H. Orri Stefánsson[1] and Richard Bradley[2]

Statewise Dominance and Changing Risk Attitudes[3]

Abstract: We show that Statewise Dominance has problematic implications given two seemingly innocuous assumptions. The first of these assumptions is that rational people can change their attitudes to risk. The second assumption is that there are possible lives that would be just as good (or preferable), for those leading them, even though non-equivalent levels of wellbeing at two points in time were permuted. In the presence of these two assumptions, Statewise Dominance implies that there are risky situations where a person's attitude to risk can have no impact on her preference between risky prospects. So, we face an inconsistency between a set of prima facie plausible assumptions about attitudes to risk and how preferences over life-stages relate to preferences over lives.

Keywords: Statewise Dominance; risk attitudes; rationality; preference change.

[1] Stockholm University and Institute for Futures Studies, orri.stefansson@iffs.se

[2] London School of Economics and Political Science, r.bradley@lse.ac.uk

[3] This paper has been presented at University of Stirling, University of Oxford, Lund University, Stockholm University and Institute for Futures Studies. We are very grateful for all the useful comments and suggestions we received. Funding from the Swedish Research Council (1-VR-041) is gratefully acknowledged. Stefánsson in addition gratefully acknowledges funding from the Knut and Alice Wallenberg Foundation (WAF 2023).

1. Introduction: Statewise Dominance

The axiom of Statewise Dominance is thought by many to be an uncontroversial constraint on the ranking of *prospects*, that is, alternatives that result in different outcomes depending on which state of the world is the actual one. Informally, any version of Statewise Dominance implies that if prospect I results in at least as preferable outcome as prospect II in every possible state of the world, then prospect I is at least as preferable as prospect II.

Statewise Dominance has however been questioned as a constraint on the ranking of *social* prospects. For instance, suppose that two people are in equal need of an indivisible resource. Further suppose that an outcome where one person gets to use the resource is just as preferable, or good,[4] from a social point of view, as an outcome where the other person does so. Finally suppose that prospect I gives the resource to one person in any state of the world, whereas prospect II gives the resource to one person if a coin comes heads up but gives the resource to the other person if the coin comes up tails. Then some find that prospect II is strictly more preferable than prospect I, even though they result in equally good outcomes in every possible state of the world (Diamond 1967).[5]

In contrast, Statewise Dominance has not been much questioned as a constraint on the ranking of *individual* prospects, that is, alternatives that result in different outcomes for a single individual (say, different wellbeing levels for that individual) depending on which state of the world is the actual one.[6] The aim of this paper is to put pressure on Statewise Dominance as a general requirement of individual rationality. We show that the axiom has problematic implications given two seemingly innocuous assumptions.

The first of these assumptions is that rational people can have different attitudes to risk, when it comes to their own wellbeing, and that these attitudes can change over time. For instance, a rational person might be risk seeking in their youth but more risk averse when they get older. (However, we need not assume that there are *no* constraints on rational attitudes to risk, as we explain in section 3.)

The second assumption is that there are possible lives that would be just as good (or preferable), for the those leading them, even though non-equivalent levels of wellbeing at two points in time were permuted.[7] For instance, there could be a life

[4] We use 'preferable' and 'good' synonymously in this paper. In section 2 we say more about how we understand 'preferable'.

[5] Some may want to respond to this by making goodness depend on whether the outcome was brought about by a lottery or not (see, e.g., Broome 1991). We briefly discuss this type of response in section 5.

[6] One exception is the literature on infinite decision theory—either with infinitely many prospects or infinite utility or wellbeing. (See for instance Peterson 2022.) Our argument only involves two prospects and finite wellbeing.

[7] It may be worth emphasising from the start how weak this assumption is. The assumption is *not* that levels of

such that the person leading it would be equally well off overall even though the wellbeing they enjoy when they are, say, 50 were switched with the wellbeing they enjoy when they are, say, 51 (and even though they are at different levels of wellbeing at 50 and 51).

These two assumptions are not, strictly speaking, inconsistent with Statewise Dominance. However, Statewise Dominance implies, in the presence of these two assumptions, that there are situations where a person's attitude to risk has no bearing on how preferable some prospects are for her, even though they clearly differ in how risky they are. And that we find implausible.

We take our result to put pressure on those who claim that Statewise Dominance is a general constraint on the ranking of prospects. But rejecting Statewise Dominance is, of course, not the only reasonable response to our result. Some may want to deny our permutation assumption, others may deny that risk attitudes can rationally change, and yet others may be willing to deny our assumption about how a person's changing attitudes to risk can permissibly affect how preferable prospects are for her. So, strictly speaking, we have simply found an impossibility result involving a set of *prima facie* plausible assumptions about, on the one hand, how to evaluate risky prospects,[8] and, on the other hand, how evaluations of life-stages relate to evaluations of whole lives.

In the next section we present the first version of this impossibility result. In the section after that (section 3), we significantly weaken our assumptions about, first, what risk attitudes are permissible, and, second, how life-stage preference relates to lifetime preference (i.e., our 'Pareto' condition). In section 4 we substantially weaken our permutation assumption. In the penultimate section (section 5), we defend our assumption that a person's attitude to risk can change. Section 6 concludes.

2. The impossibility result (first version)

The focus of our concern is a rational agent's preference ranking of distributions of their wellbeing across different stages of their life and across different states of the world. By someone's wellbeing at some life-stage and in some state of the world, we

consumptions, or the *goods* that the person has, can be permuted without affecting how good the life is. That would be a much stronger and quite implausible assumption. Nor do we take the assumption to be true of *every* life. We merely assume that there *exist* possible lives, and *two* points in time, such that the levels of *wellbeing* at those points can be permuted without affecting how preferable the lives are.

[8] In parts of decision theory and economics, 'risk' (and 'risky prospect') connotates a situation where the decision-maker has access to external (and, in some sense, objective) probabilities, whereas 'uncertainty' connotates a situation where the decision-maker must try to consult their own subjective estimates of probabilities when making a decision. However, as will become clear in the next section, our result is not limited to 'risk' thus understood; in fact, no (explicit and quantitative) probabilities are needed to derive our result.

mean how well off they are at that state and time, independently of how things are or could be at other times or in other states. More generally, we take a person's wellbeing at a state and time to fully account for everything that is relevant to how good the person's life is for them in the particular state and at the time in question.

For simplicity we assume that we can measure how well off the agent is on a cardinal scale, which means that, for instance, the difference between being at wellbeing level 20 and at wellbeing level 30, on this scale, corresponds to the same difference in how good the person's life is as the difference between being at wellbeing level 50 and at wellbeing level 60. (In section 5 we explain why this assumption is consistent with not assuming risk neutrality with respect to wellbeing.)

