



Fairness, ambiguity and dynamic consistency

Richard Bradley¹ · H. Orri Stefánsson²

Accepted: 11 November 2024
© The Author(s) 2025

Abstract

Considerations of dynamic inconsistency have figured prominently in debates over the rationality of preferences that violate the separability conditions characteristic of expected utility theory. These debates have mostly focused on risk- and ambiguity averse preferences, but analogous considerations apply to preferences for fairness. We revisit these debates in the context of a specific hypothesis regarding the violations of separability by such preferences, namely that they are potentially both explained and rationalised by non-instrumental attitudes to the chances of the various possibilities involved. Our main aim is to argue that, first, when these violations of static separability are motivated by such non-instrumental attitudes to chances, then they need not result in dynamically inconsistent behaviour. Second, and more generally, considerations of dynamic consistency do not, we argue, undermine the rationality of non-instrumental attitudes towards the distribution of chances, despite the fact that such attitudes give rise to violations of the static separability assumptions of expected utility theory.

Keywords Dynamic choice · Fairness · Ambiguity · Consequentialism · Reduction · Chances

1 Introduction

Considerations of dynamic inconsistency have figured prominently in debates over the rationality of preferences, or associated patterns of behaviour, that are in violation of one or other of the (static) separability conditions characteristic of expected utility theory, such as von Neumann and Morgenstern's (1944) Independence axiom

✉ H. Orri Stefánsson
orri.stefansson@philosophy.su.se

Richard Bradley
r.bradley@lse.ac.uk

¹ Department of Philosophy, Logic and Scientific Method, London School of Economics and Political Science, London, UK

² Department of Philosophy, Stockholm University, Stockholm, Sweden


and Savage's (1954) Sure-thing principle. Early versions of these discussions were focused on the (separability-violating) forms of risk aversion exhibited in the famous Allais paradox and specifically on the question of whether or not such risk averse preferences were shown by their implications for sequential choice to be irrational (see inter alia Buchak, 2014; Cubitt, 1996; Gauthier, 1997; Hammond, 1989). More recently, similar considerations have featured in the debate over the rationality of preferences commonly observed in the Ellsberg paradox (Epstein & Le Breton, 1993; Hill, 2020; Siniscalchi, 2009, 2011) and specifically to support the charge that ambiguity aversion is irrational (Al-Najjar & Weinstein, 2009; Fleurbaey, 2018). Less frequently discussed, but equally significant, are the preferences for fairness exhibited in a famous example from Peter Diamond (1967) and which imply a violation of the very widely endorsed weak separability condition of State Dominance, and which have similar implications for sequential choice.

In this paper we revisit these debates in the context of a specific hypothesis regarding the violations of separability observed in these famous cases, namely that they are potentially both explained and rationalised by agents' valuing intrinsically, as well as instrumentally, the chances of the various possibilities involved, in a way that results in chances across different individuals and possible events having decreasing marginal value. (By the chance of an outcome we mean its objective probability of realisation and not the agent's subjective degree of belief in it.) The focus will be on the argument that (a) violations of (static) separability must lead to dynamically inconsistent choices and/or a range of attendant forms of putatively irrational behaviour, including an aversion to information and a sensitivity to the timing of the resolution of uncertainty; and that (b) these consequences strongly support the claim that the attitudes leading to these violations (ambiguity-, unfairness- and risk-aversion) are irrational. Our main aim is to argue that these claims are false when the violations of static separability in question are motivated by certain kinds of attitudes to the chances of good or bad outcomes. In doing so we draw on existing results of a more general kind, but display their significance for when these motivations are in play.

Let's begin by giving a somewhat informal version of the kind of argument that is the object our attention, starting with a statement of a canonical version of the requirement of static separability characteristic of expected utility theory. We will throughout work with what we think is the simplest separability condition that is required for our purposes, Event Dominance, which implies the aforementioned Sure-thing principle.¹ Informally put, Event Dominance says that for any acts (i.e., choice-alternatives) f and g , if for any possible event E , f is weakly preferred to g , then f is weakly preferred to g overall; if, in addition, for some possible event E' ,

¹ In fact, the two are equivalent in the presence of Savage's (1954) principles P1 and P3. The Sure-thing principle is also conceptually very close to the aforementioned Independence axiom; the main difference between them is that the former is formulated in a framework without probabilities (used to model choice under "uncertainty") whereas the latter is formulated in a framework where probabilities are given (used to model choice under "risk").

Fig. 1 Event dominance and sequential choice



f is strictly preferred to g , then f is strictly preferred to g overall. By implication, if for any possible event E , one is indifferent between f and g , then one is indifferent between f and g overall. Event Dominance implies that how two alternatives compare in one event can be determined separately from how they compare in other events. But, as we shall see, this requirement is inconsistent with some common attitudes to ambiguity, fairness, and risk.

To study the implications of violations of Event Dominance we turn to sequential decision problems in which the agent must choose a plan for making choices at later times, after the resolution of some uncertainty. Such problems are represented by directed graphs (called ‘trees’), of the kind exhibited in Fig. 1, in which each square ‘choice’ node represents a moment of choice and the edges from them the actions that can be taken at that time, and each circular ‘chance’ node represents a moment at which some uncertainty is resolved and the edges from it the information that is thereby received. A plan is simply a tree with a single edge from each choice node. For instance, in Fig. 1, four such plans are available: to choose f if E is the case and also if E is not the case; to choose f if E is the case but otherwise choose g ; to choose g if E is the case and also if E is not the case; and to choose g if E is the case but otherwise choose f .

Two assumptions are crucial to the argument that preferences in violation of separability lead to dynamic inconsistency. The first is *Consequentialism*: in essence the principle that the choice made at any moment depends only on what could occur later. From this it follows that an agent should make the same choice at nodes rooting the same subtrees, irrespective of the larger tree that they might be embedded in.² The second assumption of *Reduction* requires preferences over plans to match

² Consequentialism is sometimes termed dynamic separability because of this implication.

preferences over the static (or reduced) acts that they correspond to, where the static act corresponding to a plan is the one which yields for each of the possible complete resolutions of uncertainty just the consequence yielded by the plan under that resolution. Finally, for an agent to be *dynamically consistent* (in a decision problem) it must be the case that at each choice node (in that problem) she implements the plan she chose *ex ante* (at the first choice node).

