodox theory, where a different subjective parameter is added. Non-expected utility theories give up linearity of the value function in the probabilities. These models offer an alternative to expected utility theory that allows for phenomena such as a preference for certainty (as observed in the Allais case) to be modelled explicitly. They give up the independence axiom/the STP, and a different set of axioms is used.

A good example of such a model is Buchak's Risk-Weighted Expected Utility (Buchak (2013), REU). Her model keeps both utilities and probabilities as subjective parameters, and it adds a third one: how much an individual cares about consequences in worse states of the world in comparison to consequences in better states of the world. This is represented by a risk function. This model violates the STP - giving up the linearity of utility in the probabilities is its defining characteristic, and this is in fact what it sets out to do -, but another axiom is added in its place, and a new representation theorem derived.[27]

Analogously, KP add a new subjective parameter: how early or late resolution of uncertainty is assessed.[28]. However, while in Buchak's case the aim is to make the overall value of alternatives depend on probabilities in a non-orthodox way, in the KP case, they want to make the overall value of alternatives depend on time in a non-orthodox way (and as a result, it is only applicable to cases of dynamic decision making). In order to make utility depend on time in this way, KP first need to add even further structure to the model: they need to make time a primitive of the model (similarly to how Savage's states of the world were needed to make probabilities a subjective parameter). Then they can - in a similar fashion to what happens in Buchak's model - have a function that weights expected utility according to the time period in which it is obtained. In what follows I describe these two steps in more detail.

The way time is represented explicitly in KP is by allowing the branches that connect decisions at different points in time to matter for the decision making process. KP's model allows one to apply a utility function throughout the branches of a tree, and as a consequence the branches that represent the passage of time influence the

27 In Buchak's model, the added axiom is the Comonotonic Sure Thing Principle, logically weaker than the STP and independence.

28 Note that in their model probabilities are not subjective, so overall their model has two subjective parameters

process of decision making. To explain how this works I would like to make use of the car park example described in the introduction, and now solve it using KP's model. Above I gave a representation of the example using the orthodox model. However, if we represented the same situation using KP, we would obtain the following decision tree:

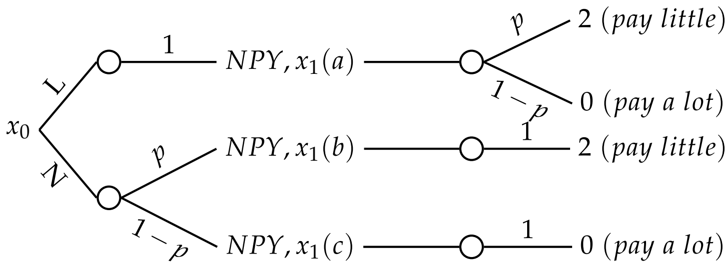

Figure 4.3: KP Model the Car Park Decision

Let us compare this representation with figure 4.1 above. In the KP representation one has one additional stop - represented by the set X_t. It is precisely this set X_t that allows one to apply different treatments to γ (the expected value of the next round) before moving on to computing the rest of the value of the decision in question. There is a stage in the decision problem - after a decision in period t is made but before one for $t+1$ needs to be made - that allows for things to be weighted differently at that point. A good analogy to consider could be the following: in the KP model there is an extra bus stop where passengers are allowed to get in or get out of the bus. Just because the bus stops there, there is more that can change for that bus in between the two stops that come before and after that additional stop.[29]

29 Another way to think about the additional structure in KP when compared to vNM is to consider the difference in terms of the roles of the branches along the decision tree. In vNM having two branches

Up to this point, the added structure does not make KP's model inconsistent with the orthodox axioms. Adding more structure to a decision problem does not of course force on to abandon the vNM theory or indeed to permit preferences that violate one of its axioms.

The second part of the change in structure - the part analogous to what happens in non-expected utility models - is what forces a deviation from the orthodox axioms. Making the overall value of actions depend on something that is not modelled as a consequence nor a probability will result in a deviation from orthodoxy. In fact, for values of the exponent of γ other than 1, KP's model will violate one of the orthodox axioms. I will discuss which axioms are and aren't violated in the next chapter.

4.5 CONCLUSION

In this chapter I introduced KP's model as an alternative to individuation. I had argued against individuation in the previous chapter, but here I identified an additional problem with the orthodox model that has to do with the fact that, due to its lack of structure, it does not allow for any elements other than consequences and probabilities to affect the overall value of an action. Therefore, if one does not want to individuate

that lead to the same consequence would be no different than having one branch only leading to that consequence with the added probabilities of the initial two branches. However, in KP, the number of branches does make a difference (as long as γ is raised to a power other than 1). It is the difference between weighing something before or after a multiplication takes place.

certain given motivations and make them consequences, then one should add more structure to the model. This is what KP do: they make time a primitive of the model. Once time is modelled explicitly, the possibility of making the choice worthiness of acts depend on time opens up as an alternative to individuation. In the next chapter, I discuss how adding time as a primitive allows KP's model to preserve dynamic consistency.

5

DYNAMIC CONSISTENCY AND INDEPENDENCE

5.1 INTRODUCTION

One of the goals of this thesis is to offer alternatives to the individuation strategy described in Chapter 3. In the previous chapter I introduced Kreps and Porteus's model as an alternative. As it deviates from the orthodoxy, their model violates at least one of the axioms of the standard theory. In this chapter I discuss which axiom or axioms are indeed violated by their model, and I assess how problematic such a violation is. This is done with the purpose of evaluating a model that was presented as an alternative to individuation, and therefore it is a relevant step in reaching the goal stated above.

In particular, this chapter aims at answering the following questions: (1) which of the orthodox axioms are violated by KP's model?, and (2) how problematic is such a violation? Question (1) needs to be answered first and it is asked mainly so that question (2) can be answered. The reason why question (2) is important is because the viability of KP's model as an alternative to individuation depends on its answer.

To answer the first question I start by explaining the orthodox conditions of rationality imposed on dynamic decision making. In chapter 2 I described the orthodox axioms, but in the context of dynamic decision making it is common to use a different set of axioms. These axioms are taken to be a more enlightening way of discussing dynamic decisions, but it can be shown that the static axioms are implied by a combination of the dynamic axioms.