A distribution of wellbeing for a person over time will be called a *life*. (So, more than one person can lead the same life, thus understood.) For instance, suppose we divide each life into four equally long life-stages, intuitively corresponding to, say, childhood, early adulthood, middle age, and old age. (As we shall soon see, we only need two life-stages for our argument. The argument, however, also works, and is arguably more compelling, with a more fine-grained division of lives into life-stages.) So, for instance, the distribution 10, 15, 15, 10 corresponds to a life, namely, the life of anyone who is at wellbeing level 10 when they are a child, at level 15 during both early adulthood and middle age, and at level 10 during old age.[9]

A distribution of wellbeing across states of the world at a life-stage is called a *life-stage prospect* while a distribution of wellbeing across both states of the world and life-stages is called a (complete) *prospect*. Thus, a prospect maps each state of the world to a life, while a life-stage prospect maps each state at a fixed life-stage to the person's wellbeing level at that life-stage. Throughout we will only need to work with two prospects at the same time, which we (initially) simply denote I and II, and only two states of the world, which we denote s1 and s2 (and which form the set **S**). I(s1) denotes the (complete) life that results from prospect I when state s1 is actual. So, I(s1) is identical to some distribution of wellbeing, such as 10, 15, 15, 10.

By an agent's preference ranking, at some moment of time, of distributions of wellbeing across different states and life-stages, we mean their fully informed, reflective preference over these distributions. It is, in other words, the preference that the person *would* have if she were informed by the available evidence (which sometimes means knowing the probabilities of the relevant states and the prospects with which she will later be faced) and that she would endorse, even on reflection, based on this

[9] It may be worth noting that for the purpose of our argument, what we are calling a 'life' need not correspond to a life as it would ordinarily be understood. For instance, our argument works just as well if the life-stages add up to, say, the years between 20 and 60. In fact, our argument works even if all the four stages add up to only one year, so that each life-stages would be only three months. Nevertheless, we will continue using the terms 'life-stages' and 'lives'.

evidence and her non-cognitive attitudes (including her attitude to risk). Such informed and reflective preferences at a time may be regarded as capturing what is preferable for the agent by her lights, taking into account all features of the distribution. Crucially, we assume that informed and reflective preferences can change over time,[10] as we shortly explain.

We will be concerned with the agent's preferences, at each life-stage, over the corresponding life-stage prospects, over lives, and over complete prospects.[11] It may be worth emphasising that we need not assume that a person's preference at a life-stage, between the corresponding life-stage prospects, is independent of what prospects the person expects to be faced with at other life-stages, nor independent of how prospects differ with regard to past life-stages. Nor do we have to assume that the agent's preferences over life-stage prospects is independent of what their wellbeing was or is predicted to be at other stages.

Moreover, for now, we make no assumption about how preferences over life-stage prospects relate to preferences over complete prospects. In a moment we introduce a ('Pareto') condition that constrains the relationship between these two types of preferences. But other than that, we remain agnostic about how these two types of preferences do, or should, relate to each other.

To set things up more formally, let $\mathbf{T} = \{T1, T2, \ldots, Tn\}$ be a set of possible life-stages, $\mathbf{L} = \{L1, L2, \ldots, Ln\}$ be the corresponding set of possible lives with, say, $L2(T3)$ being the wellbeing level of life L2 at life-stage T3. Our results are consistent with the set **L** being either finite or infinite. More generally, as will become apparent, we only need to make very minimal assumptions about the content and structure of **L**. In fact, all we require is that the set is closed under life-stage permutations (more on which below).

Let $I \gtrsim_T II$ denote that the person of interest weakly prefers prospect I to II at life-stage T, $I \sim_T II$ denote that the person is indifferent between I and II at life-stage T, and $I \succ_T II$ that they strictly prefer I to II at life-stage T.[12] In contrast, $I_T \succ_T II_T$ denotes that, at life-stage T, the person strictly prefers *life-stage prospect* I_T to life-stage prospect II_T.

Preference between lives can be identified with preferences between complete prospects that result in the same distribution of wellbeing across life-stages (that is, in the same lives) in all states of the world. However, preferences over life-stage prospects must be treated separately. After all, it makes little sense to say that someone prefers

[10] Bykvist (2006: 270) convincingly demonstrates that even idealised (e.g., informed and reflective) preferences can change.

[11] In addition, a person may of course, at a life-stage, have preferences over life-stage prospects that correspond to some *other* life stage. Such preferences will however play no direct role in our formal argument.

[12] As is typically done, we assume that $A \sim B$ just in case both $A \gtrsim B$ and $B \gtrsim A$ hold; while $A > B$ just in case $A \gtrsim B$ holds but not $A \sim B$.

some complete prospect to some life-stage prospect. Hence, we use bold font when stating preference over life-stage prospects ('$\succ$') but not when stating preferences over complete prospects ('$\succ$').

Part of our motivation is to examine how preferences between life-stage prospects should relate to preferences between (complete) prospects. One way to read the main result of our paper is that what we have identified a tension within the following triple: (i) a natural assumption about how preferences between life-stage prospects relate to preferences between (complete) prospects, (ii) a plausible assumption about attitudes to risk, and (iii) Statewise Dominance.

The version of Statewise Dominance with which we shall be working can now be formally stated as follows:

> **Statewise Dominance.** For any prospects I and II, for any life-stage T in **T**, if for all states s in **S**, $I(s) \succsim_T II(s)$, then $I \succsim_T II$ is rationally required.

Those who are committed to Statewise Dominance presumably assume it to hold both for preferences between life-stage prospects and for preferences between complete prospects (that is, if they make such a distinction). To derive our result, however, we only need to assume that it holds for preferences between complete prospects.

Note that, in light of the traditional definition of the interrelationship between indifference and weak preference (recall footnote 12), Statewise Dominance implies that if for every state s, $I(s) \sim_T II(s)$, then $I \sim_T II$. This implication plays a crucial role in our argument.

We assume that people can rationally have different attitudes to risk and that these attitudes can change. Specifically, we assume that a person need not be risk neutral with respect to their own wellbeing, and that their attitude to risk concerning their own wellbeing may be different at different points in their life. We say that a person is at time T risk neutral with respect to their own wellbeing just in case they are, for any wellbeing magnitude x,[13] indifferent between on the one hand a 'trivial' life-stage prospect that ensures that they enjoy wellbeing of magnitude x at time T with certainty (that is, a prospect such that they enjoy this wellbeing in any state of the world at time T) and on the other hand a 'risky' life-stage prospect that provides them with an *expectation*[14] of wellbeing of magnitude x at time T but results in them enjoying different levels of wellbeing depending on which state of the world turns out to be actual.[15]

[13] This is a definition of *general* risk aversion with respect to wellbeing. Sometimes it is more useful to define risk aversion with respect to a particular interval of wellbeing. We can however set such complexities aside.

[14] This expectation is calculated by multiplying the wellbeing that the person enjoys in a state with the probability of that state, and then adding together the probability weighted wellbeing levels in all the possible states.

[15] We can generalise this definition by formulating it in terms of a *mean-preserving spread* (Rothschild and Stiglitz 1970). That allows us to say which of two *risky* distributions of wellbeing a risk averse person prefers.

In contrast, the person is risk averse if they prefer the trivial prospect to the risky one but risk seeking if they have the opposite preference.