Suppose that, in violation of Event Dominance, a person (strictly) prefers f to g in both events E and $\neg E$, but nevertheless (strictly) prefers g to f overall. By Reduction and given the preference for g over f , the person should, *ex ante*—that is, before arriving at the chance node³ and thus before knowing whether E is true or false—prefer to ‘go down’ at both choice nodes in Fig. 1. However, by Consequentialism and given the preference for f to g in both events E and $\neg E$, she should, *ex post*—that is, after learning whether E is true or false—prefer to ‘go up’ at both choice nodes. Therefore, satisfaction of Consequentialism means that there is no choice node where the person implements the plan she should, by Reduction, choose *ex ante*. In other words, if she satisfies both Consequentialism and Reduction then she is dynamically *inconsistent*. (We take for granted here, as is standard in the literature, that the agent’s preferences determine her choices, so that the preferences we speak of are those revealed in choice.⁴)

Dynamic inconsistency arises here if the agent chooses myopically, at each moment of time, the course of action perceived as best at that moment. As is well documented in the literature, dynamic inconsistency can be avoided by agents who choose in a different way (see, e.g., the discussion in Strotz, 1955 and McClennen, 1990). Sophisticated agents do so by evaluating plans at each moment of time in the light of what they know about the choices they will make in the future, treating these anticipated future choices as constraints on their current ones. A sophisticated consequentialist brings their current choices into line with what they expected to choose later. For instance, the person considered in the last paragraph would, as a sophisticated consequentialist, even *ex ante*, choose to go up at both choice nodes. In contrast, the resolute agent makes their choices at each (later) moment in time in line with plans adopted *ex ante*. The person considered in the last paragraph would, as a resolute chooser, even *ex post*, choose to go down at both choice nodes. This amounts to committing to a course of action right at the beginning and sticking to it thereafter.

A general justification of violations of static separability must be calibrated to the view of what kind of agency is rational. If myopia is, then dynamic consistency must be rejected and if resoluteness is, then Consequentialism must be. Most decision theorists however consider rationality to require sophistication and specifically the choice of plans recommended by backward induction. A sophisticated agent

³ In an effort to keep the discussion as simplified as possible to begin with, we do not in Fig. 1 explicitly model the path prior to the chance node (although we will have reason to do so later, in Figs. 2 and 3).

⁴ In this regard our treatment differs from Siniscalchi (2011)’s. He takes dynamic consistency to be a property of preferences while treating sophistication as a property of choices, which means that he can allow that a sophisticated agent be dynamically inconsistent.

justifiably violates static separability only if either Reduction or Consequentialism fails. The main argument of our paper is that there are permissible ways of valuing chances and their distribution that justify violating both Reduction and Consequentialism in the right kind of circumstance.

We proceed as follows. In Sect. 2, we present the (static) choice problems that figure in Diamond's and Ellsberg's thought experiments and show how a hypothesis regarding how agents value chances, and hence their distributions across persons and states of the world, explains the separability-violating patterns of preferences Diamond and Ellsberg predict and/or defend. In Sect. 3, we study a sequential decision problem corresponding to Diamond's thought experiment and show how an agent with these attitudes to chances can avoid dynamic inconsistency. In Sect. 4 we do the same for a sequential version of Ellsberg's three-colour paradox. Finally in Sect. 5, we compare our findings about the Diamond and Ellsberg preferences to the dynamic implications of the risk- (or regret) averse preferences observed in Allais' paradox. We conclude that considerations of dynamic consistency do not militate against the rationality of the postulated attitudes towards chances of goods, despite the fact that they can motivate violations of the static separability assumptions of expected utility theory.

2 Attitudes to chances

The chances of benefit or cost associated with our actions, being objective features of our decision situation, are legitimate objects of both epistemic and practical concerns, including distributional ones. Such concerns regarding the distribution of chances, respectively across people and states of the world, can explain and rationalise the patterns of preferences exhibited in the famous thoughts experiments of Diamond and Ellsberg. Since these preferences are in violation, respectively, of the conditions of State- and Event Dominance, these concerns, if legitimate, refute the claim that rationality requires adherence to these conditions. Or so we claim; our aim now is to vindicate it.

2.1 Chances and diamond fairness

Let's begin with a case described by Peter Diamond. Suppose an (impartial) social decision maker must decide between giving a good to Robbie, giving it to Bobbie, or tossing a fair coin to decide who should get it. Many people would say that the last of these options is to be preferred to the other two because it is fairer in one respect: unlike the other two, it gives *both* Bobbie and Robbie a chance of securing the good (indeed an equal chance). Diamond took his example to serve as a refutation of Harsanyi's (1955) Utilitarianism which entails that all three options are equally preferable. Indeed, the example presents a challenge for any theory of choice that satisfies the widely endorsed condition of State Dominance (an implication of

Table 1 (a) Diamond fairness (outcomes) and (b) diamond fairness (chances)

	(a)			(b)	
	Heads	Tails		Heads	Tails
Unfair B	(0, 1)	(0, 1)	Unfair B	(0, 1)	(0, 1)
Unfair R	(1, 0)	(1, 0)	Unfair R	(1, 0)	(1, 0)
Fair	(0, 1)	(1, 0)	Fair	$(\frac{1}{2}, \frac{1}{2})$	$(\frac{1}{2}, \frac{1}{2})$

Event Dominance), which requires that two options are equally preferable if they result in equally preferable outcomes in any state of the world.⁵

To see this consider Tables 1a and b which respectively exhibit the final outcomes and the distributions of chances of outcomes associated with the three options, labelled as Unfair R, Unfair B and Fair. In Table 1a an ordered pair (x, y) denotes an outcome in which Robbie receives x units of the good and Bobbie y units; while in Table 1b such an ordered pair indicates that Robbie has a chance of x of receiving a unit of the good and Bobbie a chance y .

If we focus on Table 1a then the following argument presents itself. Suppose that you are indifferent as to whether Robbie or Bobbie has the good (i.e., between $(1, 0)$ and $(0, 1)$) perhaps because it benefits them equally and they are equally deserving recipients in all other respects as well. Then you are indifferent between the outcomes of all three options in case the coin lands heads and in case the coins lands tails. So, by State Dominance (and, by implication, Event Dominance) you should be indifferent overall between all three options.

The problem with this argument is immediately apparent when we shift attention to Table 1b, where the symmetry between the three options is broken. For it does not follow from the fact that you are indifferent between Robbie and Bobbie receiving the good that you are indifferent between various possible distributions of the chances of them doing so. On the contrary, if they cannot both get the good and dividing between them is not an option, then giving them an equal chance of receiving it is a natural way of ensuring impartiality in your treatment of them. The preference for fairness exhibited in Diamond's example—and which we will occasionally refer to as the 'Diamond preference'—is thus nothing more than a preference, *ceteris paribus*, for equality in the distribution of chances. Such a preference is not rationally required of course; our point is simply that it is rationally permissible.

Such a preference for equality in the distribution of chances is inexplicable if the value of a half-chance of the good is half the value of getting the good for sure. Its explanation thus requires that the value of the chances of a goods not (always) be a linear function of the chance itself; for short, that chances can be valued non-linearly. We suggest, specifically, that a half-chance of the good is, for Robbie and Bobbie, more than half as valuable as the good itself and, more generally, that the chances of many goods have diminishing marginal value (Stefánsson & Bradley,

⁵ State Dominance can thus be understood as a restricted version of Event Dominance, i.e., one that holds only when the 'events' are maximally specific, in the sense that they contain only one state.

