

Orii H. Stefansson and [Richard Bradley](#)
How valuable are chances?

**Article (Accepted version)
(Refereed)**

Original citation: Stefansson, H. Orii and Bradley, Richard (2015) *How valuable are chances?*
[Philosophy of Science](#) . ISSN 0031-8248

©2015 by the [Philosophy of Science Association](#)

This version available at: <http://eprints.lse.ac.uk/62780/>
Available in LSE Research Online: July 2015

LSE has developed LSE Research Online so that users may access research output of the School. Copyright © and Moral Rights for the papers on this site are retained by the individual authors and/or other copyright owners. Users may download and/or print one copy of any article(s) in LSE Research Online to facilitate their private study or for non-commercial research. You may not engage in further distribution of the material or use it for any profit-making activities or any commercial gain. You may freely distribute the URL (<http://eprints.lse.ac.uk>) of the LSE Research Online website.

This document is the author's final accepted version of the journal article. There may be differences between this version and the published version. You are advised to consult the publisher's version if you wish to cite from it.

How Valuable are Chances?

H. Orri Stefánsson and Richard Bradley[†] *

Abstract

Chance Neutrality is the thesis that, conditional on some proposition being true (or being false), its chance of being true should be a matter of practical indifference. The aim of this paper is to examine whether Chance Neutrality is a requirement of rationality. We prove that given Chance Neutrality, the Principal Principle entails a thesis called Linearity; the centrepiece of von Neumann and Morgenstern's expected utility theory. With this in mind, we argue that the Principal Principle is a requirement of practical rationality but that Linearity is not; and hence, that Chance Neutrality is not rationally required.

* To contact the authors, please write to H. Orri Stefánsson, Chair for Welfare Economics and Social Justice, Collège d'études mondiales, Paris, France; e-mail: hlynur.orri@gmail.com.

Richard Bradley, Department of Philosophy, Logic and Scientific Method, London School of Economics and Political Science, London, UK; e-mail: r.bradley@lse.ac.uk.

† We would like to thank the audience at Bristol University and Stockholm University where this paper was presented, as well as three referees for Philosophy of Science, for very helpful comments and suggestions. Stefánsson's work on this paper was supported by an AXA Research Grant (14-AXA-PDOC-222). Bradley's work on this paper was supported by grants from the Arts and Humanities Research Council (AH/J006033/1, AH/I003118/1).

1 Introduction

Suppose you have just learnt that you have won the lottery. Good news, we suppose. Now you learn that your chances of winning were just $1/10^{10}$. Does the fact that your prior chances of winning the lottery were so low, make it even better that you did win? In a similar vein, suppose your doctor tells you that you have been diagnosed with a lethal condition of a very rare kind that only afflicts one in a billion people. Is the desirability of what you have learnt any different from the desirability of simply learning that you have this particular condition, without any information about the chances? Or different from learning that you have a lethal condition that is quite common? Finally, consider the following scenario. After driving home from work, you realise that one of the wheels on your car is so loose that it was a matter of great luck that it didn't come off on the motorway, which would almost certainly have resulted in a terrible accident. Is this risk of accident in any way undesirable, after the fact, when there was no accident?

The standard answer to these three questions—embodied, for instance, in decision theory, risk analysis and orthodox economics—is 'No'. The lottery is valuable because of the prize you get from winning. Having a chance of winning is certainly valuable at the start, but once you know that you have won the lottery the desirability of that chance is 'neutralised'. Similarly, although it is clearly undesirable to learn that you have some chance of suffering from a lethal condition (or having an accident), once you learn whether you in fact have the condition (or had the accident), the chances become irrelevant.¹ On this orthodox view—which we

¹This is not entirely correct, as we point out in Section 3.3, since even on the orthodox view on chance, the chance of the condition might be practically relevant for instance if it tells you something about the probability that an advanced treatment of it has been developed. Similarly, the chance of winning the lottery might even on the orthodox view be practically

will term *Chance Instrumentalism*²—chances themselves are not carriers of value, but have value only insofar as they are instrumental to the outcomes with which they are associated.

The aim of this paper is to question this orthodox view on the value of chances and risks (which, as we are using the terms, are simply objective probabilities of gains or losses). In particular, we will examine whether it can be practically rational to take chances of good outcomes and risks of bad outcomes to play the role of ultimate goods in our practical and evaluative reasoning. Part of our motivation for exploring this thesis is the fact that seemingly rational people answer questions like those we posed above in a way that is inconsistent with the orthodoxy. For instance, many people feel that it would be worse to learn that they have a lethal condition that only afflicts one person in a billion than just learning that they have the condition, and that the former piece of news makes their condition much harder to accept.³

The paper proceeds as follows. In the next section we explain Chance Instrumentalism in relevant if it gives you evidence about how large the prize is. To avoid this problem, we formalise the orthodox view in Section 3.3 in terms of (maximally specific) possible worlds rather than (coarse grained) propositions.

²As far as we are aware, neither Chance Instrumentalism, nor its implication, Chance Neutrality, has been explicitly discussed before in the philosophical or decision-theoretic literature. However, these principles are typically *implicitly* assumed in decision theory and, in particular, applications of decision theory such as micro economics and risk analysis, where the outcomes of interest are usually non-modal ones (i.e., they concern only what actually happens).

³As we discuss in the next section (in the context of other examples), this type of example only does what we want it to do if we assume that the former piece of news is harder to accept than the latter because it is worse, rather than it being worse because it is harder to accept.

slightly more detail and show how it relates to a thesis we call *Chance Neutrality*, according to which the chance of an event should be a matter of indifference once we know whether the event took place or not. We discuss two examples that seem to suggest that Chance Neutrality is not a requirement of rationality, before proving (in Section 3) a theorem that we use to determine whether Chance Neutrality really is rationally required. The theorem says that given Chance Neutrality, David Lewis' (1980) *Principal Principle*, according to which a person who knows the objective chance of an event should set her credence (i.e., degrees of belief) accordingly, entails *Linearity*, the claim that the value of a lottery is equal to the sum of the values of the lottery's prizes discounted by their probabilities. Hence, given that the Principal Principle is (as we argue in Section 4) a requirement of rationality, Chance Neutrality cannot be rationally required unless Linearity is. But as we argue in Section 5, Linearity is not rationally required, and, therefore, neither is Chance Neutrality. Since Chance Neutrality is entailed by the instrumentalist view on chance, this gives strong grounds for rejecting the latter. Finally, in Section 6 we reflect on the relationship between our argument against Chance Neutrality and those of other critics of the orthodoxy who take issue with its treatment of risk aversion.

2 Chance Neutrality

Why do we care about chances? One obvious answer is that chances matter to us because the things that they are chances-of matter to us. One might care about the chances of cancer because one wants to avoid cancer, about the chances of rainfall because one prefers to stay indoors when it rains, and so on. This suggests that chances matter instrumentally. By improving (or reducing) our chances we make it more (or less) likely that some state of affairs

that we care about will be realised. Moreover, learning that the chances of some state of affairs that we care about are greater (or lesser) than we expected, gives us evidential reason to increase (or reduce) our expectation of that state of affairs being realised. The value of a chance that derives from the causal and/or evidential relations it stands in with respect to what it is a chance-of is what we call its *instrumental* value.

On the Chance Instrumentalist view, chances have *only* instrumental value and do not matter ‘intrinsically’⁴. If a chance of some good has only instrumental value in virtue of being a chance of that good, then it has no value independently of the realisation or otherwise of the good in question. Or to put it slightly differently, if the value of a chance that *X* is exhausted by its instrumental role in making it true that *X*, then once it is settled whether or not *X* is true, there can be no further value associated with the chance of *X*. For if it is the case that *X* then the chance is redundant and if it is not the case that *X* then the chance is inert. (These are just two sides of the same coin: A chance of some magnitude α that *X* is a chance of magnitude $1 - \alpha$ that not *X*.) So Chance Instrumentalism implies that chances are neutral in the sense that, conditional on *X*, it does not matter what the chances of *X* are.