To take a concrete example, we assume that a rational person might at some point in their life strictly prefer, say, enjoying wellbeing of magnitude 10 for sure over a 50:50 gamble[16] between wellbeing of magnitude 5 and wellbeing of magnitude 15—thus being risk averse at that time—even though they have the opposite preference at some other point in their life. Perhaps it is reasonable not to be averse to risk at early life-stages, since one can expect that any losses that result from such behaviour early in one's life can be recovered at later stages, while becoming risk averse at later stages, as the opportunities to make up for losses start to run out.

Put more formally, we use the following assumption to derive our initial result (but in the next section we prove a stronger result with a significantly weaker assumption):

> **Permissible Risk Attitudes**: Let s and s' be any two states in **S** with strictly positive probability and T and T* any two life-stages in **T**. Let e, f, g, be any three well-being levels such that f>e>g, and I and II be two corresponding prospects such that:
>
> $I_T(s)=e,\ I_T(s')=e,$ $\quad II_T(s)=f,\ II_T(s')=g,$
> $I_{T^*}(s)=f,\ I_{T^*}(s')=g,$ $\quad II_{T^*}(s)=e,\ II_{T^*}(s')=e,$
>
> Then one is rationally permitted to have preferences $I_T \succ_T II_T$ and $I_{T^*} \succ_{T^*} II_{T^*}$.

Our next assumption is structurally similar to the *anonymity* condition that is implied by most theories of distributive ethics, and which intuitively ensures that the moral value of a wellbeing distribution does not depend on *who* enjoys which wellbeing level; so, switching people's wellbeing levels does not affect how good the distribution is. (Note however that this does not imply the much stronger and quite implausible condition that switching people's *consumption* levels does not affect how good the distribution is.) The analogous assumption for individual prospects, which we will later significantly weaken, says that how preferable a life is, for the person leading it, should not depend on *when* each wellbeing level is enjoyed; so, switching the wellbeing levels at (equally long) life-stages should not affect how preferable the life is.

To state the assumption more precisely, we make use of the concept of a *permutation*, that is, a reordering of an ordered set or, more formally, a bijection from a set to itself. In particular, we say that one life, L, is a life-stage permutation of another life,

[16] In this framework, a 50:50 gamble is a prospect that results in one of two outcomes depending on which of two equally probable (sets of) states is actual.

L*, just in case there exists a permutation σ on the set of life-stages such that L(T)=L*(σ(T)), for all life-stages T.

> **Life-Stage Permutation Invariance.** If life L is a life-stage permutation of life L*, then $L \sim_T L^*$ is rationally required at all life-stages T.

Let's consider why one might accept Life-Stage Permutation Invariance. Recall that the wellbeing measure is meant to capture everything that is relevant to how good a person's life is for them at some time and state. But then it might seem that it should make no difference whether they enjoy a wellbeing level of 45 when they are fifty years old and a wellbeing level of 55 when they are sixty, rather than enjoying a wellbeing level of 55 when they are fifty years old and a wellbeing level of 45 when they are sixty. If such a permutation did make a difference, then that should arguably be accounted for in the wellbeing measure. In contrast one might reasonably reject a similar invariance condition for prospects, because differences in risks at different life-stages might well matter and wouldn't be accounted for by the wellbeing measure.

In section 4 we will nevertheless consider a reason to weaken Life-Stage Permutation Invariance, based on concern for the *pattern* of wellbeing within a life. As we shall demonstrate in response, our formal argument can be made with a (much) weaker invariance assumption that is not undermined by concern for such patterns. But to keep the discussion as simple as possible, we first illustrate our argument by using the stronger assumption.

Our final assumption is what we call Life-Stage *Ex Ante* Pareto, and which informally says that if a person strictly prefers life-stage prospect I_T to life-stage prospect II_T at every life-stage T, then they should, at any life-stage, strictly prefer (the complete) prospect I to (the complete) prospect II. More formally:

> **Life-Stage *Ex Ante* Pareto.** For any prospects I and II, for any rational person, and for any way of dividing the person's life into finitely many life-stages, T1 to Tn, and corresponding preference relations, $\succsim_{T1}$ to $\succsim_{Tn}$, if $I_{Tj} \succ_{Tj} II_{Tj}$ for all of the life-stages Tj, then at any life-stage Tk, $I \succ_{Tk} II$.

Now, there are well-known problems with *ex ante* Pareto principles in distributive ethics. One such principle, analogous to Life-Stage *Ex Ante* Pareto, says that if every person prefers prospect I to prospect II from their own individual perspective, then anyone would strictly prefer I to II from a moral or impartial perspective. That principle could be rejected by those who think we are required to be concerned with outcome-inequality (Fleurbaey and Voorhoeve 2013). However, concern for equality does not undermine Life-Stage *Ex Ante* Pareto, which only concerns one person.

On other occasions *ex ante* Pareto principles in distributive ethics can be questioned since the unanimity in people's preferences is 'spurious' in the sense that they have conflicting beliefs about the probabilities of the states of the world (Mongin 2016). To avoid the analogous problem in our framework, we could build into the antecedent of our Pareto principle a restriction to people who do not change their beliefs about the probabilities of the states of the world between the relevant life-stages. Given our previous assumption about what it means to have an informed preference, this could mean assuming that no new evidence becomes available between the relevant life-stages. In addition, we could restrict the principle to cases where people's expectations about what prospects they will be faced with at different life-stages, and what wellbeing they will enjoy, remains fixed. Neither restriction would block the impossibility results we derive below.

However, such a weakening is unnecessary, as will become apparent shortly, since our argument only requires a *permissibility* version of the Pareto assumption, which states that if one strictly prefers life-stage prospect I_T to life-stage prospect II_T at every life-stage T, then it is *permissible* to strictly prefer (the complete) prospect I to (the complete) prospect II. But to try to keep the initial discussion as simple and intuitive as possible, we will for now work with the much stronger Life-Stage *Ex Ante* Pareto, which will presumably be familiar from the literature on distributive ethics.

We can now start building the case for our first impossibility result. Consider prospects I and II in Table 1. Each prospect intuitively corresponds to a pair of gambles for a particular person at different stages in their life. To simplify the reasoning, we assume that the states, s1 and s2, are equiprobable. The two prospects differ from each other at both life stages. In fact, prospect II permutes the wellbeing that prospect I offers in each state. Prospect I results in wellbeing level 50 for sure (that is, in both states) at life-stage T1 but is a gamble between wellbeing levels 55 and 45 at life-stage T2, while prospect II is a gamble between wellbeing levels 55 and 45 at life-stage T1 but results in wellbeing level 50 for sure at life-stage T2. Finally, we do *not* assume that all uncertainty relevant to T1 has been settled at T2, in particular, at T2 the person need not yet know whether s1 or s2 is actual (as may happen for instance if they at T1 choose prospect I and observe getting 50).