2015). This hypothesis, if true, rationalises the preference for Fair over both Unfair R and Unfair B.

Whether the possibility that chances have diminishing marginal values should be accommodated by abandoning State Dominance or by refining the description of outcomes so that they include the chances with which they obtain is not a question that we need to settle here. Of more immediate interest is that an analogous concern for equality in the distribution of the chances, explicable in terms of the same hypothesis, is manifest in the apparently unrelated thought experiment of Ellsberg (1961), something to which we now turn.

2.2 Chances and ambiguity

Suppose that ten black and ten red balls have been distributed between two urns, so that there are ten balls in each but with proportions of black and red balls unknown to you. One of the urns will be selected on the basis of the toss of a fair coin and a ball drawn at random from the selected urn. You may either bet on the ball being black, bet on it being red or bet on red if the coin lands heads and black if the coin lands tails. Your decision problem is represented in Table 2 in which the alternatives Ambiguous B, **Ambiguous R** and **Risky** correspond to the three just described and the outcome pairs $(x, 1 - x)$ correspond to the quantity (either 0 or 1) of the good you receive (i.e., the prize you win for betting on the correct ball) in case the drawn ball is red (x) or black ($1 - x$).⁶

Since nothing distinguishes Ambiguous B from **Ambiguous R**, it would not be unreasonable for you to be indifferent between them. But since this is so irrespective of whether the coin lands heads or tails, Event Dominance then requires that you are indifferent between both and **Risky**. That is, in the event of Heads you are indifferent between the two possible outcomes of the three alternatives, $(1,0)$ and $(0, 1)$, and the same is true in the event of Tails; so, by Event Dominance, you should be indifferent between the three alternatives. But many think that you do have reason to prefer **Risky** to the others. Since you have no information about the proportion of black balls in the selected urn, both Ambiguous B and **Ambiguous R** amount to bets on maximally *ambiguous* possibilities (i.e., possibilities where you are as uncertain as you can be about the relevant chances). Not so for **Risky** however. For suppose that z black balls went into the urn selected when the coin lands heads and $1 - z$ into the other. Then you have a 50% chance of having a chance $z/10$ of winning the bet and a 50% chance of having a $(1 - z)/10$ chance of winning it. So your chance of winning is exactly one half, independently of z . Hence choosing **Risky** is equivalent to a bet on the draw of a black ball (or on a red one) from an urn containing 50% black balls and 50% red.

The difference is evident if we represent, as in Table 3, the three alternatives in terms of the range of chances $[x, 1 - x]$ of winning the good, in the event of the coin

⁶ Readers more familiar with the conventional presentation of ambiguity aversion via Ellsberg's three-colour paradox may be slightly confused by this way of doing so. The three-colour paradox is presented in Sect. 4, by which point the connection between the two should be clear.

Table 2 Ellsberg's two-colour paradox (quantities)

	Heads	Tails
Ambiguous B	(0,1)	(0,1)
Ambiguous R	(1,0)	(1,0)
Risky	(1,0)	(0,1)

Table 3 Ellsberg's two-colour paradox (chances)

	Heads	Tails
Ambiguous B	[0, 1]	[0, 1]
Ambiguous R	[1,0]	[1,0]
Risky	$[\frac{1}{2}, \frac{1}{2}]$	$[\frac{1}{2}, \frac{1}{2}]$

landing Heads or Tails, induced by each, with x being the chance of winning when all z balls are red and $1 - x$ the chance when all z are black.

A common and natural rationalisation of the 'Ellsberg preference' for **Risky** over the other two alternatives, is that there is a value in knowing what chances one faces. And certainty about one's chances of winning the good is indeed secured by the choice of **Risky** over both Ambiguous B and **Ambiguous R** because the former yields a chance of one-half of winning whereas the latter two leave one with chances of winning that could lie anywhere between zero and one. Now a preference for certainty is something that is predicted by the hypothesis of the diminishing marginal value of chances. For it follows from it that a half-chance of winning the prize is more than half as valuable as winning the prize for sure, and more generally that more equal distributions of the chances have greater value than less equal ones. It is this property of attitudes to chances of chances that, we suggest, rationalises the observed preference for **Risky** over both **Ambiguous R** and Ambiguous B (Bradley, 2016, 2017).

2.3 Attitudes to the resolution of uncertainty

Both Diamond's and Ellsberg's thought experiments reveal something general about attitudes to chances, namely that the separability-violating patterns of preferences or choices identified in these thought experiments are both explicable and rationalizable in terms of a desire to smooth or equalise the chances of obtaining a good across people or states of the world. Behaviourally this is manifested in a strict preference for *randomisation* between alternative acts or plans; more precisely, by a strict preference for an act or plan that amounts to the use of a randomising device to select between two equally-preferred acts or plans over either of these equally preferred acts or plans. For a smoother distribution of the chances is precisely what is achieved by use of such devices (as we saw above).

An agent who has such preferences for randomisation will also display, in the right circumstances, both an aversion to information and preferences regarding the timing of the resolution of uncertainty. Consider for example someone with a strict

preference for fairness of the kind identified by Diamond who seeks to use the toss of a coin to determine which of two persons should receive a good. If this choice is to serve the purpose of securing fairness, then the decision maker should seek to avoid knowing the outcome of the toss prior to settling on who gets the good in the event of it landing heads and who gets it if it lands tails. They should, in other words, be averse to learning the outcome of the coin toss at that point in time.⁷ Likewise, they should strictly prefer to have the opportunity to distribute the good to the two individuals as a function of how the coin lands, before the uncertainty as to whether it lands heads or tails is resolved, to having that opportunity after the uncertainty is resolved. (This amounts to a preference between what Eichberger et al. (2016) call *ex ante* and *ex post* randomisation.)

The same applies to someone with the Ellsberg (ambiguity averse) preferences, a point made by Siniscalchi (2009, 2011). Since a coin toss can be used, as demonstrated in the previous section, to fix the chances the decision maker faces, it follows immediately that someone with a preference for certainty about the chances they face will be averse to learning the outcome of the coin toss that secures such certainty. Equally they will prefer to have the opportunity of making a bet on the colour of the ball drawn from the urn prior to the resolution of the uncertainty as to which urn the draw will be from to having it afterwards. (Crucially, in the two-colour problem, the information on offer regarding how the coin lands does not enable the agent to make a better choice by making her beliefs about the chances she faces more accurate. As we shall see below, this is different in the three-colour Ellsberg paradox.)

To summarise the argument so far: information aversion and preferences regarding the resolution of uncertainty are to be expected under suitable conditions from someone with certain attitudes to the distribution of chances of goods; in particular, those that arise from assigning diminishing marginal value to these chances. If such attitudes are rationally permissible, as we have argued, then so too are these behavioural manifestations of it and to take them as an argument against ambiguity or unfairness aversion is simply to beg the question.