I start by discussing the independence axiom in the context of dynamic decisions, which takes me naturally to the remaining conditions imposed on dynamic choice that together imply independence. These conditions are plan reduction, separability, and dynamic consistency. One of the main worries with models of dynamic decision making that deviate from the standard axioms is that they might allow agents to be dynamically inconsistent. Therefore, in order to preserve the viability of KP's model as an alternative to individuation, an important goal of this chapter is to show that KP do not, in fact, violate dynamic consistency.

An agent behaves in a dynamically inconsistent way if she plans at a stage x of the decision process to choose action A at the next stage, but when facing the choice one stage later does not choose A, even though the option is still available. This is problematic for obvious reasons, and usually referred to as irrational because such an agent can make a sequence of choices that will result in a sure loss relative to the starting point.

I show that this does not happen in KP's model because their model is solved using backwards induction. I then show that separability is also preserved, but that plan reduction is violated by KP's model. I go on to discuss the importance of this violation, and conclude that plan reduction is actually not a good normative restriction on choice, especially in the context of dynamic decision making, when time is passing by, and it becomes natural for agents to have preferences regarding how decisions are made.

This chapter is organised as follows. In the next section I introduce McClennen's (McClennen (1990)) notation because I will mainly follow his definitions of the dynamic choice axioms. Section 5.3 discusses the conditions of rationality that are thought appropriate for dynamic decision making. Section 5.4 discusses the importance of dynamic consistency and it then shows that it is not violated by KP's model. The following sections discuss the violation of plan reduction, and how adding time as an explicit component of the model is what allows for dynamic consistency to be preserved. Section 5.7 discusses the status of plan reduction, and it argues that it is reasonable to give it up in certain contexts, which hints at the idea that different contexts may require different axioms. Finally, I conclude in the last section.

5.2 NOTATION

A dynamic decision problem is one where agents have to make a sequence of choices over time. Their initial choices, and sometimes also uncertain events, determine the posterior decisions that have to be made. Dynamic decision problems are usually represented by decision trees that can contain both choice nodes - a point in the decision making process where the agent has to make a decision -, and chance nodes - a point where some uncertainty gets resolved. In this section I introduce the notation used by McClennen (1990). I will use his definitions of the conditions that are usually judged appropriate to impose on dynamic choice, and therefore it is important that I introduce his terminology before moving on.

First, a bounded (in the sense that it does not go on indefinitely) sequential decision problem is represented by a decision tree, and denoted by T. Choice and chance nodes are represented by n_i and terminal nodes are named outcomes and represented by o_i. $O(T)$ is the set of such outcomes for a given decision problem T. McClennen then introduces the notion of a plan, which he defined in the following way: "A plan for a decision problem T, as I shall use the term here, consists of a complete specification of what choice is to be made at each choice node in T reachable as a result of the implementation of earlier portions of the plan" (McClennen (1990), p. 101). The set of plans available to the agent in a given decision problem T is denoted by $S(T)$ or simply S if denoting it this way causes no confusion. Elements of $S(T)$ are denoted by s, r and so on.

Trees - and subsequently plans - can get truncated. When a tree gets truncated at a given point, the truncated part of the tree is the part of the tree from the node where it was truncated and the remaining part of the tree reachable from that node. Given a tree T, the tree obtained by truncating T at node n_i is denoted by $T(n_i)$. Equivalently, $S(n_i)$ represents the set of truncated plans available at node n_i.

With regards to comparing the goodness of different plans, further terminology needs to be introduced. $D(S)$ is the set of acceptable plans, the subset of the set of plans S consisting of those plans judged by the agent to be acceptable (by her own criteria). Correspondingly, $D(S(n_i))$ is the set of acceptable truncated plans. Additionally, $D(S)(n_i)$ is the set of all plans in $D(S)$ that the agent can still implement from the node n_i.

Finally, a truncated tree considered as a *de novo* tree (i.e., considered as if what happened before the agent reached the node in question never happened) is represented by $T(n_i)^d$. This tree has no history. Correspondingly, $S(n_i)^d$ is the set of available plans in $T(n_i)^d$, and $D(S(n_i)^d)$ is the set of acceptable plans for $T(n_i)^d$. Note that the difference between $T(n_i)^d$ and $T(n_i)$ is that the former has no history and the latter has a history and agents were brought to that node through a sequence of choice and/or chance nodes.

When there are chance nodes involved in a plan, it is necessary to introduce the concept of prospects. For any tree T and associated plan s in S, a prospect is the gamble associated with plan s (and a gamble is a probability distribution over the possible prizes associated with s), and it is represented by g_s. Correspondingly, G_S is a set of prospects. Again, this carries over to the truncated case: g_s becomes $g_{s(n_i)}$ and G_S becomes $G_{S(n_i)}$. Finally, $D(G_S)$ is the set of acceptable prospects: the subset of prospects consisting of those prospects judged by the decision maker to be acceptable.

5.3 DYNAMIC POSTULATES OF RATIONALITY

As discussed in chapter 3, one problem that arises when trying to solve the Allais problem by making expected utility depend non-linearly on probabilities (instead of applying the individuation strategy) is that the independence axiom is violated. This is problematic because, arguably, there are normative reasons to want to keep the independence axiom as a requirement of rationality. Therefore, a natural question to ask after offering the Kreps and Porteus model as an alternative to individuation in the hope case, and after making utility dependent on time, is the following: does their model give up independence?

Before attempting to answer this question, it is important to clarify what exactly this independence axiom refers to. This is especially important because we are now discussing this axiom in the context of dynamic decision making, and inevitably the definition used in chapter 2 needs to be qualified in a couple of senses. An added difficulty here is that the independence axiom is described in many different ways by different authors, and this is true especially in the context of dynamic decision making. Since I will be following McClennen's terminology in this chapter, I will use his definition, but I will then comment on what his definition entails.