The instrumental value of chance is often paramount in our deliberations, which is why Chance Instrumentalism may seem attractive. But the view is incorrect: Chances can have non-instrumental value too. Consider the following two examples:

⁴By saying that chances (and risks) can matter *intrinsically*, we do not mean that they are or can be objectively or universally valuable. Rather, we mean that they can be intrinsically valuable *to some person* in that they play the role of ultimate goods in his or her evaluative reasoning. Nor do we mean that the chance’s value is in no way related to the value of what it is a chance-of; but simply that the value of the chance is not exhausted by its instrumental relationship to what it is a chance-of.

Example 1 (Taking Risks). *Ann is an experienced mountain climber. She does not directly seek death or severe injury, but nonetheless tackles climbs that puts her at risk of both. She does this because part of the value of climbing, according to her, stems from the confrontation of risk. Ann is not completely foolhardy and will not undertake climbs if the chances of death or injury are too high. But for her the activity is of little worth if there is no associated chance of death or injury, even though death and injury are outcomes that she would strongly prefer to avoid. Indeed for her there is an optimal region of risk, where the chances of death or injury are high enough to require courage of the climber, but not so high as to make the activity foolish.*

Example 2 (Equal Treatment). *Bob's father has been given a watch as a present, but since he already has one, he decides it should go to one of his children. Bob doesn't particularly want or need the watch, but nonetheless is outraged when his father announces his intention to give the watch to Bob's sibling. Seeing Bob's reaction, his father decides to toss a coin to settle who gets the watch, which leaves Bob satisfied.*

These two examples are very different, but share the common feature of a protagonist who violates Chance Instrumentalism, since s/he finds the chance (or risk) of some outcome desirable even though s/he does not desire the outcome itself. Ann does not value the risk of death and injury associated with climbing as means to these highly undesirable outcomes. Rather, she values them as intrinsic features of an activity that she considers worthy. We do not have to share Ann's penchant for danger to recognise this type of value. It is often the case that the difficulties associated with some activity or project (generically the risk of failure) are part of what makes them worthy as activities or projects. This can be as true for extreme crocheting as it is for mountaineering. And in these cases, there is an important difference between achieving some outcome when the chances of failing were high and achieving it when suc-

cess was pretty much guaranteed; in violation of Chance Neutrality. Bob, on the other hand, does not value the chance of a watch because of its instrumental relation to the outcome of receiving the watch, since he does not even want the watch. Nor does he view the chances neutrally: For him it is far worse not to get the watch when he never had a chance of getting it than to not get it when he had a fair shot of doing so. So both Ann and Bob violate both Chance Instrumentalism and Chance Neutrality.

There are two obvious responses to these examples available to the Chance Instrumentalist. Firstly, she can grant that these examples show that the value of the chance of some good G is not always exhausted by its instrumental relation to G , but argue that the value of the chance nonetheless derives from its instrumental relation to *some* good. Take our mountain climber Ann. Ann does not seek death, but she does seek goods like the mental states of excitement, exhilaration, or sense of pride, that come with confronting risk; and social goods like the respect of other mountaineers and the admiration of those who are unable to control their fear of death. Similarly, although Bob does not want the watch, he does seek the good of fair treatment and equal consideration by his parents, and being given a chance of getting the watch symbolises for him his equal status amongst the children. More generally, the Chance Instrumentalist will claim, for any chance proposition, there is some 'ultimate' good G^* (which could, of course, be a bundle of goods like excitement, respect, health and so on) such that the chance of G , given that G^* , is without value.

The instrumentalist is right in pointing to the existence of these other goods. But she has the explanation the wrong way around. Ann does, let us suppose, get excitement and pride from confronting risk, but this is not what is valuable to her about it. The excitement is a *by-product* of the activity, not what makes it worthwhile in the first place. Even if she could get as much excitement from a simulated climb as from the real thing, she would not judge

climbing ersatz mountains to be nearly as valuable as conquering real ones. The pride that she feels from controlling fear depends even more directly on the value that she attaches to facing (real) risks, and the same can be said about the respect that she receives from others. Ann values this respect just in case she believes it deserved. And those who respect her must, for the respect to be valuable to her, agree with her in valuing the confrontation of risk independently of the feelings it generates.

The second example requires a somewhat different response since Bob's desire to be treated fairly (as he sees it) is not a by-product of the value he attaches to the chance of getting the watch. Rather it is *intrinsic* to him being treated fairly (as he sees it) that he should have this chance. We could not, in other words, ensure that he gets the desired good of fair treatment without giving him the chance of getting the watch; for Bob being given this chance is part of being treated equally to his sibling. So in this case too, it is wrong to think of the value of chances as deriving from an instrumental relation to some good.

A second response would be to follow John Broome's (1991) re-individuation strategy as a way of saving Chance Instrumentalism and Chance Neutrality from our criticism.⁵ Broome points out that principles of rationality can often be saved from apparent counterexamples by individuating the relevant outcomes more finely, in particular to make sure that different outcomes get described differently. For instance, in *Taking Risks*, we can, formally speaking, save Chance Instrumentalism by describing the results of Ann's activities as conquering a worthwhile or worthless mountain depending on whether the activity is risky or not; whereas in *Equal Treatment*, we might make Bob's attitude formally consistent with Chance Instrumentalism by describing two situations where he doesn't receive the watch as 'No watch and fair treatment'/'No watch and unfair treatment'. And similarly for other apparent violations

⁵We thank a referee for *Philosophy of Science* for pressing us on this issue.

of Chance Instrumentalism/Neutrality.

One problem with this way of saving Chance Instrumentalism is that there is a clear sense in which the principle is *only* formally satisfied, but is not satisfied in a substantive sense. The reason one mountain is a worthwhile challenge but another is not, according to Ann, is that the former is risky but the latter is not; and the reason one situation where Bob gets no watch is fair but another is not, on his view, is that in the former situation he had a chance of receiving the watch but in the latter he had not. So the whole reason these re-individuations allow us to formally save Chance Instrumentalism is that they build risks and chances into the description of ultimate outcomes, which clearly goes against the intuition behind Chance Neutrality and, by implication, against Chance Instrumentalism (since the latter entails the former).

Moreover, the re-individuation strategy, if used indiscriminately, will trivialise decision theory and other theories of rationality (as Broome 1999 himself points out) since it is *always* possible to re-describe outcomes such that an apparent violation of a rationality principle is made consistent with it. Our solution, however, can be seen as a principled application of the re-individuation strategy. Building on the observation that agents sometimes legitimately take risks and chances to be part of the outcomes they ultimately care about, we are suggesting dropping Chance Neutrality and making objective probability distributions explicitly part of the outcome space. Hence, we make room for attitudes that violate Chance Instrumentalism/Neutrality without running the risk of trivialising decision theory.⁶

⁶See Bradley and Stefánsson (2015) for a similar criticism of Broome's re-individuation strategy, but in a different context.

3 Three Rationality Theses About Chance

In this section we examine the relationship between three claims about rational attitudes to chances that are the focus of this paper: Chance Neutrality, the Principal Principle and Linearity. First, we sketch a possible-worlds semantics for propositions about chance that is maximally neutral with regard to the metaphysical status of chances and in particular to the question of whether chance facts are reducible to non-chance facts. This enables us to give a formal statement of the three claims and then to prove our main theorem: That in the presence of the standard axioms of probability and desirability, Chance Neutrality and the Principal Principle together entail Linearity. In Section 4 we argue that the Principal Principle is a requirement of rationality, but in Section 5 that Linearity is not. Given our theorem, this entails that Chance Neutrality is not a requirement of rationality.