3 Fairness and dynamic consistency

In view of the violation of Event Dominance that it entails, it is to be expected that a preference for fairness of the kind identified by Diamond's example will open up the possibility of dynamically inconsistent behaviour. That this is indeed the case can be demonstrated by a three-person Diamond example, analogous to the three-colour Ellsberg one, that is discussed below. Table 4 represents such a static decision problem in which you must choose between giving the good to person R (Unfair) or, at

⁷ Siniscalchi (2011) correctly notes that 'information aversion' is a misleading term for this phenomenon which involves no dislike for information per se, but rather a disposition to avoid acquiring it when so doing interferes with the achievement of the agent's goal—in this case, the smoothing of the distribution of chances.

Table 4 (a) Three person diamond (outcomes). (b) Three person diamond (*ex ante* chances)

	(a)		(b)	
	Heads	Tails	Heads	Tails
Fair	$(0, 1 - \varepsilon, 0)$	$(0, 0, 1 - \varepsilon)$	Fair	$(0, \frac{1}{2}, \frac{1}{2})$
Unfair	$(1, 0, 0)$	$(1, 0, 0)$	Unfair	$(1, 0, 0)$

a small cost ε (borne by the person receiving the good), tossing a coin to decide whether person B or person G receives it (Fair), with table (a) displaying the final outcomes for each of the three persons in terms of quantity of the good received in each of the two possible states of the world (Heads and Tails) and table (b) displaying the chances of obtaining the good conferred on the three persons in these states. In (a) an ordered triple (x, y, z) denotes an outcome in which R receives x units of the good, B receives y units and G, z units; while in (b) such an ordered triple denotes that R has a chance of x of receiving the relevant quantity of the good, B a chance y and G a chance z .

Now an impartially benevolent decision maker of the kind postulated by Diamond will be indifferent between the outcomes $(0, 1 - \varepsilon, 0)$ and $(0, 0, 1 - \varepsilon)$. But we assume that she would strictly prefer the outcome $(1, 0, 0)$ to both $(0, 1 - \varepsilon, 0)$ and $(0, 0, 1 - \varepsilon)$, for any $\varepsilon > 0$; otherwise her choices would be *inefficient*, in that they would sacrifice total good for no gain. State Dominance therefore requires a strict preference for Unfair over Fair. A concern with equalising the chances as much as is possible, on the other hand, motivates a strict preference for Fair, so long as ε is sufficiently small.⁸

Consider now the representation in Fig. 2 of an associated sequential decision problem, displaying the outcomes in terms of the quantity of the good and the chance of obtaining it conferred by the plans *CF* and *CU* contingent on how the coin lands.⁹ (Here, and in what follows, we use italicised capital letters for plans, but we continue to use non-italicised capital letters, or capitalised names, e.g., Fair, for particular static acts.) If we assume Reduction and that Fair is preferred to Unfair in the static problem, then in the sequential problem, the plan *CF* (i.e., going up at both choice nodes) is preferred *ex ante* to the plan *CU* (i.e., going down at both choice nodes). On the other hand, *ex post*, when all relevant uncertainty has been resolved, efficiency requires the choice of *U* at both choice nodes. But this is in violation of the requirement of dynamic consistency.

What should we make of this? Although the violation of State Dominance entailed by preferences for fairness would seem to induce dynamically inconsistent behaviour in certain circumstances, it does not follow that a preference for fairness is irrational. On the contrary, it might seem that the correct conclusion to draw is

⁸ A preference for Fair in the original Diamond set-up does not of course logically entail it in three-person one. The point is simply that the sort of concerns motivating it in the former case would also do so in the latter.

⁹ Both in Figs. 2 and 3, we include a choice node (the very first choice node) that may seem superfluous since we never consider the case where the person chooses 'down' at that node. We nevertheless include those choice nodes to signal that this is where we assume that what we call the '*ex ante* preference' holds, that is, the preference before any uncertainty has been resolved.




Fig. 2 Three person sequential diamond case

that dynamic consistency is itself not a requirement of rationality. For in this example something conatively relevant changes as a result of the resolution of uncertainty, namely the chances conferred on the three individuals whose wellbeing is the concern of the decision maker! So the preference change involved here is grounded in reasons that the agent takes to be relevant to her choices, and in the fact that when the agent’s informational situation changes, these reasons support different options.

On the other hand, there is undoubtedly something myopic about such behaviour since the agent should anticipate this change in her informational state and consequent change in her preferences, and adjust her choice of plan in the light of this. More specifically, it should be apparent to her that her *ex ante* commitment to distribute the good to either B or G in accordance with the result of the coin toss will not survive in the *ex post* conditions in which efficiency requires her to give it to R. For once she has learnt the outcome of the coin toss, it will no longer be possible to use its outcome as a means to treat people fairly. A more sophisticated agent would therefore, by dint of backwards induction on her expectations about her future choices, plan on giving the good to one of the individuals independently of how the coin lands and avoid the small loss $(-\epsilon)$ induced by the use of a randomising device. Reduction then implies that in the corresponding static decision problem only Unfair is a permissible choice.

This argument concludes in the impermissibility of choosing Fair in Diamond’s static decision problem, a conclusion that is contradicted by the hypothesis that a concern with the distribution of chances is rationally permissible. It is tempting therefore to reject Reduction in order to resolve this dilemma. But although Reduction is undoubtedly a questionable condition (more on this later), there is another more pressing problem with the argument: it is intuitively too pessimistic about our ability to use devices like coin tosses to treat people fairly. After all we do in fact do this quite successfully all the time even though there is always the question, once the coin has been tossed, of whether we want to stick to our plan. This suggests that

a different diagnosis is called for, one which allows that a sophisticated agent can choose resolutely.

Note that the choice of plan *CF* benefits both R and G *before* the coin is tossed, by conferring each with a chance of receiving the good. (On some moral views it also confers an *impersonal* benefit.) Now if chances have intrinsic as well as instrumental value, then part of this benefit is independent of how the coin subsequently lands; that is, the individuals are benefited by the use of a coin even if the way it ultimately lands is not in their favour. (Similarly, the impersonal benefit (if any) remains however the coin lands.) For the decision maker to then fail to implement the plan of distributing the good in accordance with how the coin lands, by choosing *U* after coin has landed, would be to relinquish or destroy these benefits. So, contrary to what was argued above, it would not be rational for a person motivated to equalise the chances of obtaining the good to choose *U* at this point.

The upshot is that no dynamic inconsistency should arise because the decision maker will anticipate implementing plan *CF* once the coin has landed in virtue of it being the best course of action to take at that moment in time. The implications of this for the status of the separability conditions, both static and dynamic, depend however on exactly how the decision problem is modelled. If the final consequences of the pursuit of each plan are individuated narrowly in terms of the quantity of the good achieved, then the implementation of plan *CF* at the second node will require a violation of Consequentialism because the choice of *F* over *U* at this point is contingent on a benefit that was realised in the past (the conferment of chances induced by the choice of *CF*). This feature of the motivation makes the reasoning described here an instance of the kind of “sophisticated resoluteness” that Rabinowicz (1995, 2020) calls *wise* choice.