The definition McClennen uses in the context of decision making is the following:

McClennen's independence: Let g_1, g_2, and g_3 be any three gambles, let $g_{13} = [g_1, p; g_3, 1-p]$ be a gamble over g_1 and g_3 such that one stands to confront the gamble g_1 with probability p and the gamble g_3 with probability $1-p$, and let $g_{23} = [g_2, p; g_3, 1-p]$ be similarly defined. Then g_1 is in $D(g_1, g_2)$ iff for $0 < p \leq 1$, g_{13} is in $D(g_{13}, g_{23})$.

He then adds the note that, in the context of dynamic decisions, gambles are associated with plans (they are prospects), and not with single actions as in the static case.

The KP model respects independence within single time periods (i.e., when applied to the decision of what to do in a single stage) - as in each period one computes the expected value of each action - but it does indeed violate independence as defined above when one is considering the whole decision tree (i.e., it violates independence when it is applied to prospects). But why? And how relevant is this? This is what I want to answer in this chapter.

In order to do so, I will start by discussing the orthodox conditions of rationality that are imposed on dynamic decision making, and a result that relates these conditions with the independence axiom.[30] These conditions are the following:

- Dynamic Consistency (DC)
- Separability (SEP)
- Plan Reduction (PR)

These three conditions imply independence as defined above, therefore, at least one of them must be violated by the KP model. In the remainder of this chapter I explain what these three conditions mean, and I work on finding which of these conditions are violated by KP, and I will then discuss the importance of such a violation.

30 Slightly different results have been shown by different authors (Hammond (988a, 1988); Karni and Schmeidler (1991); Volij (1994)), but here I choose to use McClennen's (McClennen (1990)) terminology and result.

5.4 DYNAMIC CONSISTENCY

A major worry with unorthodox models of dynamic decision making is that they violate dynamic consistency. Therefore, the first step is to show that dynamic consistency is indeed preserved by KP's model.

5.4.1 *Definition*

Dynamic consistency establishes a relationship between planned behaviour and actual behaviour. It requires that actual behaviour does not deviate from planned behaviour. There is no situational decision-making - the decision is perceived in the same way no matter where the agent is at. I will again state McClennen's formulation of the principle here. Note that $D(X)$ is defined as above.

Dynamic Consistency (DC): For any choice point n_i in a decision tree T, if $D(S)(n_i)$ is not empty and $s(n_i)$ is in $D(S(n_i))$, then $s(n_i)$ is in $D(S)(n_i)$; and if $s(n_i)$ is in $D(S)(n_i)$, then $s(n_i)$ is in $D(S(n_i))$.

In natural language, the above says that, for any choice node in a decision tree, a plan is in the set of acceptable truncated plans if and only if it is also in the set of all plans in the initial set of acceptable plans that the agent can still implement from the choice node where she is currently at (if the latter set is not empty).

Informally, this requirement says that if a plan is in the set of acceptable plans at the beginning of the decision process, then it will be in the set of acceptable plans at every stage of the process (provided it is still available). In terms of decision trees, this means that upon reaching a new decision node, an agent will stick to her original plan, even if reconsidering her choices is a possibility. This can be well understood if we focus on Figure 5.1.

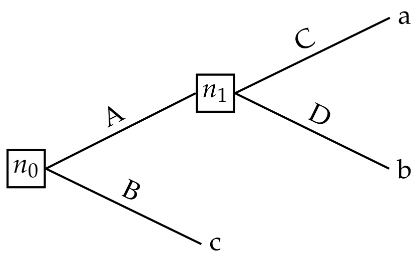

Figure 5.1: Dynamic Consistency

One would violate DC in this situation if one plans at n_0 to head for b and at n_1 to head for a.

5.4.2 *The Harms of Dynamic Inconsistency*

Why is a violation of dynamic consistency problematic? In general, it does not seem like a good idea to plan to do something and yet not carry through with this plan (I am here excluding the existence of obvious reasons not to do so, such as preference change). Strotz (1955) justifies this by saying that a dynamically inconsistent agent will find "himself continuously repudiating his past plans" and "may learn to distrust his future behaviour..." (p. 173). He also mentions the *mañana effect*, which is seen as problematic as it entails that unpleasant things will keep on being postponed. In fact, people do sometimes pay a positive amount to restrict their available options in the future in order to avoid dynamic inconsistency.

Additionally, dynamic inconsistency could cause an agent to be exploited in a manner akin to how an agent with intransitive preferences can be money-pumped. If someone's preferences are dynamically inconsistent, then she can end up willingly making a book against herself through a sequence of choices (for a more in depth discussion see Schick (1986), Green (1987), and Cubitt and Sugden (2001)). The money-pump argument runs in the following way: focusing on 5.1, suppose a decision maker plans at n_0 to head for b (to choose the sequence of choices A and D) and at n_1 to head for a (to choose C). Assume the following three things: (1) that at $t = 0$ she only has access to the plan consisting of A and C, (2) that she is willing to pay a positive amount to trade the plan she currently has for one she prefers, and (3) that she does not anticipate nor does she factor in her later preference changes. Then the following problematic case might occur: at $t = 0$, she would be willing to pay a fee so that she could trade the plan that allows her to go to a at $t = 1$ for the plan that allows her to go to b at $t = 1$ (because this is her preferred plan at $t = 0$). However, when she gets to n_1 she will want to head to a, and therefore she would again be willing to pay to return to the original plan, given the assumption that at $t = 1$ she prefers $a + 2 \times fee$ to b.[31] Therefore, the agent would pay this fee twice just to get the plan she would anyhow have, had she not made any trades nor paid anything. This could be taken much further if we considered more time periods.

5.4.3 *The KP Model and Dynamic Consistency*

We defined DC above by saying that, when considering a decision tree, an agent will stick to the original plan upon reaching every new decision node. Is this going to be necessarily the case in KP? The answer is yes. KP's model doesn't distinguish between a decision made before any actual decisions need to be made, and one that is made in the middle of the decision-making process. This is because utility is defined recursively and the KP solution is obtained by backwards induction (i.e., by starting at the end nodes, and deciding what to do then, and then moving on to the previous stage keeping in mind what will be the result of the last stage).