3.1 Semantics

First, some notation. Let W be the set of base worlds, $\{w_1, w_2, \dots, w_n\}$, and $\wp(W) = \Omega$ the set of all factual propositions $\{A, B, C, \dots\}$, each proposition being a set of worlds in W . Then let $\Pi = \{ch_1, ch_2, \dots\}$ be the set of all logically possible probability distributions over Ω , and $\Delta = \wp(\Pi) = \{Ch_1, Ch_2, \dots\}$ be the set of all chance propositions, i.e., propositions about the chances of truth of propositions in Ω . Intuitively, the probability distributions serve as truth-makers for claims about chances, i.e., as the basic chance facts. For instance, the proposition that X has a chance greater than one-half is given by, or made true by, the set of chance functions ch such that $ch(X) > 0.5$.

Next, we define a set of extended worlds, $\mathcal{W} = W \times \Pi$, the set of all pairs $\langle w, ch \rangle$, such that $w \in W$ and $ch \in \Pi$. Intuitively, an extended world gives a complete description of both

the base (non-modal) facts and the chance facts, so that \mathcal{W} serves, in this framework, as the set of elementary possibilities. In the interests of metaphysical neutrality, we will remain agnostic about what possible worlds and probability distributions are co-possible, though clearly the framework is compatible with all kinds of restrictions. Strict reductionist theories require each base world to be paired with a unique probability distribution; for instance, on the finite frequency interpretation of chance, base facts about the proportions of outcomes that belong to a certain class will determine the chances of these outcomes. Strict anti-reductionist views, on the other hand, allow that any pair of base worlds and chance worlds are co-possible. All kinds of intermediate positions are also conceivable of course.

The set of extended worlds \mathcal{W} induces a corresponding set of extended propositions, $\Gamma = \wp(\mathcal{W}) = \{\mathcal{A}, \mathcal{B}, \mathcal{C}, \dots\}$, the set of all sets of $\langle w, ch \rangle$ pairs. Extended propositions serve as the possible contents of agents' attitudes. In the interests of simplicity, for any $A \in \Omega$ and $Ch \in \Delta$ we denote the element $A \times \Pi$ of Γ by A and the element $W \times Ch$ of Γ by Ch , thus allowing us to denote element $A \times Ch$ by $A \cap Ch$, this being the set of extended worlds $\langle w, ch \rangle$ such that $w \in A$ and $ch \in Ch$.

In what follows, we use $Ch(B) = \alpha$ to denote the proposition that is true just in case $ch(B) = \alpha$; that is, the proposition that the chance of B is α (and the chance of $\neg B$ is $1 - \alpha$). And we use $\langle Ch(A_i) = \alpha_i \rangle$ to denote the conjunction of the propositions $Ch(A_1) = \alpha_1$, $Ch(A_2) = \alpha_2$, ..., $Ch(A_n) = \alpha_n$, where $\{A_i\}_{i=1}^n$ is an n -fold partition of Ω and the α_i are real numbers such that $\alpha_i \in [0, 1]$ and $\sum_{i=1}^n \alpha_i = 1$. Each such proposition $\langle Ch(A_i) = \alpha_i \rangle$ expresses the chances of realising each of the 'prizes' represented by the base propositions A_i . The $\langle Ch(A_i) = \alpha_i \rangle$ thus serve as the propositional equivalents, in this framework, of the lotteries that appear in the orthodox theory of rational choice under objective uncertainty; in particular, in von Neumann and Morgenstern's (1953) expected utility theory. Let \mathcal{L} be the

set of all such lottery propositions.

The final step is to characterise rational attitudes to these propositions. To do so we draw on the version of decision theory developed by Richard Jeffrey (1983). It has two salient advantages in this context. First, because Jeffrey's measures of rational degrees of belief and desire are defined on the same space—the set of propositions—it is relatively simple to examine the relationship between rationality principles about practical and epistemic attitudes to chances. Secondly, his basic theory depends on rather weak rationality assumptions relative to its main alternatives, i.e., compared to those based on Leonard Savage's (1972) decision theory. (In particular, Jeffrey's theory contains a much weaker separability axiom than its alternatives.) This enables us to evaluate proposals for constraints on rational attitudes to chances without taking too much for granted. It is worth noting that the differences between Jeffrey's (evidential) decision theory and the newer 'causal' decision theory do not matter here, since we will not be discussing the choice-worthiness of actions. So any inadequacies of Jeffrey's theory in this regard are irrelevant to the present discussion.

In Jeffrey's theory the degrees of belief of a rational agent are measured by a subjective probability (or credence) function, P , on a Boolean algebra of propositions, satisfying the standard probability axioms. Her degrees of desire are measured by a corresponding desirability function, V , defined on the same algebra but with the logically contradictory proposition removed and satisfying:

Desirability: If $A = \bigcup(A_i)$, and if for any i and j , $A_i \cap A_j = \emptyset$, then:

$$V(A) = \sum_i V(A_i) \cdot P(A_i|A)$$

We will take these two measures, P and V , to be defined on the entire Boolean algebra of

propositions based on the set Γ .

3.2 The Chance Theorem

We are now in a position to state the three theses of interest to us. The first is a version of David Lewis' (1980) famous Principal Principle. The principle says that you should set your degrees of belief in the truth of a proposition, conditional on its objective chance being some numerical value, to that numerical value. This entails, for instance, that if you believe that a coin has 0.5 chance of coming up heads, then you should believe to degree 0.5 that it will do so. And, more generally, that your degree of belief in the coin coming up heads should equal its expected chance of landing heads, relative to your degrees of belief in the various relevant chance hypotheses. More formally:

Thesis 1 (Principal Principle). *For any $A \in \Omega$ and $\langle Ch(A_i) = \alpha_i \rangle \in \mathcal{L}$, if $P(\langle Ch(A_i) = \alpha_i \rangle) \neq 0$, then:*

$$P(A_i | \langle Ch(A_i) = \alpha_i \rangle) = \alpha_i$$

This statement of the Principal Principle differs from Lewis's in a number of ways. Firstly, our formulation is in some ways more simplistic than Lewis'. For instance, it ignores the fact that the principle is not supposed to hold for propositions for which the agent in question has *inadmissible information*, this being information whose relevance to the truth of the proposition does not come entirely from its relevance to its chance of truth. It also ignores the fact that chances evolve over time, which means that for a chance proposition to have a fixed truth value, it must be time-indexed. In the next section we explain why our argument still holds for less simplistic versions of the Principal Principle.

Secondly, while our version of the Principal Principle makes reference to an arbitrary par-

tition of propositions $\{A_i\}$ with respect to which chances are specified, Lewis's is restricted to a two-element partition. Our principle obviously implies Lewis's for, given a partition $\{A, \neg A\}$, it implies that:

$$P(A \mid Ch(A) = \alpha) = \alpha$$

And it seems evident that our version is no less plausible than Lewis'.⁷

The second thesis to be considered is the claim that chances are neutral in a practical or desirabilistic sense; for example that the desirability of learning that you have a very rare lethal condition is just the desirability of learning that you have the condition. Formally:

Thesis 2 (Chance Neutrality). *For any $A \in \{A_i\}$ and $\langle Ch(A_i) = \alpha_i \rangle \in \mathcal{L}$:*

$$V(\langle Ch(A_i) = \alpha_i \rangle \cap A_i) = V(A_i)$$

The zero point in a Jeffrey-desirability measure is typically reserved for propositions, such as the tautology \top , whose truth is neither desirable nor undesirable (see e.g. Jeffrey 1983, Section 5.6). In this case the standard definition of conditional desirability (Bradley 1999) entails that $V(\langle Ch(A_i) = \alpha_i \rangle \mid A_i) = V(\langle Ch(A_i) = \alpha_i \rangle \cap A_i) - V(A_i)$. Then Chance Neutrality

⁷In fact, Thesis 1 can be derive from Lewis' principle if we make the natural assumption that when one has already conditioned a subjective probability function on a proposition that says that the chance of A_i is α_i , then conditioning the function on another proposition that says that the chance of A_j is α_j , where $A_i \cap A_j = \emptyset$, won't make a difference to the function's assignment of probability to A_i (provided that the two chance propositions are consistent). In other words, given this assumption, $P(A_i \mid Ch(A_i) = \alpha_i \cap Ch(A_j) = \alpha_j) = P(A_i \mid Ch(A_i) = \alpha_i)$; which, given Lewis' Principal Principle, is equal to α_i .

implies that:

$$V(\langle Ch(A_i) = \alpha_i \rangle | A_i) = 0$$

In other words, given the truth of A_i , the desirability of a chance distribution over a partition of which A_i is a member is zero. This implication of Chance Neutrality is perhaps even closer to the intuitive idea of the neutrality of chances than Thesis 2 is, and could serve as an alternative formalisation of it in contexts in which the zero-normalisation can be assumed.