Table 1. Illustrating first impossibility result

	I		II	
	s1	s2	s1	s2
T1	50	50	55	45
T2	55	45	50	50

Life-Stage Permutation Invariance here implies that the life offered by I in state s1 is exactly as preferable as the life offered by II in state s1; more formally, for any T, $I(s1) \sim_T II(s1)$. The same is true for state s2, that is, $I(s2) \sim_T II(s2)$. Statewise Dominance therefore implies that the person should be indifferent between these two prospects, $I \sim_T II$.[17]

But now suppose that the person in question is risk averse at life-stage T1 but risk seeking at stage T2. And recall that we assumed that the two states are equiprobable (and not necessarily known at T2). Then the person could, we think, reasonably strictly prefer life-stage prospect I_{T1} to life-stage prospect II_{T1} at life-stage T1 and also strictly prefer life-stage prospect I_{T2} to life-stage prospect II_{T2} at life-stage T2. In the next section we explain why we think this preference is perfectly reasonable. It should be evident that Permissible Risk Attitudes implies that the preference in question is indeed permissible. But if the person strictly prefers I_{T1} to II_{T1} at T1 and strictly prefers I_{T2} to II_{T2} at T2, then by Life-Stage *Ex Ante* Pareto, the person should strictly prefer the complete prospect I to the complete prospect II. In contrast, we saw above that Statewise Dominance together with Life-Stage Permutation Invariance implies that the person should be indifferent between these two complete prospects. So, we have arrived at the following conclusion:

> **Impossibility result.** Statewise Dominance, Life-Stage Permutation Invariance, Life-Stage *Ex Ante* Pareto, and Permissible Risk Attitudes form an inconsistent set.

We think that, from a philosophical point of view, the most interesting aspect of the above result is the fact that, in some circumstances where risk attitudes change, Statewise Dominance and Life-Stage Permutation Invariance together require that the person's attitude to risk (at a time) be *irrelevant* to how they compare some prospects, even though these prospects differ in how well they fit with the person's attitude to risk. The only difference between prospects I and II is that the life-stage prospects that correspond to I suit the person's risk attitudes better at both life-stages T1 and T2. In

[17] This result is analogous to a result that Bradley (2022) proves for social prospects. Nebel (2020), Blessenohl (2020), and Stefánsson (2023) contain related results, also involving social prospects. The literature in theoretical economics contains structurally similar results, also for social prospects, the most general of which seems to be Mongin and Pivato (2015). (For a more recent overview of results of this kind, see Li et al. 2024. See also McCarthy et al. 2020 for related results and discussion.) Mongin and Pivato's result implies that *ex ante* Pareto, Statewise Dominance, and a condition that ensures the state independence of preferences over lives (which is built into the notion of wellbeing in our framework), ensures that the preference has an additively separable (e.g., expected utility) representation (given some additional technical assumptions). But they don't discuss a condition analogous to our Life-Stage Permutation Invariance. (All of these results can be seen as generalisations or applications of Harsanyi's 1955 classical result for social prospects.)

all other respects I and II are identical. So, I is better aligned with the person's preferences than II in one respect, namely, in virtue of better fitting the person's attitude to risk at two life-stages, while II is not in any respect better aligned with the person's preferences than I. But then there is nothing that could offset the fact that the life-stage prospects that correspond to I suit the person's risk attitude better at these two life-stages; similarly, there is nothing that could explain why it is impermissible to prefer the complete prospect I to the complete prospect II despite the former being preferable in the aforementioned respect. Therefore, the person's attitude (at a time) to risk can, according to Statewise Dominance and Life-Stage Permutation Invariance, have *no* impact on the comparative preferability of these two prospects. And that, we think, should seem odd even to those who are generally sceptical of principles such as Life-Time *Ex Ante* Pareto (e.g., Bykvist 2024).

3. Generalising the result

We will now generalise the above result, by significantly weakening the Pareto assumption, as well as our assumption about what risk attitudes are permissible. In the next section we generalise the result in a different way, by significantly weakening the permutation invariance assumption.

The first weakening of our assumptions starts from the observation made earlier that for our argument to succeed, we did not have to assume that the person imagined in the last section is *required* to strictly prefer the life-stage prospects corresponding to prospect I to those corresponding to II at both life-stages T1 and T2. We only have to assume that this preference is rationally *permissible*, for instance, for a person who is risk averse at T1 and risk seeking at T2. Here is an explanation for why the preference is permissible. Suppose the person is not myopic,[18] knows which sequences of gambles the two prospects offer over time, and moreover correctly predicts how their risk attitudes will change. At T1 they might then think to themselves: 'I'd rather take the sure thing now and take the risky gamble at T2, since I am averse to risk now but will (or at least may) be risk seeking at T2.' (Recall our assumption that at T2, the person does not yet know which state is actual.) Still, at T2, when they have become risk seeking, they may feel *some* regret for not having been able to take the risk at the earlier state, T1. However, that negative feeling is, we can assume, not sufficiently strong for them to wish that they had taken the risk at T1 (i.e., chosen or gotten II), since that would have meant that they would have to accept the sure thing 'now', when they are risk seeking. So, at T2, they are, all things considered, happy with

[18] If the person is myopic and at each time only cares about the gamble that they face at that time, then they will of course also strictly prefer I to II at both T1 and T2. But some may find such myopia irrational.

having prospect I rather than II. This reasoning strikes us as being perfectly reasonable.

The above explanation of our assumption only makes sense if risk attitudes don't apply only to whole lives (and complete prospects), but can also apply to life-stages (and life-stage prospects). Those who think that risk attitudes apply only to whole lives (and complete prospects) may thus not be convinced by the above justification of our assumption. That is an issue we will have to set aside for now. However, it is worth noting that even those who agree with us that risk attitudes do also apply to life-stage prospects may worry that there is a direct tension between Life-Stage *Ex Ante* Pareto and our assumption that risk attitudes can apply to life-stage prospects, for instance, since together they may seem to rule out that one can hedge against risk by accepting anti-correlated risks at different life-stages. So, perhaps one should never assume *both* Permissible Risk Attitudes *and* Life-Stage *Ex Ante* Pareto.

Although the above is an important objection to our previous use of Life-Stage *Ex Ante* Pareto, it does not undermine the assumption we actually need to derive our result. Since that assumption only concerns what is *permissible*, it merely rules out that a risk averse person is *required* to hedge against risk by accepting anti-correlated risks at different life-stages.

To state the assumption we actually need, we again make use of the concept of a permutation, this time defined on prospects. In particular, we say that prospect II is a life-stage permutation of prospect I just in case there exists a permutation σ on the set of life-stages such that $II_T=I_{\sigma(T)}$, for all life-stages T. Then we can state our next condition as follows:

> **Permissible Strict Preference.** There exist prospects I and II that are life-stage permutations of each other, such that, for some life-stage T in **T**, one is rationally permitted to have preference $I \succ_T II$.

In other words, it can be permissible to strictly prefer a prospect (I) over a life-stage permutation of it (II). Our motivation for Permissible Strict Preference is that, first, we think that everyone should accept a weakened version of Permissible Risk Attitudes, which says that it is *sometimes* permissible to reverse one's preference between two life-stage prospects (e.g., due to a changed attitude to risk). That is, we think that even if one considers Permissible Risk Attitudes from the last section too strong—in particular, if one rejects the idea that for *any* sure wellbeing level and for *any* non-trivial prospect one could rationally prefer the prospect to the sure wellbeing level at one point in one's life while having the opposite preference at some other point—one should still accept this to hold for *one* (non-trivial) prospect plus *one* wellbeing level. More precisely:

> **Permissible Risk Attitudes** (**weaker version**)**.** There exist two states s and s' in **S**, three wellbeing levels, e, f, g, two stages T and T* in **T**, and two prospects, I and II, where
>
> $I_T(s)=e, I_T(s')=e$, $I_{T^*}(s)=f$, $I_{T^*}(s')=g$,
> $II_T(s)=f$, $II_T(s')=g$,
> $II_{T^*}(s)=e$, $II_{T^*}(s')=e$
>
> such that one is rationally permitted to have preferences $I_T \succ_T II_T$ and $I_{T^*} \succsim_{T^*} II_{T^*}$.