Backwards induction in one's approach to dynamic choice is sufficient for dynamic consistency. To see this consider a two-period decision situation: one decision to be made today, and one to be made tomorrow. If the agent's choice tomorrow does

31 A similar example could be constructed without the need for this assumption, and in any case this fee can be an infinitesimal small value.

not affect her choice today (which is the case in KP's model, as each periods' expected utility does not depend on the other periods), then she starts by maximising tomorrow's expected value. And tomorrow, she also maximises tomorrow's expected value. Thus, what the agent chooses tomorrow is what she would choose to choose tomorrow if she was forced to make tomorrow's choice today. In other words, the agents sticks to her original plan.

5.5 SEPARABILITY

In this context separability establishes a condition on how the history of a decision can or cannot affect the present-moment decision. What it says is that the history can be ignored, and that, when facing a particular choice node, one should ignore what happened before that choice node and just focus on the future. The rest of the tree can be truncated, as it is done below in Figure 5.2.

Separability (SEP): For any tree T and any node n_i within T, let $T(n_i)^d$ be a separate tree that begins at a node that corresponds to n_i but otherwise coincides with $T(n_i)$, and let $S(n_i)^d$ be the set of plans available in $T(n_i)^d$ that correspond one to one with the set of truncated plans $S(n_i)$ available in $T(n_i)$. Then $s(n_i)$ is in $D(S(n_i))$ iff $s(n_i)^d$ is in $D(S(n_i)^d)$.

McClennen also describes this condition informally, by saying that a plan considered as a continuation is choice worthy if and only if it is choice worthy considered as if *de novo*. Where *de novo* simply means that the tree is truncated and one chooses from there.

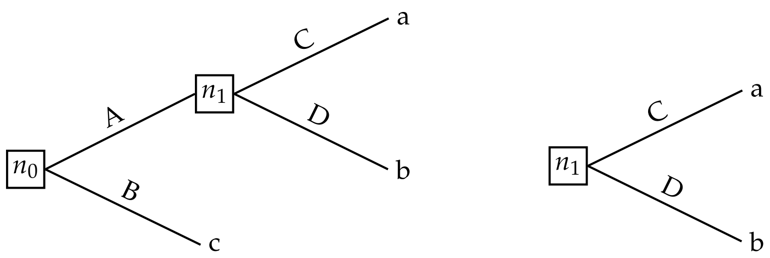

Figure 5.2: Separability

In Figure 5.2 if one would plan -from the position of n_1 - to head for b when facing the left-hand side tree, but plans to head for a when facing the truncated tree on the right-hand side, then one would violate SEP.

5.5.1 *KP Model and SEP*

For very similar reasons to DC, SEP also holds in KP. SEP says that truncating the tree should make no difference for the decisions that are yet to come. Again, since the solution in KP is found by backwards induction, the decision made from where the tree is truncated will obviously be the same, because when the agent reached a decision for that particular node, she hadn't even considered the nodes up to the point, but only the nodes after that point.

5.6 PLAN REDUCTION

Plan reduction can be divided into two components. I here explain the two components separately, and I then define PR.

5.6.1 *Simple Reduction*

I will start by giving McClennen's definition of this condition, and then I will move on to try and make it clearer by using informal language. This is a very weak requirement, and it is only together with the following one - NEC - that it makes more sense in terms of practical implications. I will argue below that these two requirements together (SR and NEC) amount to the well known reduction of compound lotteries requirement, but applied to a dynamic setting. I believe this connection will make the role of the following two requirements clearer, but for now I simply describe them individually. I start by stating the definition of SR.

Simple Reduction (SR): Let T be any decision tree with associated set of plans S such that each plan s in S requires for its implementation a single choice "up front" by the agent, and let G_S be the set of prospects associated with such plans. Then for any plan s in S and associated prospect g_s in G_S, s is in $D(s)$ iff g_s is in $D(G_S)$.

First of all, it is important to note that this condition only applies to very simple decisions that, although dynamic, only ask of the decision maker one choice. The dynamic component comes from the fact that the uncertainty associated with the one-off decision is sequentially resolved. This is what is meant by the claim that only a single choice is required, and this condition imposes a restriction on these types of decisions.

To better understand this condition, let us focus on an example: the decision represented by the tree depicted in Figure 5.3. The agent makes one initial decision between A and B - as indicated by the decision node at the left-hand side of the tree -, and the upper part of the tree represents uncertainty being sequentially resolved.

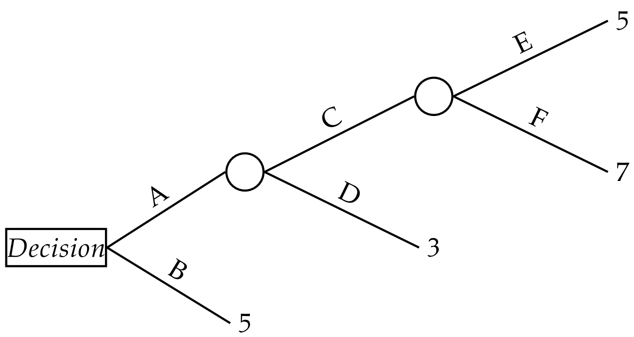

Figure 5.3: One-Decision Tree

This means that the agent can choose between a certain outcome (B, which guarantees the agent a payoff of 5) or a lottery, which gives the agent 3 with probability $1 - p$, 7 with probability $p \cdot q$, and finally 5 with probability $p \cdot (1 - q)$. Using McClennen's terminology, the agent has to decide between two plans: plan B that has an associated prospect that amounts to the degenerate lottery that gives the payoff 5 with probability 1, and plan A that has an associated prospect that amounts to the lottery described above.