The final thesis of interest, Linearity, states that the value of a lottery proposition is a chance-weighted average of the values of its outcomes (or ‘prizes’):

Thesis 3 (Linearity). *For any $A \in \Omega$ and $\alpha_i \in [0, 1]$ such that $\sum_i \alpha_i = 1$:*

$$V(\langle Ch(A_i) = \alpha_i \rangle) = \sum_i V(A_i) \cdot \alpha_i$$

Linearity can be thought of as a transcription of von Neumann and Morgenstern’s expected utility theory into this propositional framework. It says that, for instance, the desirability of getting 10 if a coin comes up heads but 0 if it comes up tails, is equal to the desirability of 10 times the probability of heads plus the desirability of 0 times the probability of tails.

The three theses we have been discussing are independent of each other and none are entailed by Jeffrey’s basic framework. But as we now prove, within Jeffrey’s framework, Chance Neutrality together with the Principal Principle implies Linearity.

Theorem 1. *Assume Chance Neutrality. Then the Principal Principle entails Linearity.*

Proof. Since $\bigcup\{A_i \cap \langle Ch(A_i) = \alpha_i \rangle\}$ is a partition of $\langle Ch(A_i) = \alpha_i \rangle$, we have:

$$V(\langle Ch(A_i) = \alpha_i \rangle) = \sum_i V(A_i \cap \langle Ch(A_i) = \alpha_i \rangle) \cdot P(A_i | \langle Ch(A_i) = \alpha_i \rangle) \quad (1)$$

By Chance Neutrality, (1) becomes:

$$V(\langle Ch(A_i) = \alpha_i \rangle) = \sum_i V(A_i) \cdot P(A_i | \langle Ch(A_i) = \alpha_i \rangle) \quad (2)$$

Now assume the Principal Principle. Then (2) becomes:

$$V(\langle Ch(A_i) = \alpha_i \rangle) = \sum_i V(A_i) \cdot \alpha_i \quad (3)$$

In other words, Chance Neutrality and the Principal Principle entail Linearity. \square

3.3 Objections and Refinements

There are two types of objections we need to consider. First, as noted before, the version of the Principal Principle stated above is arguably too simplistic. As Lewis (1994) himself recognised, this version of the principle leads to contradiction if we allow for the possibility that a chance function ch assigns positive chance to an ‘undermining future’; that is, to a sequence of events such that if they all take place, then the chances are different from what ch says they are. Let F be a proposition that undermines ch in the aforementioned sense, and suppose that $ch(F) > 0$. Now let Ch be the proposition that the chances are what ch says they are. By Lewis’ original Principal Principle, $P(F | Ch) > 0$. But since F undermines ch , we have $F \supset \neg Ch$, so $Ch \supset \neg F$, which means that $P(F | Ch) = 0$, in contradiction to what the Principal Principle entails since F has a positive chance according to ch .

Various popular theories of chance, according to which chances supervene on what actually occurs, entail that undermining futures can have a positive chance. In particular, Lewis’ own Best System theory of chance entails this possibility. Hence, Lewis admitted that there

was a need to amend his own principle. The most natural way to do so is as follows:

Thesis 4 (New Principle). *For any $A \in \Omega$ and $\langle Ch(A_i) = \alpha_i \rangle \in \mathcal{L}$, if $P(\langle Ch(A_i) = \alpha_i \rangle) \neq 0$, then:*

$$P(A_i | \langle Ch(A_i) = \alpha_i \rangle) = ch(A_i | Ch)$$

For then if F undermines ch , the New Principle entails that $P(F | \langle Ch(A_i) = \alpha_i \rangle) = ch(F | Ch) = 0$, which means that we don't reach the contradiction we reached above.⁸

Replacing the Principal Principle with the New Principle makes no difference to our argument however as we can prove much same theorem as above. For notice that the New Principle entails that $P(A_i | \langle Ch(A_i) = \alpha_i \rangle) = ch(A_i | Ch) = \alpha_i$,⁹ which means that a simple variant of the proof of Theorem 1 shows that Chance Neutrality and the New Principle entail Linearity.

Another reason why our original statement of the Principal Principle is too simplistic, is that it ignores the fact that if one knows whether or not A is true, then that knowledge should 'trump' one's knowledge of A 's chance. To take a typical but unrealistic example: Suppose a perfectly reliable oracle informs you that a coin that is about to be tossed has 0.5 chance of coming up heads and actually will come up heads. To what degree should you then believe that the coin comes up heads? If you really think that the oracle is perfectly reliable, it seems you should ignore what you know about the chances, and fully believe that the coin will come up heads.

⁸The New Principle was independently suggested by Lewis (1994) and Ned Hall (1994) as a correction of the Principal Principle.

⁹Recall that Ch is true just in case the chances are what ch says they are. Hence, for any F that undermines ch , by conditioning ch on Ch we exclude F from the domain of ch ; and, in particular, make sure that ch is correct about A_i .

Information of the above kind is what Lewis called ‘inadmissible’. Although he did not give a precise definition of inadmissibility, he suggested that evidence is inadmissible whenever it affects your beliefs about the outcomes of events or the truth of propositions without affecting your beliefs about the chances of these outcomes or propositions (Lewis, 1980, 265). And as Lewis rightly pointed out, the Principal Principle only holds when an agent does not have evidence that is inadmissible with respect to the proposition that is being evaluated.¹⁰

Now it follows from the above observation that what is entailed by the Principal Principle and Chance Neutrality is Linearity limited to propositions for which the agent in question does not have inadmissible information. But on reflection, it should be evident that even the proponent of Linearity should in general take the principle to be limited to such propositions. Recall that Linearity states that the value of a lottery is a chance-weighted average of the values of its outcomes. Suppose lottery L has only two possible outcomes, O_1 and O_2 . Now the oracle tells you that the two outcomes have the same chance and that the lottery will actually result in O_2 . By Chance Neutrality, the chance of O_2 (and O_1) then becomes irrelevant from a practical point of view. Moreover, since you know that the lottery won’t result in O_1 , it seems your evaluation of the lottery should no longer be affected by the instrumental value it previously had in virtue of possibly realising O_1 . So the value of the lottery is certainly no longer a chance-weighted average of the values of O_1 and O_2 . Thus we have an intuitive explanation of why Linearity should be limited to a situation where the agent has no inadmissible information about the propositions in question.