But then by a weak *permissibility*-implication of Life-Stage *Ex Ante* Pareto,[19] one can permissibly have the corresponding strict preference over the complete prospects. In other words, Permissible Strict Preference follows from the weaker version of Permissible Risk Attitudes plus a permissibility-implication of Life-Stage *Ex Ante* Pareto.

It should be evident that we can replace Life-Stage *Ex Ante* Pareto and Permissible Risk Attitudes with Permissible Strict Preference in the impossibility result from the last section. Using the same argument as before, we derive from Statewise Dominance and Life-Stage Permutation Invariance that one should be indifferent between the complete prospects I and II, whereas Permissible Strict Preference implies that one can permissibly strictly prefer I to II. So, we have arrived at the following more general (albeit perhaps less intuitive) result, of which the impossibility result from the last section is an implication:

> **General impossibility result** (first version)**.** Statewise Dominance, Life-Stage Permutation Invariance, and Permissible Strict Preference form an inconsistent set.

In this impossibility result, Statewise Dominance and Life-Stage Permutation Invariance state that some preference is *required*, whereas Permissible Strict Preference merely states that some preference is rationally *permissible*, which we take to be analogous to claiming that the opposing preference is *not* rationally required. In the next section, we relax the permutation assumption, such that it too only claims that some preference is permissible.

[19] Recall that this version of the principle says that for any prospects I and II, for any rational person, and for any way of dividing the person's life into finitely many life-stages, T1 to Tn, and corresponding preference relations, $\succsim_{T1}$ to $\succsim_{Tn}$, if $I_{Tj} \succ_{Tj} II_{Tj}$ for all of the life-stages Tj, then at any life-stage Tk *one can permissibly* strictly prefer I to II.

4. Life-Stage Permutation Invariance weakened

Those who accept our argument so far, and agree that it is implausible that the person imagined in previous sections is rationally required to be indifferent between prospects I and II, could of course conclude that we should give up Life-Stage Permutation Invariance. There may in fact be a good reason to reject Life-Stage Permutation Invariance, as formulated above, since the pattern of a wellbeing distribution over time could perhaps affect how good the corresponding life is (see, e.g., Slote 1983, Velleman 1991 and Liu 2024). For instance, it could be argued that a life that starts off badly but ends well is better than a life that starts well but ends badly, even if one life is obtained from the other by life-stage permutation. Or vice versa: perhaps it is more important to be well off as a child than as an adult. Either way, we would have a reason to reject Life-Stage Permutation Invariance.

But now consider *Weak Life-Stage Permutation Invariance* which says that there are *possible* distributions of wellbeing—that is, possible *lives*—such that for *some* pair of consecutive life-stages, we can permute the (non-equivalent) wellbeing levels in those stages but still the corresponding lives are rationally permissibly deemed equally preferable by *someone*. (We could relax the 'consecutive' part, but we include it merely to illustrate how plausible the condition is; more on which below.)

To state this weak invariance condition more formally, now let life L be a *consecutive two life-stage inversion* of life L* just in case one is obtained from the other by switching wellbeing levels at two consecutive life stages. More precisely, L is a consecutive two life-stage inversion of L* just in case for some pair of consecutive life-stages, Ti and Ti+1, L(Ti)=L*(Ti+1) and L*(Ti)=L(Ti+1), while L(T)=L*(T) for all other life-stages T. (So, a consecutive two life-stage inversion is a particular type of life-stage permutation.) Now consider:

> **Weak Life-Stage Permutation Invariance**. There are at least two pairs of distinct lives, (L, L*) and (M, M*), such that the members of the pairs are consecutive two life-stage inversions of each other, and such that one is rationally permitted to have preferences $L \sim_T L^*$ and $M \sim_T M^*$, for all life-stages T.

For instance, take a person who lives to eighty, and let each year in their life correspond to a life-stage; further suppose that the distribution of wellbeing during these eighty life-stages is *random*, in the sense that there is no part of the distribution such that knowing that part makes the rest of the distribution predictable.[20] Finally, suppose that this person is at wellbeing level 50 when 40 years old but at level 55 when

[20] Eagle (2005) prespecifies and defends this notion of randomness as unpredictability, which has precursors in the work of Kyburg (1974) and Suppes (1984). Note that 'randomness', in this sense, does not necessarily involve

41 years old. Then, plausibly, the life would have been just as good, for the person leading it, if the person had instead been at level 55 when 40 years old but at level 50 when 41 years old (leaving all other stages unaffected). And similarly for levels 50 and 45. At the very least, it would seem plausible that for *some* such random distribution and for some consecutive stages of the corresponding life, the life would be just as good even though we permuted non-equivalent wellbeing levels at these life-stages.[21] Note that the claim is not that for *every* person, there are some life-stages for which this permutation assumption holds; only that it could hold for *some* person.

Table 2: Illustrating stronger impossibility result

		I'		II'	
		s1	s2	s1	s2
...		...	...	...	...
T40		50	50	55	45
T41		55	45	50	50
...		...	...	...	...

Now suppose that the person imagined in the last paragraph changes their attitude to risk the night before their 41st birthday, from being risk averse at 40 to being risk seeking at 41. And consider prospects I' and II' in Table 2, where either can be obtained from the other by switching the wellbeing in each state of the world at T40 and T41, where T40 denotes the person's life-stage from age 40 and up to 41 and analogously for T41. (Again, we assume, to simplify, that the two states are equiprobable.) In this case, Weak Life-Stage Permutation Invariance and Statewise Dominance could[22] together imply that the person should be indifferent between I' and II', by an argument analogous to that in section 2. But that seems very odd, for reasons already discussed. In particular, this means that the person's attitude to risk at T40 and at T41 should have *no* impact on how they compare these two prospects' preferability.

Weak Life-Stage Permutation Invariance seems hard to reject. However, perhaps someone may object that even though Weak Life-Stage Permutation Invariance is

a probability distribution.

[21] Here is a more concrete example (inspired by McKerlie 1989: 483) that we think motivates Weak Life-Stage Permutation Invariance. Suppose that you need to go to the dentist for a check-up. As it happens, you don't need anything fixed. So, delaying the visit for a day or two won't make a difference to your teeth. Nor will it make a psychological difference to you (for instance, you aren't waiting anxiously for the appointment). Still, you find a dentist check rather unpleasant. Finally, suppose that the life-stages are rather fine-grained: one stage corresponding to 24 hours. It seems plausible that postponing the visit for a day does not affect how good your life is overall. Nor does moving the visit to one day earlier than planned. But then Weak Life-Stage Permutation Invariance is true.