Which restriction does SR place on this choice? It requires that if the lottery associated with A is preferred to the lottery associated with B, then A must also be preferred to B. Coming back to McClennen's terminology, there are two plans, A and B (since a single choice is made, a single action defines a plan), and there are two prospects, the lotteries associated with both A and B, and SR claims that $A(B)$ is in the set of acceptable plans if and only if $G_A(G_B)$ is in the set of acceptable prospects.[32]

5.6.2 *Normal-form/extensive-form coincidence*

The NEC condition creates a bridge between dynamic decisions and static ones. To discuss this condition, one needs to be able to reduce a dynamic-choice representation (a decision tree) to a static normal-form representation (a decision matrix) that rep-

32 Note that this is a weak requirement because the argument against reduction conditions is usually that there might be some utility in making several choices rather than one, or vice-versa, but here only one decision is made, so that argument does not play a role in arguing against this condition.

resents all existing strategies in the dynamic problem (i.e., all possible combinations of actions). Then NEC basically says that these two representations are equivalent in terms of the decisions prescribed. I will again adopt McClennen's terminology in stating the condition, and I will then try to make it clearer by explaining it informally. Let s^n be the normal-form version of s (a single plan), T^n be the normal-form version of T, and S^n be the set of plans in T^n (or, equivalently, the normal-form version of S).

Normal-form/extensive-form coincidence (NEC): Let T be any decision tree with associated set of plans S, and let T^n be the decision problem that results by converting each s in S into its normal form, so that each s in S is mapped into s^n in S^n. Then for any plan s in S, s is in $D(S)$ iff s^n is in $D(S^n)$.

To explain what this condition requires, it is useful to compare the same decision problem represented both as a decision tree and in the form of a decision matrix, as I have done in Figure 5.4. For NEC to be respected, both representations should prescribe the same choice of plan (although nothing is said about which decision should be made). The idea behind this condition is that the representation of a dynamic problem using a decision tree does not add any relevant information to the decision-making process when compared to a normal-form representation (in other words, the introduction of choice nodes does not change the preference ordering over strategies).

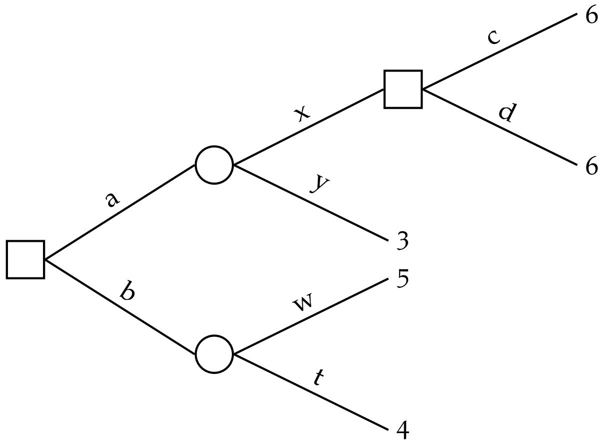

(A) Decision Tree

	xw	yw	xt	yt
ac	6	3	6	3
ad	6	3	6	3
bc	5	5	4	4
bd	5	5	4	4

(B) Decision Matrix

Figure 5.4: Decision Tree and Matrix for the Same Decision Problem

NEC says that these two representations are equivalent in terms of which decision should be made, therefore implying that the representation of different branches with choice nodes in a dynamic decision problem adds no information. All the information that is lost going from a decision tree to a decision matrix must be irrelevant (according to NEC). What information is lost then? The different branches with choice and chance nodes are lost. What information do these represent? They tell us in which period a decision is being made, and they also tell us in which of those periods how much uncertainty was resolved or not.

5.6.3 *Plan Reduction*

McClennen shows that SR and NEC together imply what he calls Plan Reduction (PR), which he defines in the following manner:

Plan Reduction (PR): Let T be any decision tree with associated set of plans S, and let G_S be the set of prospects associated with such plans. Then for any plan s in S and associated prospect g_s in G_S, s is in $D(s)$ iff g_s is in $D(G_S)$.

This condition states that a plan and the prospect associated with that plan must be indifferent, even in a decision context where multiple sequential decisions need to be made. Why is this implied by the former two conditions? SR gives us the rule, that plans should be indistinguishable from their associated prospects, and NEC guarantees that this is true for any type of decision, and not only those where only one decision needs to be made, since it says that there is no information held in the branches or number of decision nodes.[33]

33 Now that the three conditions have been stated, it is important to notice that they are independent from each other. This result is part of one of McClennen's propositions, that none of these conditions is logically implied by the conjunction of the other two, nor by any subset of the other conditions. In order to show this he shows several examples when all the conditions but one are satisfied throughout pages 123 to 125, where he states the final result.

5.6.4 *KP Model and PR*

Plan reduction is violated in the KP model. This shouldn't come as a surprise if we consider that KP's model allows one to have preferences regarding **how** a decision is made. The idea behind PR is precisely that the decision process is irrelevant for the decision maker. The idea behind KP's model is just the opposite. The number of branches in KP makes a difference, as discussed above, but this disappears if one tries to represent the problem using a decision matrix, where only probabilities and consequences can play a role. Therefore, NEC is violated.

Let us consider the coin toss example introduced in the previous chapter. If NEC is respected, then it follows that one must be able to represent this problem using a decision matrix and arrive at the same decision. Therefore, if representing this problem using a decision matrix leads to a different solution, then by *reductio* it follows that NEC is indeed violated by KP. Recall that the solution for the problem, for the functions and values assumed, was that the decision maker should choose $d_0(b)$. Now let us represent the same problem using a decision matrix as follows:[34]

34 Following convention, I represent the period 1 decision points that will not be reached.

	heads	tails
$d_0(a), d_1(a)$	$U(0)$	$U(10)$
$d_0(a), d_1(b)$	$U(0)$	$U(10)$
$d_0(a), d_1(c)$	$U(0)$	$U(10)$
$d_0(b), d_1(a)$	$U(0)$	$U(10)$
$d_0(b), d_1(b)$	$U(0)$	$U(10)$
$d_0(b), d_1(c)$	$U(0)$	$U(10)$

Figure 5.5: Decision Matrix for the Coin Toss Problem

It is in truth not possible to directly apply KP's solution concept when the problem is represented in this way; however, when represented in this way, the two available actions (in period 0 - the decision the agent actually has to make) become indistinguishable (because the represented states do not involve the time at which uncertainty is resolved), and therefore the solution prescribed by KP - namely that the decision maker should choose $d_0(b)$ - can no longer be justified. This happens because KP use an additional function to the utility function, that is applied between different stages of a compound lottery. When compound lotteries are merged into one single lottery, as it happens in a decision matrix, then this stage does not occur. By explicitly modelling this stage, KP violate NEC - if we want to use KP, we cannot reduce a decision tree into a decision matrix.