On the other hand, if the final consequences are modelled in a holistic way so as to reflect the conferment of chances on each individual by the choice of plan, then the choice of *F* at the second node has a consequentialist rationalisation: it enables the realisation of the benefit of conferring equal chances on B and G (in contrast to the choice of *U* which would see these benefits withdrawn). So, Consequentialism is at least formally satisfied, even if perhaps the philosophical view that Consequentialism is meant to capture is violated. In this case however, on pain of a violation of Reduction, the static version of the Diamond decision problem must also be remodelled with the consequences more finely individuated. Such a remodelling would, as noted before, render the use of a coin to settle the distribution of the good (at least formally) consistent with State Dominance (contrary to Diamond’s own diagnosis). In any case, whichever modelling strategy is adopted, the argument from the threat of dynamic inconsistency to the impermissibility of the postulated attitudes to chances fails.

A final note. The diagnosis offered here, however modelled, also serves to dispel a frequently raised puzzle regarding the benefits of randomisation (see, for instance, Machina, 1989). If it is better *ex ante* to distribute a good by means of a coin toss, surely this is true *ex post* as well? So once the coin has landed heads say, shouldn’t one toss the coin again rather than give it, as planned, to person B? An answer suggested by our diagnosis is: No, because to toss the coin again would make it the case that you treated G more favourably than B by conferring an overall chance on G of

Table 5 Ellsberg's three-colour paradox

	R	B	G
B0	0	1	0
R0	1	0	0
B1	0	1	1
R1	1	0	1

three-quarters of obtaining the good, compared to the overall chance of one-quarter conferred on B.¹⁰ To fail to live up to the commitment of giving the good to B in the event of the coin landing heads would thus have the consequence of degrading the benefit of a fair chance initially conferred on B and G.

4 Ambiguity aversion and dynamic consistency

We noted earlier that considerations of dynamic consistency have figured prominently in debates regarding the rational permissibility of ambiguity averse preferences of the kind exhibited in the Ellsberg paradoxes. The static version of Ellsberg's three-colour decision problem is represented in Table 5, where the two pairs of options displayed in this decision matrix yield a prize (1 unit of a good) contingent on the draw of a ball of a particular colour from an urn containing 30 red balls and 60 black or green ones in an unknown proportion. The prize is won if the ball drawn is black in the case of option B0, if it is a red or green in the case of R1, and so on. The frequently observed ambiguity averse pattern of preferences for R0 over B0 and for B1 over R1 are in violation of the Sure-thing principle because they should not be sensitive in this way to the outcome associated with the draw of a green ball.

Similarly, one of these preferences must violate Event Dominance. As already alluded to, this is a separability condition between events, meaning that it requires that subjects be able to evaluate alternatives event-by-event. So, conditionally on event G being true, subjects should be indifferent between R0 and B0 and also between B1 and R1. After all, in this event R0 and B0 result in the same outcome and so do B1 and R1. But note that conditionally on the disjunction R or B being true, B0 is the same alternative as B1 and R0 is the same alternative as R1. Therefore, one should, by Event Dominance, prefer R0 to B0 just in case one prefers R1 to B1. So, either the preference for R0 over B0 or that for B1 over R1 violates Event Dominance.

As with the static two-colour Ellsberg paradox, such ambiguity aversion is readily explained by a concern with the distribution of chances across possible states of the

¹⁰ This is not the only plausible explanation for why one should not toss the coin again. Another, and we think compatible, explanation is that once the coin has landed heads, B owns (and thus has a right to) the good, and therefore G has no claim to it anymore (see Nissan-Rozen 2019: pp. 490–492).

world, the relevant ones here taking the form of possible proportions of black and green balls in the urn. For the acts R_0 and B_1 both imply a fixed chance (one-third and two-thirds respectively) of getting the good. In contrast, for B_0 these chances are known only to lie between 0 and two-thirds and for R_1 between one-third and one. So someone who prefers to smooth the chances across events will be ambiguity averse in the sense of preferring to face chances that are independent of the state of the world (and thus ‘known’). They will, therefore, have grounds for preferring R_0 to B_0 and for preferring B_1 to R_1 .

To examine the argument that dynamic consistency is incompatible with ambiguity aversion in the static three-colour Ellsberg paradox, consider Fig. 3 which gives an extensive form version of it. Uncertainty is now resolved in two phases. At the first chance node, a selection is made at random between an urn containing 30 red balls and an unknown number n of black ones (the Mixed urn) and an urn containing $60 - n$ green ones (the Green urn), with the chance of the latter being selected equaling the proportion of the population of 90 balls that are green. At the second chance node (if it is reached), a ball is drawn at random from the Mixed urn containing red and black balls and the agent receives a quantity of the good (0 or 1) depending on the action chosen at this point in time.

At the initial choice node, the agent must decide between two plans: CB_X and CR_X (with X equal to 0 or 1, depending on the quantity of the good they receive in the event of the Green urn being selected). At the second choice node, if reached, the agent must choose between B_X and R_X (i.e. between B_0 and R_0 or between B_1 and R_1 depending on the case). Note that the plans CR_0 , CB_0 , CR_1 , and CB_1 correspond to the acts R_0 , B_0 , R_1 and B_1 in the static three-colour Ellsberg paradox. So, on the assumption of Reduction, at the initial node an agent with ambiguity averse preferences will prefer the plan CR_0 to the plan CB_0 but the plan CB_1 to the plan CR_1 . But once it has been revealed that the Mixed urn has been selected (when this is the case), the agent will either prefer (‘plan’) R_0 to B_0 and R_1 to B_1 or the other way




Fig. 3 Sequential Ellsberg paradox

around, since at this point R_0 and R_1 have exactly the same outcomes, as do B_0 and B_1 . But this amounts to a failure of dynamic consistency: *ex post* the agent does not prefer the continuation of her *ex ante* preferred plan.