Finally, the above statement of the Principal Principle ignores the fact that chances seem

¹⁰It should be noted, however, that Hall (1994), amongst others, argues that unlike the Principal Principle, the New Principle does not need an admissibility clause.

to evolve over time. For instance, remembering in the last minute to put an extra sweater on before leaving the house reduces the chance that one catches a cold. Moreover, once a chancy event has taken place (or not), its chance becomes 1 (or 0) according to most theories of chance. (For a notable exception, see Hoefer 2007.) If chances change over time, chance propositions must, strictly speaking, be time-indexed if they are to have a fixed truth-value. But then all the rationality principles that we have been discussing should be reformulated in the same way, by time-indexing the chance propositions that occur in their statements. Hence, our proof of the above theorem still goes through after we have made the appropriate changes. But to keep the notation as simple as possible, we will continue without the time-indexing.¹¹

The second type of objection we need to deal with has to do with our formalisation of the intuitive idea of chance neutrality. One might argue that even if chances only have instrumental value, Chance Neutrality as formulated above should be rejected, since the desirability of a chance of *A* is not exhausted by the value it has in virtue of being instrumental in making it true that *A*. Learning that the chance of *A* is high might affect our beliefs about the chance of some other proposition *B* and so have evidential value. For example, even if I know that I have some disease, learning what were my chances of contracting it might be important if it tells me something about my susceptibility to disease more generally or about the availability of treatments for it. (For example, the less common the disease, the more likely it perhaps is that there is something seriously wrong with my immune system, and the less likely it is that

¹¹The idea that chances evolve over time might give rise to worries about the diachronic consistency of an agent who violates Chance Neutrality (we thank a referee for *Philosophy of Science* for raising this issue). But as we point out later, giving up Chance Neutrality does not imply giving up the Independence condition that ensures dynamic consistency.

an advanced treatment has been developed for the disease.)

To avoid this objection, the Chance Neutrality thesis could be stated in terms of possible (base) worlds rather than propositions, so that it says that once we know which base world is actual, chances should not matter. More formally, let $W_i = \{w\} \times \Delta$ be the proposition that world w_i is the case. We will call W_i a *world proposition* and any proposition $\langle Ch(W_i) = \alpha_i \rangle$ expressing the chances of the world propositions W_i a *lottery proposition*. The new thesis can then be stated as:

Thesis 5 (Chance Neutrality*). *For any lottery proposition $\langle Ch(W_i) = \alpha_i \rangle$ and world proposition W_j :*

$$V(\langle Ch(W_i) = \alpha_i \rangle \cap W_j) = V(W_j)$$

Chance Neutrality* doesn't imply Linearity as originally formulated, but together with the Principal Principle it does imply a version that is similarly restricted to worlds. Formally:

Thesis 6 (Linearity*). *For any lottery proposition $\langle Ch(W_i) = \alpha_i \rangle$ with $\alpha_i \in [0, 1]$ such that $\sum_{i=1}^n \alpha_i = 1$:*

$$V(\langle Ch(W_i) = \alpha_i \rangle) = \sum_i \alpha_i \cdot V(W_i)$$

Linearity* is even closer than Linearity to being a statement of von Neumann and Morgenstern's expected utility theory within the Jeffrey framework, since the former lacks the partition-independent feature of the latter. To bring out the parallel, note that the subset \mathcal{L} of Γ consisting of the lottery propositions $\langle Ch(W_i) = \alpha_i \rangle$ forms a mixture space that is closed under linear combinations, with the linear combination of lottery propositions $\langle Ch(W_i) = \alpha_i \rangle$ and $\langle Ch(W_i) = \beta_i \rangle$, with respect to a real number $\lambda \in [0, 1]$, being defined by:

$$\lambda \langle Ch(W_i) = \alpha_i \rangle + (1 - \lambda) \langle Ch(W_i) = \beta_i \rangle := \langle Ch(W_i) = \lambda \alpha_i + (1 - \lambda) \beta_i \rangle$$

Moreover, note that $\langle Ch(W_i) = \lambda\alpha_i + (1 - \lambda)\beta_i \rangle$ is itself a lottery proposition, ensuring closure. Linearity* then implies that preferences on the lottery propositions satisfy the von Neumann and Morgenstern Independence axiom, which can be rendered here as the requirement that for all $L, L', M \in \mathcal{L}$ and $\lambda \in [0, 1]$:

Independence Axiom. $\lambda L + (1 - \lambda)M \succeq \lambda L' + (1 - \lambda)M \Leftrightarrow L \succeq L'$

Chance Neutrality* and Independence are conceptually distinct ideas. In particular, an agent can perfectly well view the desirability of the chances of two different goods as independent of each other, but value the chances of both goods intrinsically. But, as we prove below, in the presence of the Principal Principle, Chance Neutrality* implies Linearity* and hence Independence. The other direction does not hold however: Linearity* is not sufficient to derive the Principal Principle or Chance Neutrality*, even in the presence of the other condition. Hence Independence does not imply Chance Neutrality* even given the Principal Principle.

Nor, it should be noted, does Independence alone imply Linearity or Linearity*. This is important because we will argue against the latter conditions, and hence for the rejection of Chance Neutrality. But these arguments are not in themselves arguments for rejecting Independence. The Independence axiom is quite controversial as a principle of rational preference, with proponents arguing that it is a condition of dynamic consistency and opponents pointing to putatively permissible violations of it in the Allais paradox and elsewhere. But it is Chance Neutrality that is our target, not Independence.

Theorem 2. *Given the Principal Principle, Chance Neutrality* implies Linearity*.*

Proof. By the Desirability axiom:

$$V(\langle Ch(W_i) = \alpha_i \rangle) = \sum_i V(\langle Ch(W_i) = \alpha_i \rangle \cap W_i) \cdot P(W_i \mid \langle Ch(W_i) = \alpha_i \rangle) \quad (4)$$

But the Principal Principle entails that $P(W_i \mid \langle Ch(W_i) = \alpha_i \rangle) = \alpha_i$ and Chance Neutrality* that $V(\langle Ch(W_i) = \alpha_i \rangle \cap W_i) = V(W_i)$. So adding these two principles to (5) gives us:

$$V(\langle Ch(W_i) = \alpha_i \rangle) = \sum_i \alpha_i \cdot V(W_i) \quad (5)$$

In other words, the Principal Principle and Chance Neutrality* entail Linearity*. □

4 The Principal Principle and Practical Rationality

David Lewis (1980) thought the Principal Principle (PP) “captures all we know about chance” (266). Although not everyone agrees with him that the principle captures *everything* we know about chance, most concur that it captures the essence of how our degrees of belief, or credence, should relate to our knowledge of (or beliefs about) chances. In other words, most people agree that some version of the principle is a requirement of theoretical rationality.¹²

It is sufficient for our argument that the Principal Principle is a norm of theoretical rationality. For if people ought to align their credences with their beliefs about chance, but are not rationally required to satisfy Linearity, as we will argue, then they cannot be required (by rationality) to satisfy Chance Neutrality. Nevertheless, in this section we will seek to answer the question: Is the PP a requirement of *practical* rationality? Answering this question is independently interesting, we think, but particularly so in the present context, since if Linearity and Chance Neutrality are rationality requirements, then they must be requirements of practi-

¹²In addition to seeming intuitively very plausible as a principle of theoretical rationality, the epistemic importance of the Principal Principle has been justified on the grounds that it maximises epistemic utility (or accuracy). See in particular (Pettigrew 2012, 2013).

cal rationality, since they are principles about rational desire rather than belief.¹³

The best way to evaluate the Principal Principle as a norm of practical rationality, is to envision a situation where we are choosing between ‘lotteries’ with known (and fixed) chances (i.e., objective probabilities) for prizes. For this is exactly the type of situation where, first, knowledge of chances are practically important, and, second, the Principal Principle has something to say. (In contrast, the Principal Principle is silent in situations where one does not know the chances, and it has no practical relevance in situations where the chances are not relevant for decision-making.) But it should be emphasised that ‘lotteries’ is a technical term, that can stand for any prospect that gives rise to a known chance distribution over possible outcomes. ‘Chances’ can also be interpreted in different ways, the only requirement being that they play the role chances do in the Principal Principle.