[22] In the next paragraph we explain why we say that the two principles *could* imply this. Shortly we add a condition which *ensures* that this is in fact implied.

true, it is *not applicable* to I' and II' since these are not instances of the type of distribution that figures in the antecedent of the condition. (Hence, we only said above that the indifference *could* be implied.) In other words, the objection is that even though there are *some* distributions and *some* life-stages for which Life-Stage Permutation Invariance holds, that doesn't mean that it holds for prospect I' and life-stages T40 and T41. Moreover, to ensure that no results like those we have been discussing can be derived, it would have to be claimed that the antecedent of Weak Life-Stage Permutation Invariance is violated by *any* permutation that would be found objectionable by anyone who changes their risk attitude between the life-stages whose wellbeing levels are permuted. This objection illustrates that, strictly speaking, the above impossibility results no longer hold when Life-Stage Permutation Invariance is replaced by *Weak* Life-Stage Permutation Invariance. (Shortly we however prove an impossibility result that involves Weak Life-Stage Permutation Invariance but not the stronger invariance condition.)

Upon examination, the above objection should be found unconvincing. That is, although the above impossibility result, as stated, no longer holds when Life-Stage Permutation Invariance is replaced by *Weak* Life-Stage Permutation Invariance, that does not undermine the philosophical importance of the conflict we have identified. What determines whether wellbeing levels at two life-stages can be permuted without affecting how good or preferable the life is, should be how wellbeing is distributed *within* the life (and possibly the person's attitude to how wellbeing is distributed within their life). After all, what puts pressure on Life-Stage Permutation Invariance, as a general condition, is the importance of patterns of wellbeing within a life but over time. So, we should be able to determine whether Life-Stage Permutation Invariance can be applied by looking separately at each *column* in tables like those above. In contrast, how wellbeing is distributed *across states* (and the person's attitude to that) is irrelevant to the question of whether Life-Stage Permutation Invariance holds with respect to two life-stages. So, we don't need to consider *rows* in tables like those above to determine if Life-Stage Permutation Invariance holds with respect to two life-stages. Similarly, since risk attitudes are about 'row distributions', not 'column distributions', the person's attitude to risk is irrelevant when determining whether a particular permutation affects how good their life is. Therefore, it is hard to see how someone could plausibly maintain that Weak Life-Stage Permutation Invariance is true but nevertheless cannot be applied in a way that gives rise to problematic results like those we have been discussing in this paper.

Recall that Weak Life-Stage Permutation Invariance only claims that it is *possible* that Life-Stage Permutation Invariance holds for some stages of a person's life. We can therefore informally state the above result as follows. Suppose that, as a matter of

fact, Ann thinks her life would be just as good even though we permutated her well-being at two life-stages. As it happens, her attitude to risk changes between these two life-stages. Then Statewise Dominance implies that she is not permitted to strictly prefer one prospect over another even though she weakly prefers the (life stage prospect corresponding to the) one over the (life stage prospect corresponding to the) other at any time and sometimes strictly prefers the one prospect over the other. That, we think, is quite implausible.

However, it is worth noting that we can prove a formal impossibility result with Weak Life-Stage Permutation Invariance by strengthening Permissible Strict Preference slightly. We again make use of consecutive two life-stage inversions, but now applied to prospects, and say that prospect II is a consecutive two life-stage inversion of prospect I just in case for some pair of consecutive life-stages, Ti and Ti+1, $I_{Ti}=II_{Ti+1}$ and $I_{Ti+1}=II_{Ti}$, while $I_T=II_T$ for all other life-stages T. Less formally, a prospect I is a consecutive two life-stage inversion of prospect II if only two life-stage prospects are switched, and where the life-stages involved are consecutive.

Moreover, let a 'non-trivial' life-stage permutation of a prospect be one where the switched life-stage prospects are not equivalent (which is in fact true of all the examples we have considered). For instance, consider again the inversion from the last paragraph, where $I_{Ti}=II_{Ti+1}$ and $I_{Ti+1}=II_{Ti}$, while $I_T=II_T$ for all other life-stages T. This permutation is non-trivial just in case $I_{Ti}\neq I_{Ti+1}$. So, I' and II' in Table 2 are examples of *non-trivial consecutive two life-stage inversions*. Now we postulate:

> **Strong Permissible Strict Preference.** For any prospects I and II that are non-trivial consecutive two life-stage inversions of each other, and for any life-stage T in **T**, one is rationally permitted to have preference $I \succ_T II$.

The motivation behind Strong Permissible Strict Preference is the same as that behind Permissible Strict Preference, namely, liberalism about what attitudes to risk are permissible, and in particular the thought that one is permitted to change one's attitude to risk. Indeed, Strong Permissible Strict Preference is implied by Permissible Risk Attitudes (from section 2) in combination with the permissibility implication of Life-Stage *Ex Ante* Pareto (recall fn. 19).

It may however be worth emphasising that Strong Permissible Strict Preference is far from the claim that *any* risk attitude is rationally permissible—nor of course does it claim that one is permitted to prefer *any* prospect over another—since it merely says that one can prefer any prospect over a *non-trivial consecutive two life-stage inversion* of it. So, despite its name, Strong Permissible Strict Preference is not particularly strong.

Now let lives I(s1) and II(s1) be consecutive two life-stage inversions of each other and the same for lives I(s2) and II(s2). (As before, we assume only two states, but the argument can be easily generalised to multiple states.) Moreover, suppose that prospects I and II are non-trivial consecutive two life-stage inversions of each other. Further suppose that, for some person with preference $\succcurlyeq$, we have $\text{I(s1)} \sim_T \text{II(s1)}$ and $\text{I(s2)} \sim_T \text{II(s2)}$ at life-stage T, as is indeed permissible according to Weak Life-Stage Permutation Invariance. But then Statewise Dominance says that, for this person, $\text{I} \sim_T \text{II}$ is required (at life-stage T). However, Strong Permissible Strict Preference implies that the person is rationally permitted to have preference $\text{I} \succ_T \text{II}$ (at any life-stage T). So, we have arrived at:

> **General impossibility result** (second version)**.** Statewise Dominance, Weak Life-Stage Permutation Invariance, and Strong Permissible Strict Preference form an inconsistent set.

In this impossibility result, only Statewise Dominance is a principle stating that some preference is *required*; the other two merely state that some preferences are rationally *permissible*.

5. Two arguments against allowing different risk attitudes

Throughout this paper we have asserted that, intuitively, it seems hard to deny that a rational person may change their attitude to risk, for instance, becoming more risk averse with age. Before concluding, we nevertheless address two theoretical objections to this claim. The first objection is that the most common way to cardinalize wellbeing (or 'utility') implies risk neutrality with respect to wellbeing. The second objection is that if risk can matter in different ways at different points in time, then that must be in virtue of the fact that it affects wellbeing differently at these different times, in which case it should somehow be accounted for in the wellbeing measure.