Does this argument support the conclusion that ambiguity aversion is irrational? Our contention is that it does not. Dynamic inconsistency arises here because, for an agent with the kinds of attitudes to chances we have postulated, the ambiguity they face changes over time and specifically as a result of how the chances of obtaining the good conferred by the plans depend on which urn is selected: the Green one or the Mixed one. These chances are displayed in Table 6, with 6(b) giving the range of *ex ante* chances associated with each plan and 6(c) the range of *ex post* ones (for the case where it is the Mixed urn that is selected). For instance, at the first node the choice of CR_0 (the bet on a red ball being drawn) induces a chance of one-third of receiving the good, while the choice of CB_0 (the bet on a black ball being drawn) induces a chance of receiving it that ranges between zero, the case where all balls in the population are either green or red, and two-thirds, the case where they are either black or red. At the second node (when the Mixed urn has been selected) the choice of R_0 induces a chance of receiving the good that ranges between one-third and one, respectively the case when there are only black and red balls in the population of balls (hence when the Green urn is empty) and when there are only green and red ones. On the other hand, the choice of B_0 continues to induce a chance of receiving the good that ranges between zero, the case of only red and green balls, and two-thirds, the case of only red and black balls. So the chances associated with implementation of the plan CR_0 at the second node are different from the *ex ante* chances associated with it, while those associated with the plan CB_0 do not change. Similarly while the chances associated with the plan CR_1 do not change when it is implemented at the second choice node, those associated with the plan CB_1 do. At the first choice node CB_1 induces a chance of two-thirds, but at the second choice node B_1 induces a chance that ranges between zero and two-thirds.

What we observe here is that the first resolution of uncertainty *increases* the ambiguity associated with plans CR_0 and CB_1 . *Ex ante* no ambiguity is associated with either, but learning that only the non-green balls are in play changes the information situation significantly. Now the proportion of red balls in the urn from which the draw will be made (the Mixed urn) depends on the proportion of green balls in the original population, something that is unknown. Likewise the proportion that are

Table 6 Ellsberg’s three-colour paradox

	(a) Quantities of the good (R, B, G)	(b) <i>Ex ante</i> chances (all B, all G)	(c) <i>Ex post</i> chances (all B, all G)
CB_0	(0, 1, 0)	$\left[\frac{2}{3}, 0\right]$	$\left[\frac{2}{3}, 0\right]$
CR_0	(1, 0, 0)	$\left[\frac{1}{3}, \frac{1}{3}\right]$	$\left[\frac{1}{3}, 1\right]$
CB_1	(0, 1, 1)	$\left[\frac{2}{3}, \frac{2}{3}\right]$	$\left[\frac{2}{3}, 0\right]$
CR_1	(1, 0, 1)	$\left[\frac{1}{3}, 1\right]$	$\left[\frac{1}{3}, 1\right]$

either black or green. The upshot is that if the grounds for *ex ante* ambiguity aversion lies in a preference for known chances over unknown ones, the *ex post* chances associated with these plans are not similarly distinguished. So it is perfectly reasonable for agent's preferences over these plans to change.

A myopic agent will respond to these preference changes by deviating from the plan chosen *ex ante*. Sophisticated decision makers will however avoid such dynamic inconsistency by anticipating that the *ex post* ambiguity connected to each plan will be different from the *ex ante* ambiguity associated with them and by adjusting their choice of plan in the light of this. As is evident from Table 6 the *ex post* chances of winning induced by the plans R_0 and R_1 are the same, as are those induced by plans B_0 and B_1 . Consequentialism therefore requires the agent to prefer R_0 to B_0 iff they prefer R_1 to B_1 . If the agent is sophisticated therefore she will, on pain of a violation of Consequentialism, choose plan CR_0 over plan CB_0 iff she will choose plan CR_1 over plan CB_1 . Reduction then of course requires her to be ambiguity neutral in the static Ellsberg paradox, contrary to our claim that ambiguity aversion based on a preference for known chances over unknown ones is perfectly reasonable.

This impasse can be resolved by rejecting either Consequentialism or Reduction. In contrast to the Diamond case, our argument in this paper does not provide strong grounds for rejecting Consequentialism here. In particular, there are no enduring gains to be had by a sophisticated agent from choosing the less ambiguous alternative when she knows that she will acquire information later on that will affect how ambiguous these alternatives are. Indeed, for an agent to act resolutely in the sequential version of the Ellsberg paradox would amount to a refusal to take account of the change in their informational state that they anticipate, something that is hard to square with the idea that the agent is motivated to act on what she knows about the chances conferred by her alternatives.

In contrast too to the Diamond case, it is Reduction that seems least compelling as a condition of rational choice in the Ellsberg case. This is for the reasons already anticipated in the earlier discussion of preferences for the timing of the resolution of uncertainty. For as Hill (2020) argues, the epistemic situation of the agent facing the sequential problem is quite different from the one in the static problem because in the former (and not the latter) she knows that she will receive ambiguity relevant information before having to choose whether to bet on black or red. This information, as we have seen, does not dispel the ambiguity, but it does transform it. And the rational agent should surely in such situations take account of this fact – something that they cannot do in the static choice problem. Reduction must therefore be rejected.¹¹

¹¹ It is worth noting that this point does not depend on chances having intrinsic value. For instance an agent with rank-dependent preferences will anticipate, in the sequential game, that she will acquire information affecting the ranking outcomes on which her valuation of option depends; through the elimination of certain possibilities in particular. (We are grateful to an anonymous referee for this point.)

5 Risk aversion and Allais' paradox

In this penultimate section, before concluding, we consider whether the defence we have made of the dynamic implications of an aversion to unfairness and to ambiguity extends in a natural way to the treatment of the kind of aversion to risk (or regret) exhibited by many subjects in the Allais paradox.

Consider Table 7. Here subjects are offered a choice between on the one hand lotteries L1 and L2 and on the other hand lotteries L3 and L4. The events, e.g. '2–11', denote the numbers of the tickets which, if drawn, determine the prize of each lottery; for instance, a draw of any ticket numbered 2–11 results in a prize of 5 million dollars if L1 is chosen.

When faced with lotteries like these, many people prefer L2 to L1 and L3 to L4, as indeed predicted by Maurice Allais (1953) after whom this type of choice 'paradox' is named. One explanation, or rationalisation, for this pattern of preference is that people want to avoid the *regret* (Bell, 1982; Bradley & Stefánsson, 2017; Loomes & Sugden, 1982) that they would feel if they ended up with nothing ('0') when they could have guaranteed themselves a very good outcome ('\$1M'), which is a reason for preferring L2 over L1. No such reason is present when comparing L3 and L4, however, where the increased monetary expectation might tempt them to L3. An alternative explanation, or rationalisation, is that when making the first choice, the possibility of completely eliminating the risk of winning nothing gives reason to prefer L2 over L1 and carries greater weight than the possibility of merely reducing—albeit by the same magnitude (1/100)—the probability of winning nothing by choosing L4 over L3 (see, e.g., Buchak, 2014; Stefánsson & Bradley, 2019; Weirich, 1986). For the present purposes, we do not have to choose between these rationalisations of the 'Allais preference'. (In fact, as we shall see, our diagnosis is independent of which rationalisation one favours.) What matters is, first, that the preference violates Event Dominance, and, second, the dynamic implications it gives rise to.