SR is not violated. This is because EMU is only applied at state nodes (the stop points at the decision tree named $x_i(j)$), which only occur just before a decision node. Therefore, there is no special stop in between lotteries, and the reduction of lotteries without decisions in between can be done in an orthodox fashion.

However, given the equivalence between SR and NEC on the one side, and PR on the other, PR is still violated by KP model (when γ is raised to a power other than 1).

It is important to note how this links back to the definition of independence used by McClennen. He uses a definition that applies to the whole decision tree, and as such it assumes that one can merge lotteries across time into one composite lottery. In order to apply the independence axiom across periods, McClennen has to assume the reduction axiom. In fact, Segal (1992) provides a very good description of the confusion between independence on its own, and independence which includes the reduction axiom.[35] The former he calls compound independence, and the latter mixture independence (which, not surprisingly can be derived from compound independence together with the reduction principle). McClennen's independence, since it is applied to the entire tree, inevitably includes reduction. Therefore, independence without reduction (applied within a single period) is not violated by KP's model.[36] This discussion will be taken further in the next section.

35 He refers to reduction of compound lotteries, but these two principles are closely related, as I explain below.

36 I discussed in chapter 2 that the STP is in some senses weaker than the independence axiom; however, reduction is implied by the STP due to the assumption that all partitions of the states of the world are allowed. As a result, there is no STP without reduction and therefore, the STP would be violated by KP's model, but only due to the impossibility of separating the STP from the partition assumption.

5.7 THE STATUS OF PR

It comes as good news that KP's model does not violate DC. This means that the model is not subject to the usual criticisms that are made against models that do not subscribe to DC. However, before dismissing all concerns, it is important to keep in mind that KP's model does violate PR. So we need to discuss how problematic this violation is, and in doing so assess if what is gained by adding more explicit elements to the model compensates for what is lost by foregoing PR.

Recall the definition of plan reduction stated above: this condition entails that in a dynamic decision scenario one must be indifferent between a plan and the prospect associated with that plan. A plan is the sequence of choices made throughout the decision tree and it includes the path from the beginning of the decision making process to the terminal nodes, while a prospect is the lottery that attributes the combined probabilities of the chance nodes in a path to the final outcomes of that same path. In informal terms, plan reduction entails that there is nothing more to plans - in terms of decision making - than the prospect associated with it. Therefore, whatever is not represented in the prospect cannot be said to influence the decision maker in her decision. Since prospects only include probabilities and final outcomes, the idea behind this axiom is that the process by which probabilities and outcomes are determined does not matter for the decision to be made (or if it does it needs to somehow be represented as a final outcome).

This axiom's desirability lies mainly in the simplicity and tractability it gives to orthodox theory. However, there are several reasons why it can be considered inappropriate as a condition of rationality. I will here discuss the main objections to this type of rationality constraint, and argue how these objections get exacerbated in certain contexts of decision making. PR is closely related to the well known Reduction of Compound Lotteries (RCL) axiom.[37]. The reason why I note this is because RCL has been extensively discussed in the literature, and it is the subject of many objections in the decision theory literature. The main objection is that RCL does not allow for a decision maker to have preferences regarding how many lotteries she participates in. An agent might love gambling and therefore get utility from the fact that there are several lotteries involved in the decision making process rather than only one, or conversely, she might dislike gambling and much prefer one single lottery to take place rather than many. These are certainly acceptable motivations for a decision maker and therefore an argument against RCL and PR by extension.

Especially in the context of Humean decision theory, where what an agent desires or not is not discussed, it would be rather strange to preclude the agent from having preferences towards more or less gambling. In fact, as mentioned in the introduction, vNM were aware of this shortfall of their theory, and they referred to the fact that their model could not account for a utility for gambling as paradoxical (Von Neumann and Morgenstern (1944), p. 28), but nonetheless there was little that could be done. They

37 The general idea of RCL is very similar in nature to PR, but applied to static situations. Informally it says that an agent should be indifferent between a simple lottery and a compound lottery for which the multiplication of the probabilities results in the same values as those in the simple lottery.

were aware of the difficulties the RCL axiom resulted in, but they believed it to be the price to pay for a concise model with a representation theorem that gives us expected utility maximisation.

In fact, there is much empirical evidence that decision makers do not always obey this axiom (for example, Kahneman and Tversky (1963), and Conlisk (1989)): agents do not perceive compound lotteries as the same as their simple lottery equivalent.[38] In fact, Segal (1992) claims that most criticisms made of the independence axiom (and, indeed, of expected utility theory) are directed towards the reduction axiom. This is because this criticism is directed towards mixture independence, which, as mentioned in the end of the previous section, is implied not only by independence, but also by reduction.

Having the process of decision making impact the agent is more plausible in some contexts than in others. For example, if one is deciding for oneself it makes sense that one also considers the decision process itself. In contrast, for models where the decision is being made for someone else (for example, when a social planner is the one making the decision), PR is quite a reasonable axiom. It might be argued that such a principle must only be taken as normatively valid for a decision maker who acts in an essentially public capacity, i.e., who makes decisions as an agent for other

38 I am here assuming that there is no appeal to a mode with highly individuated outcomes, because even if there was, agents' behaviour would still contradict the spirit of the reduction axiom, and it is the reasonableness of non-reductive behaviour that I am discussing in this section.

persons. The reason to do so is that although individuals can have direct preferences for forms of gambles, these preferences must be understood as purely personal.[39]

Worries originating from the fact that the decision making process can affect the agent get exacerbated in the context of dynamic decision making, where the decision making process occurs throughout time and takes longer. In fact, the violations discussed in the static case necessarily appeal to a process of sequential betting, which is precisely what characterises dynamic decision contexts. With regards to preferences for the timing of the resolution of uncertainty, this is certainly a concern that only arises in dynamic decisions. Therefore, it does seem that it is not the case that PR should be thought of as a rationality requirement, especially in the contexts described in this and the previous paragraph.