Would it be practically irrational to violate the Principal Principle in a situation like the one just described? It might be tempting to argue that since one is guaranteed to lose money in the long run if one violates the PP when choosing between lotteries with known chances, adherence to the principle is a requirement of practical rationality. The problem with this argument is that the ‘long run’ is very long indeed. Someone who violates the principle, in a situation like the one we are imagining, is only *guaranteed* to win less than a person who satisfies the principle if the number of lotteries they partake in is infinite. It is true, of course, that as the number of lotteries increases, it becomes more and more likely that the person who satisfies the principle wins more than the person who does not. But the important question is:

¹³It might be worth emphasising that the aim here is not to give a complete argument for the Principal Principle as a requirement of rationality. We will take it for granted that the principle (in some form) is a requirement of theoretical rationality, and try to show that it is no less plausible as a requirement of practical rationality.

Why should practical rationality require that someone who only partakes in a single game of chance satisfies the Principal Principle?

A stronger reason to take the Principal Principle to be a requirement of practical rationality, is the fact that given one additional assumption, a person satisfies the Principal Principle just in case she satisfies an objective version of the well-known *Stochastic Dominance* principle (SD) in the situation we are imagining. The objective Stochastic Dominance principle says that if for any outcome some lottery L_1 has at least as high chance as lottery L_2 of resulting in that or a better outcome, then one should not prefer L_2 to L_1 . Expected utility maximisers satisfy a subjective version of this principle, where ‘chance’ has been replaced with ‘credence’, and so do those who maximise risk-weighted expected utility, as formulated by Lara Buchak (2013), and those who maximise Jeffrey’s (1983) desirability.

In a situation like the one we are imagining, where the decision-maker knows the chances, it would certainly seem to be practically irrational to violate objective Stochastic Dominance. Suppose lotteries L_1 and L_2 either result in a prize of \$100 or nothing, and that Alice really wants the money, which is the only good she considers relevant when choosing between these lotteries. Moreover, she knows that lottery L_1 has an objective 0.3 chance of yielding the prize, but lottery L_2 gives a 0.2 chance of winning the prize. (We are assuming that playing the lotteries is cost-free.) Then surely it would be practically irrational of Alice to choose lottery L_2 over L_1 .

Someone might object that Ann the mountain climber violates objective Stochastic Dominance, despite being (we claimed) rational: She prefers a mountain climb with some chance of not succeeding over one where she is guaranteed to succeed. But this objection fails to take into account the fact that Ann considers the chance distribution to be part of the ultimate goods she cares about. When we take this into account, we should see that Ann is guaranteed

to satisfy objective Stochastic Dominance, if we assume that she satisfies the Principal Principle and is practically rational in all other respects, meaning that she always prefers what she believes to be the best means to her ends. More precisely, the assumption is that, in addition to satisfying the PP, her preferences can be represented as maximising the value function of some normative theory of rational choice (e.g. expected utility theory, risk-weighted expected utility theory, Jeffrey's evidential decision theory, causal decision theory, etc.).

It is easy to see that given the above assumption, if a decision-maker satisfies the Principal Principle, then she satisfies objective Stochastic Dominance when choosing between lotteries with known and fixed chances. Since she is a maximiser her degrees of belief satisfy subjective SD. And if she satisfies the PP, then her beliefs about how likely each lottery is to yield each prize will match the chances associated with that lottery. Hence, she satisfies objective SD. So a maximiser who satisfies the PP satisfies objective SD when she knows the chances.

A simple example will suffice to show that a person who violates the Principal Principle, but is some sort of maximiser, violates objective Stochastic Dominance when choosing between lotteries with known chances. For if there is one pair of such lotteries where she violates objective SD, then she of course does not satisfy objective SD as a general principle in this type of situation. Suppose now that although Alice knows that the chance lottery L_1 has of yielding the prize is 0.3, but lottery L_2 gives a 0.2 chance of winning the prize, her credence in winning the prize from these two lotteries is reversed: She believes to degree 0.2 that lottery L_1 will result in the prize but to degree 0.3 that lottery L_2 will. In other words, Alice violates the Principal Principle. Then assuming that she prefers winning the prize to not, she will prefer lottery L_2 to L_1 , since the former *subjectively* stochastically dominates the latter.¹⁴

¹⁴Recall that we are assuming that the only good Alice cares about, in this situation, is money. However, we should emphasise again that this assumption is not essential for our ar-

But that means that she prefers the option that is *objectively* stochastically dominated. This will generally be true: If a person violates the Principal Principle, then given a large enough set of lotteries with known chances, we can always offer her a pair of lotteries with respect to which she violates objective Stochastic Dominance.

To sum up: It seems that in the situation we are envisioning—namely, one where a person is choosing between lotteries with known chances for prizes—it would be practically irrational to violate the objective version of the Stochastic Dominance principle. But that is exactly the type of situation that is relevant for evaluating whether the Principal Principle is a requirement of practical rationality. Hence, since an otherwise practically rational person satisfies objective SD in this situation just in case she satisfies the PP, we conclude that the latter is a requirement of practical rationality.

5 Linearity is Not a Requirement of Rationality

In Section 2 we proved that, given the Principal Principle, Chance Neutrality implies Linearity and Chance Neutrality* implies Linearity*. Moreover, we have argued that the Principal Principle *is* a requirement of practical rationality. Hence, to show that Chance Neutrality (Chance Neutrality*) is not rationally required, it is sufficient to show that Linearity (Linearity*) is not a requirement of rationality. That is the aim of this section.¹⁵ The difference between Linearity and Linearity* will not make any difference to the arguments in this section,

argument. In particular, we need not assume Chance Neutrality for the argument to go through. All that is required, is the assumption that the person satisfies subjective Stochastic Dominance, which she will do if she is a maximiser.

¹⁵But we should emphasise that we reject Linearity only with respect to chance, not credence.

so we will simply refer to the former.

Linearity is a very powerful condition. It makes it possible to identify the value of a lottery with the sum of the probability weighted values of its prizes. Hence, it makes it easy to calculate the value of lotteries of the type we discussed in the last section, once the value of each prize has been determined. More generally, Linearity ensures that for any prospect that results in a (objective) probability distribution over possible outcomes, the value of the prospect is fully determined by this probability distribution and the values we assign to the outcomes.

Just how strong a constraint it places on an agent's attitudes is exemplified by the fact that it implies *risk neutrality* with respect to desirability value. To take an example, Linearity entails that a person cannot but be indifferent between a good that has desirability value 1 and a 50/50 gamble that either results in a good with desirability value 0 or one with a desirability value of 2. However, if a person prefers \$30 for sure to a lottery that has 0.5 chance of resulting in the person winning \$100 but an equal chance of her receiving nothing, then Linearity entails that this person evaluates the difference between winning \$30 and winning \$0 to be greater than the difference between winning \$100 and winning \$30. In other words, a person with the preference in question can only satisfy Linearity if she thinks that the desirability of each extra dollar when she has already won \$30 is (on average) less than the desirability of each dollar before she wins \$30.

This feature of the von Neumann and Morgenstern theory has been a common object of criticism, with critics complaining that the theory cannot fully account for other rational attitudes to risk (see, for instance, Hansson 1988 and Buchak 2013). The critics ask: What if a person insists that she values each extra dollar equally, no matter how much she has won before, but nevertheless prefers the prospect of \$30 for sure to the 50/50 gamble on winning

\$100 or nothing? Their preference for the sure prospect over the gamble has nothing to do with how she evaluates marginal increases in monetary outcomes, the person may insist, but is simply due to her aversion to accepting the risk of something bad in exchange for a chance of something good.