The explanation for why the preference violates Event Dominance is very similar to the one we saw above for Ellsberg's Three-Colour Paradox (Table 5). Recall that Event Dominance requires subjects to be able to evaluate alternatives event-by-event. So, conditionally on the event that one of tickets 12–100 is drawn, subjects should be indifferent between L1 and L2 and also between L3 and L4, since in this event, the first pair results in the same outcome and so does the second pair. But, again, conditionally on the disjunction, ticket 1 or one of tickets 2–11 being drawn, L1 is the same alternative as L3 and L2 is the same alternative as L4. Therefore, one

Table 7 Allais paradox

	1	2–11	12–100
L1	0	\$5M	\$1M
L2	\$1M	\$1M	\$1M
L3	0	\$5M	0
L4	\$1M	\$1M	0

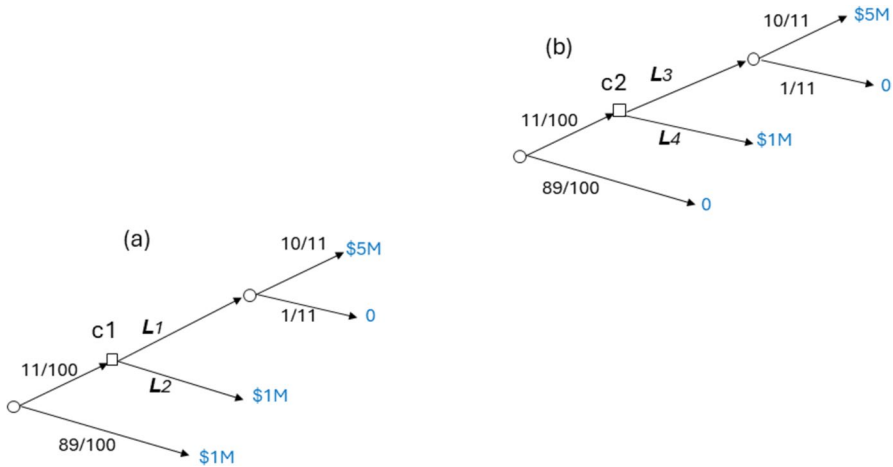


Fig. 4 Dynamic Allais paradox

should, by Event Dominance, prefer L2 to L1 just in case one prefers L4 to L3. So, either the preference for L2 over L1 or that for L3 to L4 violates Event Dominance.

To examine the dynamic implications of this violation of Event Dominance, we turn now to Fig. 4. Figure 4a can be seen as a dynamic version of the first choice in the Allais paradox, except that this time the agent knows whether one of tickets 12–100 was drawn before she makes the choice (at choice node c1). Figure 4b can analogously be seen as a dynamic version of the second choice in the Allais paradox where the agent makes the choice (at choice node c2) after learning whether one of tickets 12–100 was drawn.

By Reduction, an agent who prefers L2 to L1 and L3 to L4 in the static Allais paradox should, in the dynamic version, go ‘down’ at c1, thus choosing plan L2 rather than L1, but go ‘up’ at c2, thus choosing plan L3 rather than L4. And that might seem to be required by dynamic consistency: if one really prefers L2 to L1 in the first static decision problem, then it might appear that one should plan to go down at c1, and if one really prefers L3 to L4 in the second static decision problem, then one should plan to go up at c2. But that violates Consequentialism: the tree starting at choice node c1 is exactly the same as that which starts at c2. So, we see that someone with the Allais preference cannot be dynamically consistent while satisfying both Reduction and Consequentialism.

One potential response to this would be to give up Consequentialism by suggesting that the prospect of eliminating the ‘0’ outcome is evaluated differently at c1 than at c2.¹² This would allow someone with the Allais preference to be dynamically consistent. But this suggestion seems hard to square with either the motivation to eliminate future regret

¹² As previously discussed, this strategy could, formally speaking, be consistent with Consequentialism, if the outcomes are appropriately redescribed. But that would require describing outcomes in a way that builds a past chancy process into the description, which arguably goes against the spirit of Consequentialism.

or the motivation to eliminate risk. The fact that some chance-event happened in the past, whose outcome one now knows, does not change the fact that one can eliminate both this type of regret and risk by going down at both $c1$ and $c2$. So, if either regret aversion or risk aversion is to rationalise the Allais preference, then someone with this preference will have to prefer to go down at both $c1$ and $c2$.

Another potential response, which we favour, would be to give up Reduction. We saw before that both ambiguity averse and fairness seeking decision makers may care about when uncertainty is resolved. The same is true, we suggest, for both risk and regret averse decision makers. In particular, someone motivated to eliminate risk or future regret should plan to go down at both $c1$ and $c2$, *if* all uncertainty is resolved before they make the choice. That is a way, and indeed the only way, to eliminate both risk and future regret in these scenarios. Moreover, someone who is averse to risk or regret, in this sense, should *prefer* making the latter choice after the resolution of uncertainty rather than before the resolution of uncertainty (i.e., should prefer to choose between $L3$ and $L4$ rather than between $L3$ and $L4$).

This second response requires giving up on Reduction: although one prefers $L3$ to $L4$ in the static version of the Allais paradox, one prefers $L4$ over $L3$ in the dynamic version. That is, since one knows that the latter choice (if made) will be made in an information state where one can ensure winning the million dollars, one should even *ex ante* prefer $L4$ to $L3$ in this dynamic version. So, we suggest that a rational Allais preference has dynamic implications similar to a rational Ellsberg preference (violating Reduction), as opposed to a rational Diamond preference (which violates Consequentialism).

Moreover, we think that the view that Allais preferences are rationally permissible is no more undermined by dynamic considerations than are the views that Ellsberg and Diamond preferences are rationally permissible. *If* one can rationally permissibly be concerned with avoiding risk and future regret in the way that gives rise to the Allais preference, then one can rationally plan to go down at both $c1$ and $c2$ while displaying the Allais preference in a static decision (in violation of Reduction). Some of course think that one *cannot* be rationally concerned with avoiding risk or future regret in this way. But to vindicate that claim requires a more philosophically substantial argument than is offered by only observing the dynamic implications of such aversion to risk or future regret.

6 Concluding remarks

Let us summarise our findings and the lessons we draw from them. We have seen that people with the Allais, Ellsberg, and Diamond preferences can all avoid dynamic inconsistency, even though they violate the static separability assumptions of expected utility theory (in particular, Event Dominance). In fact, we have seen that if—as we have suggested—the Ellsberg and Diamond preferences are the result of the kind of attitudes to chances that support a concern with their distribution then we should expect people with these preferences to be dynamically consistent. Similarly for the Allais preference, where we saw that both of the traditional rationalisations of this preference lead to dynamic consistency.

However, we have also seen that the most plausible explanation for why people with these preferences avoid dynamic inconsistency is not the same across these three types of preferences. Those with the Diamond preference can avoid dynamic inconsistency because they *predictably* violate Consequentialism: what happened in the past matters to them (and they will plan accordingly). In contrast, those with the Ellsberg and Allais preferences can avoid dynamic inconsistency because the timing of the resolution of uncertainty *predictably* affects their preferences in a way that violates Reduction (and they will plan accordingly).