The standard way of justifying a cardinal measure of wellbeing, in particular in economics, is to appeal to the method developed by von Neumann and Morgenstern (1944). They pointed out that if a person's preferences satisfy a few axioms, then we can derive a cardinal measure of the extent to which the person values different outcomes by determining their indifference points between lotteries and risk-free outcomes. For instance, if a person is indifferent between on the one hand getting outcome A for sure and on the other hand a prospect that gives them outcome B with (a known) 0.75 probability and outcome C with (a known) 0.25 probability, then we can infer that A is three quarters on the way between B and C on any measure that

correctly represents the extent to which the person values these three outcomes. But crucially, this measure is by construction such that the person is risk neutral when it comes to its quantities. So, if we use this method to cardinalize a person's wellbeing, then the person can neither be risk seeking nor risk averse with respect to wellbeing.

Although in many ways ingenious, the von Neumann and Morgenstern method, and in particular the preference axioms to which it corresponds, has come under heavy criticism—both normative and empirical—over the last decades. Discussing this criticism in detail would take us too far from the core argument of our paper, so let us just mention in broad outlines two strands within this criticism. One set of objections to von Neumann and Morgenstern's theory is that it cannot represent preferences that are not only common but that seem perfectly rational even on reflection (see, e.g., Allais 1953, Buchak 2013, Bradley and Stefánsson 2017). Another set of objections concerns the method's conflation of attitudes that arguably should be kept distinct (Hansson 1988, Buchak 2013, Stefánsson and Bradley 2019). By using indifference points between lotteries and risk-free outcomes to construct a measure of the value of outcomes, the method cannot account for attitudes to risk *per se*. For instance, it cannot distinguish between on the one hand turning down a gamble because the goods that can be won do not bring sufficient additional benefit to outweigh the goods that can be lost, and on the other hand turning down a gamble because one is averse to risk *per se*.

So, how do we then get a cardinal measure of wellbeing if we cannot help ourselves to the von Neumann and Morgenstern method? More importantly, are there ways of generating such a measure without assuming Statewise Dominance? Yes, there are. For instance, we can employ risk-free difference orderings (Alt 1936, Krantz et al. 1971, ch. 4). In other words, we can take as primitive a person's ordering of how much alternatives differ (in a sense that could be interpreted as 'preferability'). That a person can order differences is of course a strong assumption, but it is still far from *assuming* a cardinal measure. However, as long as the ordering satisfies certain axioms, we can *construct* a cardinal measure that corresponds to this ordering without making any assumptions about how the person evaluates prospects. Alternatively, we can use a version of Jeffrey's (1965) decision theory, with some of Jeffrey's propositions playing the role of chance distributions, in a way that generates a cardinal measure but allows for a great deal of flexibility in how the person evaluates prospects (Stefánsson and Bradley 2015, 2019, Bradley 2017). Either way, we get a cardinal measure of wellbeing that allows for changing attitudes to risk and that does not assume Statewise Dominance.

The second objection to our argument is that if a person's attitude to risk at T1 is different from that at T2, then this must be explainable by some difference that should be incorporated into the wellbeing measure. For instance, the 'bad' of risk, according

to someone who is risk averse, should be accounted for in the outcomes of a risky prospect (see, e.g., Weirich 1986). But then it should be possible to account for that in the wellbeing measure in a way that makes the person at all times risk neutral with respect to their own wellbeing. The conflict with Statewise Dominance is thereby avoided.

If this objection is to stand without assuming the von Neumann and Morgenstern approach, then some account must be given of how wellbeing so construed is to be measured. For the sake of argument let's assume that it can in principle be done in an acceptable way. Nonetheless we think that implementing this idea would have serious downsides. The essential point is that by allowing that the wellbeing measure on outcomes incorporates any value that the agent attaches to risk, one seriously undermines both the scope and the appeal of conditions like Statewise Dominance. In addition, it then becomes difficult to see what we are doing when we ask whether a person should satisfy Statewise Dominance (and, more generally, when we ask how a person should evaluate prospects).

In the standard framework in which conditions like Statewise Dominance acquire their typical meaning, a sharp distinction is presumed between the outcomes and the prospects (or 'lotteries') that yields them so that it is possible for any outcome to be yielded with any probability. But if outcomes are in part individuated by the prospect in which they appear then this will no longer be so. The outcome of 'getting \$10 with a high chance' is not one that can be yielded with probability 0.1 in some prospect for instance.[23] Nor can the outcome of 'getting \$10 when there was a 0.5 chance of getting \$100' appear in any prospect that yields \$100 with anything other than probability one-half.

This has several consequences. Firstly, this individuation method mixes up the description of the options that the agent faces with their attitudes to them; in particular, their attitudes to their riskiness. This in turn makes it difficult to say exactly what options the agent is choosing between or to compare and evaluate different agents' risk attitudes (Stefánsson 2023). Secondly, it leads to a narrowing of the scope of the theory since it will only apply to prospects in which the outcomes are sufficiently finely individuated as to take account of any wellbeing effects of the entire probability distribution. And while no formal theory has *universal* scope, we think that this *unnecessarily* narrows the scope *too much*, given that there is an alternative (as pointed out above).

Thirdly, and relatedly, the usefulness of conditions like Statewise Dominance is undermined. In the standard framework, Statewise Dominance promises to seriously

[23] But a *prospect* that results in \$10 with high chance can of course be yielded by another *prospect* with probability 0.1.

simplify some decision-problems, for instance, those where one can see that one prospect results in a more (or less) preferable outcome in each state of the world than any other available prospect. In such cases, Statewise Dominance supposedly has the great practical advantage of allowing us to fully compare the available prospects *state by state*. In other words, there is in such cases no need for 'holistic' evaluations of prospects, where one evaluates prospects *across states*. However, if one must determine how the pattern of risk affects one's evaluation of outcomes, before applying Statewise Dominance, decision-problems can never be simplified in this way. To figure out what one feels about the different patterns of risk that a prospect gives rise to, and how these patterns affect the value of each of the prospect's outcomes, one has to evaluate each prospect holistically, or across states. So, the supposed practical advantage of Statewise Dominance is lost.

In the light of these methodological considerations, we think that it is better to allow free substitution of final wellbeing quantities in lotteries and hence that the value of a prospect may depend not just on expected wellbeing but also on the spread of possible wellbeing levels. But then one should think of the wellbeing measure not as capturing everything that the person cares about, but instead as accounting for everything that is relevant to how good the person's life is in a particular state of the world. (Of course, if the person experiences say anxiety or excitement as a result of the risk involved then *that* should be part of the wellbeing measure, since then that affects how good the person's life is in a state of the world.) But then our arguments involving Statewise Dominance stand, given how Statewise Dominance is traditionally formulated (that is, as we formulated it above).

It is worth noting, however, that our argument does not challenge a principle that is structurally similar to Statewise Dominance, but where the outcomes in the standard formulation have been replaced by outcome-chance pairs (see, e.g., Stefánsson and Bradley 2015, 2019). However, it would be quite misleading to call such a principle *Statewise Dominance*, since it would not be a dominance principle applied to *states* as standardly construed in decision theory, that is, as complete resolutions of uncertainty (see, e.g., Savage 1954, and all the literature in decision theory and economics that uses Savage's theory). It is moreover important to note this is *not* just a terminological issue. The relevant normative question is whether dominance reasoning can safely be applied to *states*, as these are traditionally defined. We interpret our result as showing that, in some cases, the answer to that question is 'no'.