The main weakness of this objection is that it assumes that people have a direct access to how desirable an outcome is, not just by itself, but also in comparison to all other outcomes; and can figure this out without consulting her preferences for lotteries, or uncertain prospects, involving these outcomes. And that may seem questionable. Can we really say for sure, and in exact terms, how desirable it is to win the 61st dollar, compared to how desirable it is to win the 30th dollar, without considering what lotteries we would accept?¹⁶

It might seem that in moral evaluations we *do* have direct access to the types of values that are required to make an argument of the above kind. Various moral theories, for instance, entail that if both Ann and Bob are dying of cancer, and if their situation is exactly the same in all morally relevant respects, then the moral value of treating both of them is exactly twice the value of treating only one of them. But then Linearity entails that we should be indifferent between being able to treat Ann (or Bob) for sure, and a lottery that has a 50/50 chance of resulting in either a treatment for both or for neither. But many people will reject that entailment. Some will say that we should prefer treating one for sure, since the risk of treating neither is not worth the chance of treating both, while others will say that reasons of fairness require us to prefer the lottery since it gives both Ann and Bob some chance of being treated. Nevertheless, on the face of it, it seems that these people need not deny that treating both Ann and Bob is twice as valuable as treating one of them.

This argument too rests on the assumption that direct ascertainment of moral value, inde-

¹⁶For an interesting debate of this issue, see (Watkins 1977) and (Harsanyi, 1977).

pendently of considerations of risk, is possible. And so it is not likely to persuade adherents of the von Neumann and Morgenstern theory. Fortunately, we do not have to be persuaded by either of the above arguments to see the benefits of giving up Linearity as a rationality requirement. For a theorem proved by Matthew Rabin (2000), which makes no assumptions about people's abilities to evaluate outcomes without consulting their preferences for lotteries involving these outcomes, seriously weakens the plausibility of the claim that Linearity is a requirement of rationality. The theorem itself is rather convoluted, so we won't state it formally, but in the next two paragraphs we provide two examples of what the theorem entails, which should suffice to show how damaging it is for the claim that Linearity is rationally required. (In the paragraph after that, we give an illustration of Rabin's proof.) The lesson from the theorem is that if all departures from risk neutrality must be explained in terms of the agent's utility function over actual outcomes—as we must do if we assume Linearity—and if a person is generally risk averse (more precisely, her utility function is concave at all points in a specified range), then a modest risk aversion with respect to small stakes entails absurdly high levels of risk aversion for larger stakes.¹⁷

¹⁷One might perhaps question the rationality of modest risk aversion, for instance by pointing out that in the long run, such an attitude is likely to make the agent worse off than if she were risk neutral, since she will, for instance, forgo gambles with a positive monetary expectation. (We thank a referee for *Philosophy of Science* for pressing us on this issue.) Our claim is not that moderate risk aversion is *always* rational, nor that such an attitude is rationally required. However, it is a widely observed fact that people very often do display a modest risk aversion when it comes to gambles with low stakes, and most people have the intuition that there is nothing irrational about that. Hence, we think that decision theory should have room for such attitudes without thereby entailing that the agents in question must be absurdly risk

Here is an example of what Rabin's theorem entails: If from any total wealth level,¹⁸ a person turns down a 50/50 bet that results in her either losing \$100 or gaining \$110, then she must turn down 50/50 bets that result in her either losing \$1,000 or gaining *any* sum of money—including an infinite sum! But that seems implausible. One could reasonably doubt that the antecedent of this conditional is very often satisfied (Armendt 2014), since most people would, when they are sufficiently rich, perhaps accept the bet in the antecedent. However, we can easily *imagine* a person who does reject the former bet for all total wealth levels, and it seems that this person would not be rationally required to turn down 50/50 bets of losing \$1,000 or gaining an infinite sum of money.

There are other less extreme implications of Rabin's theorem whose antecedents are much more likely to be satisfied. Here is one such example: A person who turns down, when her total wealth is less than \$300,000, a 50/50 gamble that results in her either losing \$100 or winning \$125, must, when her wealth is \$290,000, turn down a 50/50 gamble that results in averse with respect to larger stakes. For it is our belief that decision theory should in general be consistent with considered preferences of otherwise rational people.

¹⁸One might worry that we are here introducing the wealth of a person as an important ingredient in explaining her choices, without explicitly representing wealth levels in the formal framework we developed in Section 3. (We thank a referee for *Philosophy of Science* for reminding us of the need to address this worry.) The worry is unnecessary. A Jeffrey-desirability function is always taken to be implicitly conditioned on whatever the agent takes to be her status-quo. So in particular, the function always takes into account the wealth of the person whose desires the function represents. Hence, given that our framework is an extension of Jeffrey's, there is no need to explicitly account for an agent's wealth level in our framework.

her either losing \$600 or winning \$36,000,000,000. But that is absurd. Very many people would turn down the first bet at the wealth level in question, but hardly anybody would turn down the second bet, and there is nothing irrational about this pair of choices. But if that is so, then there are attitudes to lotteries that are inconsistent with Linearity but are nevertheless not irrational. Hence, Linearity cannot be rationally required.

Although we will not give a general proof of Rabin's theorem, the following illustration (based on Rabin 2000, 1282) of how the proof works might be instructive, since it shows how little has to be assumed to prove the theorem. Let's focus on a person who at all wealth levels turns down a 50/50 gamble that results in her either winning \$11 or losing \$10. Since she turns down the bet for her actual wealth level, she must, given Linearity, value the 11th dollar above her current wealth at most $10/11$ as much as she values the 10th to last dollar of her current wealth. (Recall that Linearity entails that people are risk-neutral with respect to desirability, and that risk averse behaviour with respect to money must be due to the decreasing marginal worth of money.) Similarly, since she would turn down the bet if she were \$21 richer, she must value the 32nd dollar above her current wealth at most $10/11$ as much as the 11th dollar above her current wealth, and hence at most $(10/11)^2$ as much as the 10th to last dollar of her current wealth. By the same reasoning, since she would turn down the bet if she were \$189 richer, she values the 200th dollar from her current wealth level at most $(10/11)^{10}$ as much as she values the 10th to last dollar of her current wealth. Iterating this procedure allows us to figure out the rate at which the value of extra dollars decrease, according to this person, which allows us to determine what bets she will accept. And it turns out that since she is unwilling to accept the bet in question at any wealth level, she should, by Linearity, be unwilling to accept a 50/50 bet that results in her either losing \$145 or gaining an infinite

amount of money!¹⁹

The theorem's strongest assumption is that the agent's value function is concave everywhere in the specified wealth range. That assumption may not very often be satisfied in practice, since the value of, say, an extra dollar might *increase* at certain points; for instance, when an agent is close to having enough money to buy a new house or make other investments that would in some way change her life. However, the point is that again, we can imagine a person that satisfies the conditions of the theorem. And we think that it would be unfortunate if our theory of practical rationality *required* that person to satisfy the conditionals that Rabin's theorem entails.

Once we give up Chance Neutrality we can avoid Rabin's result. For instance, although Jeffrey's theory entails risk neutrality with respect to desirability (since it satisfies Linearity for credence), his theory, after it has been extended to chance propositions (and if we abandon Chance Neutrality), makes it possible to explain people's preferences for gambles in terms of their (conative) attitudes to chances. In other words, unlike what von Neumann and Morgenstern's theory—and, more generally, Linearity—requires, and contrary to the assumption of Rabin's theorem, we need not, once we have extended Jeffrey's theory to chance propositions, explain all risk attitudes with respect to, say, money, in terms of people's attitudes to monetary outcomes. Therefore, people whose utility function for money is generally concave, and who satisfy the antecedents of Rabin's conditionals, need not satisfy their consequents.

¹⁹A proof of the last claim can be found here: <http://merlin.fae.ua.es/iturbe/RabinEcon00.pdf>.