Philosophically, the most important lesson we want to draw from this is that considerations of dynamic consistency do not undermine the rationality of adopting attitudes to chances, both instrumental and intrinsic, that are ruled out by expected utility theory despite the fact that such attitudes give rise to violations of static separability assumptions such as Event and State Dominance.

Acknowledgements We are grateful to the participants at Issues in Dynamic Decision Theory at the University of Konstanz in July 2023 and to two anonymous reviewers for very helpful comments and suggestions. Stefánsson gratefully acknowledges funding from Riksbankens Jubileumsfond and the Knut and Alice Wallenberg Foundation.

Funding Open access funding provided by Stockholm University.

Data availability No data was used in this article.

Declarations

Conflict of interest No competing interests.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Allais, M. (1953). Le comportement de l'homme rationnel devant le risque: Critique des postulats et axiomes de l'école Américaine. *Econometrica*, 21(4), 503–546.
- Al-Najjar, N. I., & Weinstein, J. (2009). The ambiguity aversion literature: A critical assessment. *Economics and Philosophy*, 25(3), 249–284. <https://doi.org/10.1017/S026626710999023X>
- Bell, D. E. (1982). Regret in decision making under uncertainty. *Operations Research*, 30(5), 961–981. <https://doi.org/10.1287/opre.30.5.961>
- Bradley, R. (2016). Ellsberg's paradox and the value of chances. *Economics and Philosophy*, 32(2), 231–248. <https://doi.org/10.1017/S0266267115000358>
- Bradley, R. (2017). *Decision theory with a human face*. Cambridge University Press.
- Bradley, R., & Stefánsson, H. O. (2017). Counterfactual desirability. *The British Journal for the Philosophy of Science*, 68(2), 485–533. <https://doi.org/10.1093/bjps/axv023>
- Buchak, L. M. (2014). *Risk and rationality*. Oxford University Press.

- Cubitt, R. P. (1996). Rational dynamic choice and expected utility theory. *Oxford Economic Papers*, 48(1), 1–19. <https://doi.org/10.1093/oxfordjournals.oep.a028555>
- Diamond, P. A. (1967). Cardinal welfare, individualistic ethics, and interpersonal comparison of utility: Comment. *Journal of Political Economy*, 75(5), 765–766. <https://doi.org/10.1086/259353>
- Eichberger, J., Grant, S., & Kelsey, D. (2016). Randomization and dynamic consistency. *Economic Theory*, 62(3), 547–566. <https://doi.org/10.1007/s00199-015-0913-8>
- Ellsberg, D. (1961). Risk, ambiguity, and the Savage axioms. *The Quarterly Journal of Economics*, 75(4), 643–669. <https://doi.org/10.2307/1884324>
- Epstein, L. G., & Le Breton, M. (1993). Dynamically consistent beliefs must be bayesian. *Journal of Economic Theory*, 61(1), 1–22. <https://doi.org/10.1006/jeth.1993.1056>
- Fleurbaey, M. (2018). Welfare economics, risk and uncertainty. *Canadian Journal of Economics/revue Canadienne D'économie*, 51(1), 5–40. <https://doi.org/10.1111/caje.12314>
- Gauthier, D. (1997). Resolute choice and rational deliberation: A critique and a defense. *Noûs (Bloomington, Indiana)*, 31(1), 1–25. <https://doi.org/10.1111/0029-4624.00033>
- Hammond, P. J. (1976). Changing tastes and coherent dynamic choice. *The Review of Economic Studies*, 43(1), 159–173. <https://doi.org/10.2307/2296609>
- Hammond, P. J. (1988). Consequentialist foundations for expected utility. *Theory and Decision*, 25(1), 25–78. <https://doi.org/10.1007/BF00129168>
- Hammond, P. J. (1989). Consistent plans, consequentialism, and expected utility. *Econometrica*, 57(6), 1445–1449. <https://doi.org/10.2307/1913716>
- Harsanyi, J. C. (1955). Cardinal welfare, individualistic ethics, and interpersonal comparisons of utility. *The Journal of Political Economy*, 63(4), 309–321. <https://doi.org/10.1086/257678>
- Hill, B. (2020). Dynamic consistency and ambiguity: A reappraisal. *Games and Economic Behavior*, 120, 289–310. <https://doi.org/10.1016/j.geb.2019.12.012>
- Loomes, G., & Sugden, R. (1982). Regret theory: An alternative theory of rational choice under uncertainty. *The Economic Journal (London)*, 92(368), 805–824. <https://doi.org/10.2307/2232669>
- Machina, M. J. (1989). Dynamic consistency and non-expected utility models of choice under uncertainty. *Journal of Economic Literature*, 27(4), 1622–1668.
- McClennen, E. F. (1990). *Rationality and dynamic choice: Foundational explorations*. Cambridge University Press.
- Nissan-Rozen, I. (2019). The value of chance and the satisfaction of claims. *Journal of Philosophy*, 116(9), 469–493. <https://doi.org/10.5840/jphil2019116929>
- Rabinowicz, W. (1995). To have one's cake and eat it, too: Sequential choice and expected-utility violations. *Journal of Philosophy*, 92(11), 586–620. <https://doi.org/10.2307/2941089>
- Rabinowicz, W. (2020). Between sophistication and resolution—wise choice. In R. Chang & K. Sylvan (Eds.), *The Routledge handbook of practical reason* (pp. 526–540). Routledge.
- Savage, L. J. (1954). *The foundations of statistics*. Chapman & Hall.
- Siniscalchi, M. (2009). Two out of three ain't bad: A comment on “The ambiguity aversion literature: A critical assessment.” *Economics and Philosophy*, 25(3), 335–356. <https://doi.org/10.1017/S0266267109990277>
- Siniscalchi, M. (2011). Dynamic choice under ambiguity. *Theoretical Economics*, 6(3), 379–421. <https://doi.org/10.3982/TE571>
- Stefánsson, H. O., & Bradley, R. (2015). How valuable are chances? *Philosophy of Science*, 82(4), 602–625. <https://doi.org/10.1086/682915>
- Stefánsson, H. O., & Bradley, R. (2019). What is risk aversion? *The British Journal for the Philosophy of Science*, 70(1), 77–102. <https://doi.org/10.1093/bjps/axx035>
- Strotz, R. H. (1955). Myopia and inconsistency in dynamic utility maximization. *The Review of Economic Studies*, 23(3), 165–180. <https://doi.org/10.2307/2295722>
- Von Neumann, J., & Morgenstern, O. (1947). *Theory of games and economic behaviour* (2nd ed.). Princeton University Press.
- Weirich, P. (1986). Expected utility and risk. *The British Journal for the Philosophy of Science*, 37(4), 419–442. <https://doi.org/10.1093/bjps/37.4.419>