In the particular case of the KP model, modelling time as a primitive allowed one to make utility dependent on something not classified as a consequence while not violating compound independence nor dynamic consistency. This offers a nice alternative to the re-individuation strategy, as it allows one to model a preference for the timing of the resolution of uncertainty explicitly, and to avoid the problems with individuation described in Chapter 3. There is a violation of PR; however, as mentioned above, it is reasonable to violate this axiom in a dynamic situation where decision makers are deciding for themselves, and do care about how the decision is made.

39 I associate this view with Harsanyi, but I was unable find a reference to support this.

However, given that foregoing PR makes an axiomatic model less tractable (all things equal), in a static context where the agent does not really care about the process of decision, the tractability gained with PR (or its static counterpart, RCL) might very well surpass the accuracy in representing behaviour that violating PR allows for. Regarding this precise trade-off, Segal says the following:

> Of course, it is nicer to have a model in which this axiom is accepted. But it is my belief that releasing the reduction axiom is the lesser of the many other evils we may choose.
>
> Segal (1992) (p. 167)

I agree with Segal's claim, but I do think it does not have to be an all-or-nothing decision. I believe that a class of models that increases structure and does not subscribe to PR should be used in contexts that justify foregoing PR. That is why I would argue that no model should be used all the time across all decision contexts. For example, in a dynamic setting where the decision maker is deciding for herself it would seem that KP's model would be much more appropriate than vNM. However, in the very simple decision of taking an umbrella or not, vNM should be used, as the considerations above would not make a difference and vNM's model is the most tractable. Additionally, since there are many global properties that, like certainty and hope, may influence decision making, the possibility of using different models according to context would allow one to model these factors explicitly while preserving tractability.

5.8 CONCLUSION

In this chapter I discussed which standard axioms of rationality does KP's model violate. I showed that the added structure that results from modelling time explicitly allows the model to preserve dynamic consistency, and instead violate plan reduction. Plan reduction can be justifiably foregone in certain contexts, one of them being the situations KP are trying to model. Therefore, KP's model allows to model concerns such as hope explicitly without foregoing what might be considered essential rationality conditions.

6

CONCLUSION

6.1 CONCLUSION

The goal of this thesis was to argue against the use of orthodox decision theory in specific contexts, and to offer alternative ways to model decision making. In the previous chapters, I discussed the disadvantages of the use of the orthodox model in cases where agents have a preference for global properties of actions, and the advantages of an alternative model of choice that models a global property of an action explicitly. As I stressed throughout this thesis, my goal was never to fully abandon orthodox decision theory, but rather to identify certain contexts in which the advantages of using a more complex but more realistic model surpass the disadvantages of doing so. I argued that the use of different models might be justified (or not) depending on the context in which we want to use them: in certain contexts one might be justified in dropping the reduction axiom, but in others it might be reasonable to keep it (for example, if the decision maker hires a planner to execute her decision plan). I ultimately suggested that a pluralist approach to decision theory should be adopted.

I started by identifying a decision situation where the only way to preserve orthodox decision theory is via a finer individuation of consequences. I put forward several different concerns with the use of the individuation strategy in the situation described and in similar cases (those cases where the property that is being added to the consequences via individuation is a global property of the action). My main argument was that certain properties (e.g., hope) are such that defining them as consequences modifies them. I then offered alternatives to individuation that in my view have less disadvantages and more advantages.

In specific, I introduced Kreps and Porteus's model of decision making, which models explicitly - as opposed to via a finer individuation of the consequence set - a preference for the timing of the resolution of uncertainty. In order to do so they add time as a primitive of the model. Then they add a function that represents agents' preferences with regards to one aspect of time - namely, when uncertainty is resolved. Therefore, there is an increase in structure similar to the one of vNM's framework relative to Savage's. The latter introduces states of the world as a primitive, and makes the probability function a subjective assessment of the likelihood of these states of the world. KP do something similar with time.

I then showed that the added structure that results from adding time as a primitive allows the model to preserve dynamic consistency, which is one of the main concerns with models of dynamic decision making that deviate from orthodoxy. However, one of the orthodox axioms is indeed violated: plan reduction. I argued that this axiom can be justifiably foregone in certain contexts, one of them being precisely the situations KP are trying to model (i.e., dynamic decisions where an agent is implementing her preferred plan herself). As a result, the disadvantages of the added complexity (such as foregoing important axioms) are very limited.

Next steps: Pluralist Decision Theory

In the same way that certainty and time may influence our decision making, there are many other factors that may as well and that, for similar reasons, one would not want to represent as an orthodox consequence. For example, an agent might care about the worst possible consequence an alternative might lead to by having a preference for the alternative with the highest lowest value. However, having a model that includes all of these factors explicitly would result in a very complex and not at all tractable model of decision making. A pluralist approach to decision theory could be the solution: one that encodes, as a first step, a mapping from contexts to decision theoretical models. On the other hand, having different models to deal with all the different possible motivations inevitably raises the question: when to use one rather than another? I do not think this is a reason to abandon pluralism, but rather a reason to explore the question of what decision contexts map best to which decision theories.

BIBLIOGRAPHY

Allais, M. (1953). Le comportement de l'homme rationnel devant le risque: critique des postulats et axiomes de l'école américaine. *Econometrica*, 21:503–546.

Anscombe, F. and Aumann, R. (1963). A definition of subjective probability. *The Annals of Mathematical Statistics*, 34:199–205.

Backus, D., Foresi, S., and Wu, L. (2004). Accounting for biases in black-scholes. *Available at SSRN 585623*.

Bernoulli, D. (1954/1738). Exposition of a new theory on the measurement of risk. *Econometrica*, 22:23–36.

Broome, J. (1991). *Weighing Goods: Equality, Uncertainty and Time*. Blackwell.

Broome, J. (2006). Weighing lives.

Buchak, L. (2013). *Risk and Rationality*. Oxford University Press.