6. Concluding remarks

We have argued that Statewise Dominance is in tension with the idea that a person can permissibly change their attitude to risk. This tension is most sharply exhibited

in the two impossibility results presented in sections 2 and 3, results that assume Life-Stage Permutation Invariance and that it is permissible to strictly prefer one prospect to another when each of the life-stage prospects of the former are preferred to those of the latter. But in section 4 we showed that the tension remains even if we use only Weak Life-Stage Permutation Invariance, a permissibility condition so weak that it is hard to deny.

The lesson we draw from this tension is that there is reason to question Statewise Dominance as a general constraint on the ranking of individual prospects. But questioning Statewise Dominance is of course not the only reasonable response to our result. As we have previously emphasised, some may reject our view on risk attitudes, while others may perhaps question even our weak permutation invariance assumption.

Still, we think that this new reason to give up Statewise Dominance—albeit far from decisive—should be welcome news to those who wish to give up Statewise Dominance for *social* prospects (e.g., for reasons raised by Diamond 1967). For some find it hard to see how individuals could be rationally required to satisfy Statewise Dominance while the 'social planner' is not.[24] So, although we have in this paper focused on individual rationality and individual prospects, we take our result to be important for ethics and social prospects too.

References

Adler, Matthew (2012). *Wellbeing and Fair Distribution: Beyond Cost-Benefit Analysis*. Oxford University Press.

Allais, Maurice (1953). Le Comportement de l'Homme Rationnel devant le Risque: Critique des Postulats et Axiomes de l'Ecole Americaine. *Econometrica* 21(4): 503–546.

Arrow, Kenneth (1950). A Difficulty in the Concept of Social Welfare. *Journal of Political Economy* 58(4): 328-346.

Alt, Franz (1936/1971). On the Measurability of Utility. In Chipman et al. (eds.), *Preference, Utility and Demand*. New York: Harcourt Brace Jovanovich, 424–31.

Blessenohl, Simon (2020). Risk Attitudes and Social Choice. *Ethics* 130 (4): 485-513.

Bradley, Richard (2017). *Decision Theory with a Human Face*. Cambridge University Press.

[24] Adler (2012: 480) for instance asks: "Why should the norms for evaluating actions, in light of an ordering of outcomes, be different in the case where the ordering is a moral rather than non-moral ordering?"

Bradley, Richard (2022). Impartial Evaluation under Ambiguity. *Ethics* 132(3): 541-569.

Bradley, Richard, and Stefánsson, H. Orri (2017). Counterfactual Desirability. *British Journal for the Philosophy of Science* 68(2): 485-533.

Broome, John (1991). *Weighing Goods*. Wiley-Blackwell.

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

Bykvist, Krister (2006). Prudence for Changing Selves. *Utilitas* 18(3): 264–283.

Bykvist, Krister (2024). Wellbeing and Changing Attitudes Across Time. *Ethical Theory and Moral Practice* 27(3): 429-443.

Diamond, Peter A. (1967). Cardinal Welfare, Individualistic Ethics, and Interpersonal Comparison of Utility: Comment. *Journal of Political Economy*, 75(5): 765–766.

Eagle, Anthony (2005). Randomness is Unpredictability. *The British Journal for the Philosophy of Science* 56(4): 749-790.

Fleurbaey, Marc and Voorhoeve, Alex (2013). Decide As You Would With Full Information! An Argument Against Ex Ante Pareto. In Ole Norheim, Samia Hurst, Nir Eyal and Dan Wikler (eds.), *Inequalities in Health: Concepts, Measures, and Ethics*. Oxford University Press.

Jeffrey, Richard C. (1965). *The Logic of Decision*. University of Chicago Press.

Krantz, David, Luce, Duncan, Suppes, Patrick and Tversky, Amos (1971). Foundations of Measurement, Vol. I: *Additive and Polynomial Representations*. New York Academic Press.

Kyburg, Henry. E. (1974). *The Logical Foundations of Statistical Inference*, Dordrecht: D. Reidel.

Hansson, Bengt (1988). Risk Aversion as a Problem of Conjoint Measurement, in Peter Gärdenfors and Nils-Eric Sahlin (eds), *Decision, Probability, and Utility*. Cambridge University Press, pp. 136–58.

Harsanyi, John. C. (1955). Cardinal Welfare, Individualistic Ethics, and Interpersonal Comparisons of Utility. *Journal of Political Economy*, 63(4): 309-321.

Li, Chen, Rohde, Kirsten I. M., and Wakker, Peter P. (2023). The Deceptive Beauty of Monotonicity, and the Million-Dollar Question: Row-First or Column-First Aggregation? Erasmus School of Economics working paper, available at: https://personal.eur.nl/wakker/pdf/deu.pdf.

Liu, Ying. (2024). Prudential Redemption and Its Significance. Philosophers' *Imprint* 24(14): 1-15.

McCarthy, David, Mikkola, Kalle and Thomas, Joaquin Teruji (2020). Utilitarianism With and Without Expected Utility. *Journal of Mathematical Economics* 87: 77-113.

McKerlie, Dennis (1989). Equality and Time. *Ethics* 99 (3): 475-491.

Mongin, Philippe (2016). Spurious Unanimity and the Pareto Principle. *Economics and Philosophy* 32(3): 511–532.

Mongin, Philippe, and Pivato, Marcus (2015). Ranking Multidimensional Alternatives and Uncertain Prospects. *Journal of Economic Theory* 157: 146-171.

Nebel, Jacob M. (2020). Rank-Weighted Utilitarianism and the Veil of Ignorance. *Ethics* 131 (1): 87-106.

Peterson, Martin (2022). The St. Petersburg Paradox. In Edward N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy* (Summer 2022 Edition), available at https://plato.stanford.edu/archives/sum2022/entries/paradox-stpetersburg/.

Rothschild, Michael, and Stiglitz, Joseph (1970). Increasing Risk: I. *A Definition. Journal of Economic Theory* 2(3): 225–243.

Savage, Leonard (1954). *The Foundations of Statistics*. John Wiley & Sons.

Slote, Michael (1983). *Goods and Virtue*. Clarendon Press.

Suppes, Patrick (1984). *Probabilistic Metaphysics*, Oxford: Blackwell.

Stefánsson, H. Orri and Bradley, Richard (2015). How Valuable are Chances? *Philosophy of Science* 82(4): 602-625.

Stefánsson, H. Orri and Bradley, Richard (2019). What Is Risk Aversion? *British Journal for the Philosophy of Science* 70(1): 77–102.

Stefánsson, H. Orri (2023). In Defence of Pigou-Dalton for Chances. *Utilitas* 35(4): 292-311.

Velleman, J. David (1991). Well-Being And Time. *Pacific Philosophical Quarterly* 72(1): 48-77.

Von Neumann, John and Morgenstern, Oskar (1944). *Theory of Games and Economic Behavior*. Princeton, NJ, USA: Princeton University Press.

Weirich, Paul (1986). Expected Utility and Risk. *British Journal for the Philosophy of Science* 37(4): 419-442.