6 Concluding Remarks on Risk Aversion

There is a substantial decision-theoretic literature challenging the orthodox theory of decision-making under objective uncertainty deriving from the seminal work of von Neumann and Morgenstern, and in particular challenging the Independence axiom that is its cornerstone. Much of it takes Maurice Allais' (1953) infamous paradox, and the now-abundant evidence drawn from experimental studies of choices amongst lotteries, to show that people can exhibit apparently-rational patterns of preference that cannot be represented as maximising expected utility. And in its wake a number of models of decision-making under uncertainty have been proposed, of which cumulative prospect theory and rank-dependent utility theory (or anticipated utility) are perhaps the most established (see, for instance, Quiggin 1982, Tversky and Wakker 1995, Wakker 2010), all of which allow for violations of Linearity.

A full evaluation of this literature is beyond the scope of this paper, but in view of the fact that our argument shares some common features with it, a few remarks are in order. First, a central motivation of this literature is a perceived inadequacy of the treatment of risk aversion in the orthodox theory (something we alluded to before). Linearity, while perfectly consistent with risk aversion with respect to undesirable outcomes, requires risk neutrality with respect to the *utilities* of these outcomes. But experimental evidence suggests that people often give more weight to the difference between the certainty of avoiding loss and a very low probability of loss than they do to equal differences in the probabilities of losses when these probabilities are quite high. In rank-dependent utility theory these 'non-linear' attitudes to the chances are modelled by a transformation of the probabilities into non-probabilistic decision weights on the utilities of the outcomes. Our treatment suggests instead that these attitudes should be incorporated into the agent's desirability function by taking chances to be features of the world that the agent cares about. These two approaches may predict much the same

behaviour, which makes a transformation of the probabilities into non-probabilistic decision weights unnecessary. Nevertheless, the two approaches are distinct in their understanding of the attitudes that underpin the behaviour in question. Whereas rank-dependent utility theory and theories based on it (such as Buchak's (2013)) suggests that risk attitudes are a type of attitude that is distinct from the desires that utility (or desirability) functions represent, our approach assumes that risk attitudes are just a particular type of desire.

Secondly, it has not been generally recognised in this literature that if the aforementioned arguments against Linearity are accepted, then it is not just the typical preferences displayed in the Allais paradox that may be rational, but also the pattern of preferences often observed in the Ellsberg paradox. The typical Ellsberg preferences have widely been taken to challenge the normative adequacy of Bayesianism. But as Bradley (2015) has shown, these preferences are perfectly compatible with the standard preference axioms of Bayesian decision theory, provided that agents who have these preferences take the chances in Ellsberg's lotteries to be ultimate goods with decreasing marginal utility. Indeed, the Ellsberg preferences may be regarded as further evidence against Chance Neutrality, since rejecting the principle would allow for a rationalisation of preferences that many consider to be intuitively rational.

We can put the point in another way. The Ellsberg preferences are commonly thought to result from *ambiguity aversion*; i.e., from a preference for lotteries with known rather than unknown probabilities. Many of those who reject Linearity nonetheless think that such aversion is irrational, perhaps because they implicitly accept Chance Neutrality.²⁰ But since the rejection of Linearity means that Chance Neutrality is not a requirement of rationality, the

²⁰To take just one example, see Buchak (2013), who rejects Linearity because of what it entails about rational risk attitudes, but nonetheless develops a theory of rational choice that is inconsistent with the Ellsberg preferences.

door is open to accommodating ambiguity aversion within Bayesian decision theory by allowing rational agents to have non-neutral attitudes to chances.

Finally, the main aim of this paper has been critical and focused on developing an argument against the deeply entrenched assumption of Chance Neutrality. But, as the above remarks suggest, we think of the rejection of Chance Neutrality as the first step in a positive programme of developing a decision theory that explicitly represents both objective and subjective uncertainty and which allows for a wide range of rational desirability attitudes to the former. Such a programme holds out the promise of a unified and principled treatment of risk attitudes and a resolution of associated decision-theoretic paradoxes, as well as opens a door to domain-specific studies of the relationship between agents' attitudes to quantities of relevant goods and their attitudes to chances of them (for there seems to be no good reason to think that the relationships between quantities and chances of money should be the same as between, say, health outcomes and their chances). But evidently we have only taken the first steps in that direction in this paper.

References

- Allais, Maurice. 1953. "Le comportement de l'homme rationnel devant le risque: Critique des postulats et axiomes de l'école Américaine." *Econometrica* 21(4): 503-46.
- Armendt, Brad. 2014. "On Risk and Rationality." *Erkenntnis* 79(6): 1119-27.
- Bradley, Richard. 1999. "Conditional Desirability." *Theory and Decision* 47(1): 23-55.
- Bradley, Richard. 2015. "Ellsberg's Paradox and the Value of Chances." *Economics and Philosophy*, forthcoming.
- Bradley, Richard, and H. Orri Stefánsson. 2015. "Counterfactual Desirability." *The British Journal for the Philosophy of Science*, forthcoming.
- Broome, John. (1991). *Weighing Goods*. Oxford: Basil Blackwell.
- . 1999. "Can a Humean be Moderate?" In *Ethics Out of Economics*, 68-87. Cambridge: Cambridge University Press.
- Buchak, Lara. 2013. *Risk and Rationality*. Oxford: Oxford University Press.
- Hall, Ned. 1994. "Correcting the Guide to Objective Chance." *Mind* 103(412): 505-17.
- Hansson, Bengt. 1988. "Risk Aversion as a Problem of Conjoint Measurement." In *Decision, Probability, and Utility*, ed. Peter Gärdenfors, and Nils-Eric Sahlin, 136-58. Cambridge: Cambridge University Press.
- Harsanyi, John C. 1977. "On the Rationale of the Bayesian Approach: Comments on Professor Watkins's Paper." In *Foundational Problems in the Special Sciences*, ed. Robert E. Butts and Jaakko Hintikka, 381-92. Dordrecht: D. Reidel Publishing.

- Hoefler, Carl. 2007. "The Third Way on Objective Probability: A Sceptic's Guide to Objective Chance." *Mind* 116(463): 449-96.
- Jeffrey, Richard. 1983. *The Logic of Decision*. Chicago, IL: The University of Chicago Press.
- Lewis, David. 1980. "A Subjectivist's Guide to Objective Chance." In *Studies in Inductive Logic and Probability*, ed. Richard Jeffrey, 263-93. Berkeley, CA: University of California Press.
- Lewis, David. 1994. "Humean Supervenience Debugged." *Mind* 103(412): 473-90.
- Pettigrew, Richard. 2012. "Accuracy, Chance, and the Principal Principle." *Philosophical Review* 121(2): 241-75.
- Pettigrew, Richard. 2013. "A New Epistemic Utility Argument for the Principal Principle." *Episteme* 10(1): 19-35.
- Quiggin, John. 1982. "A Theory of Anticipated Utility." *Journal of Economic Behavior and Organization* 3(5), 323-43.
- Rabin, Matthew. 2000. "Risk Aversion and Expected Utility Theory: A calibration Theorem." *Econometrica* 68(5): 1281-92.
- Savage, Leonard. 1972. *The Foundations of Statistics*. New York, NY: Dover Publication (second revised edition).
- Tversky, Amos, and Peter Wakker. 1995. *Risk Attitudes and Decision Weights*. *Econometrica* 63(6): 1255-80.

von Neumann, John, and Oskar Morgenstern. 1953. *Games and Economic Behavior*.
Princeton, NJ: Princeton University Press.

Wakker, Peter. 2010. *Prospect Theory: For Risk and Uncertainty*. Cambridge: Cambridge
University Press.

Watkins, John W. N. 1977. "Towards a Unified Decision Theory: A non-Bayesian Ap-
proach." In *Foundational Problems in the Special Sciences*, ed. Robert E. Butts and
Jaakko Hintikka, 245-80. Dordrecht: D. Reidel Publishing.