Cantor, G. (1895). Beiträge zur begründung der transfiniten mengenlehre. *Mathematische Annalen*, 46:481–512.

Conlisk, J. (1989). Three variants on the allais example. *The American Economic Review*, 79:392–407.

Cubitt, R. and Sugden, R. (2001). On money pumps. *Games and Economic Behavior*, 37:121–160.

Debreu, G. (1954). Representation of a preference ordering by a numerical function. In Thrall, D. and Coombs, editors, *Decision Processes*, pages 159–1657. John Wiley.

Diamond, P. (1967). Cardinal welfare, individualistic ethics, and interpersonal comparison of utility: Comment. *Journal of Political Economy*, 75:765–766.

Dreze, J. and Modigliani, F. (1972). Consumption decisions under uncertainty. *Journal of Economic Theory*, 5:308–335.

Epstein, L. and Zin, S. (1989). Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework. *Journal of the Econometric Society*, 57:937–969.

Epstein, L. and Zin, S. (1991). Substitution, risk aversion, and the temporal behavior of consumption and asset returns: An empirical analysis. *Journal of Political Economy*, 57:236–286.

Fishburn, P. (1970). *Utility Theory for Decision Making*. John Wiley.

Fleurbaey, M. (2008). *Fairness, Responsibility, and Welfare*. Oxford University Press.

Gaifman, H. and Liu, Y. (2015). Context-dependent utilities. In Wiebe van der Hoek, W. H. H. and fang Wang, W., editors, *Logic, Rationality, and Interaction*, pages 90–101. Springer.

Green, J. (1987). Making book against oneself, the independence axiom and nonlinear utility theory. *Quarterly Journal of Economics*, 102:785–795.

Hammond, P. (1988). Consequentialism and the independence axiom. In Munier, B. R., editor, *Risk, Decision and Rationality*. D. Reidel Publishing Company.

Hammond, P. (1988a). Consequentialist foundations for expected utility. *Theory and Decision*, 25:25–78.

Jeffrey, R. (1965). *The Logic of Decision*. McGraw-Hill.

Kahneman, D. and Tversky, A. (1963). Prospect theory: An analysis of decision under risk. *Econometrica*, 47:263–292.

Karni, E. and Schmeidler, D. (1991). Utility theory with uncertainty. In Hildenbrand, W. and Sonnenschein, H., editors, *Handbook of Mathematical Economics, Vol. IV*, pages 1763–1831. Elsevier.

Kreps, D. and Porteus, E. (1978). Temporal resolution of uncertainty and dynamic choice theory. *Econometrica*, 46:185–200.

Luce, R. and Raiffa, H. (1989/1957). *Games and Decisions. Introduction and Critical Survey*. Dover Publications.

Machina, M. (1981). Rational decision making versus rational decision modelling. *Journal of Mathematical Psychology*, 24:163–175.

Machina, M. (1989). Dynamic consistency and non-expected utility models of choice under uncertainty. *The Journal of Economic Literature*, 27:1622–1668.

Mas-Colell, A., Whinston, M., and Green, J. (2006). *Microeconomic Theory*. Oxford University Press.

McClennen, E. (1990). *Rationality and Dynamic Choice: Foundational Explorations*. Cambridge University Press.

Mossin, J. (1969). A note on uncertainty and preferences in a temporal context. *The American Economic Review*, 59:172–174.

Pettit, P. (1991). Decision theory and folk psychology. In Bacharach, M. and Hurley, S., editors, *Essays in the Foundations of Decision Theory*, pages 147–175. Basil Blackwell.

Pratt, J. (1974). Some comments on some axioms for decision making under uncertainty. In M. Balch, D. M. and Wu, S., editors, *Essays on Economic Behavior Under Uncertainty*, pages 82–92. Routledge.

Ramsey, F. (1926). Truth and probability. In Braithwaite, R. B., editor, *Foundations of Mathematics and Other Logical Essays*, pages 156–198. Routledge.

Russell, B. (2009/1950). *Autobiography*. Routledge.

Samuelson, P. (1952). Probability, utility and the independence axiom. *Econometrica*, 20:670–678.

Saramago, J. (2013/2004). *Seeing*. Random House.

Savage, L. (1972/1954). *The Foundations of Statistics*. Dover Books on Mathematics.

Schick, F. (1986). Dutch bookies and money pumps. *Journal of Philosophy*, 83:112–119.

Segal, U. (1992). The independence axiom versus the reduction axiom: Must we have both? In Edwards, W., editor, *Utility Theories: Measurements and Applications*, pages 165–183. Springer.

Seidenfeld, T. and Schervish, M. (1983). A conflict between finite additivity and avoiding dutch book. *Journal of Philosophy*, 50:398–412.

Sen, A. (1993). Internal consistency of choice. *Econometrica: Journal of the Econometric Society*, 61:495–521.

Shafer, G. (1986). Savage revisited. *Statistical Science*, 1:463–501.

Spence, M. and Zeckhauser, R. (1972). The effect of the timing of consumption decisions and the resolution of lotteries on the choice of lotteries. *Econometrica: Journal of the Econometric Society*, 40:401–403.

Steele, K. (2014). Choice models. In Cartwright, N. and Montuschi, E., editors, *Philosophy of Social Science: A New Introduction*, pages 185–207. Oxford University Press.

Strotz, R. (1955). Myopia and inconsistency in dynamic utility maximization. *The Review of Economic Studies*, 23:165–180.

Tallarini, T. (2000). Risk-sensitive real business cycles. *Journal of Monetary Economics*, 45:507–532.

Volij, O. (1994). Dynamic consistency, consequentialism and reduction of compound lotteries. *Economics letters*, 46:121–129.

Von Neumann, J. and Morgenstern, O. (1944). *Theory of Games and Economic Behavior*. Princeton University Press.

Weil, P. (1989). The equity premium puzzle and the risk-free rate puzzle. *Journal of Monetary Economics*, 24:401–421.

Weil, P. (1993). Precautionary savings and the permanent income hypothesis. *The Review of Economic Studies*, 60:367–383.