Incommensurability meets risk

Wlodek Rabinowicz

The problem to be discussed in this paper concerns interaction between value incommensurability and risk. More specifically, it focuses on value comparisons between risky actions whose outcomes are guaranteed to be mutually incommensurable in value: they will be incommensurable whatever state the world is in. It might seem that the actions compared should themselves in all such cases be incommensurable. But this intuition, as we shall see, might well be challenged; indeed, it should be challenged.

The problem in its main outline is originally due to Caspar Hare (2010). Later it was taken up by Miriam Schoenfield (2014) and by Bales, Cohen & Handfield (2014). While Hare views it as a problem for rational preferences and rational choice, I here present it as primarily a challenge for formal axiology – a general account of value relations. Schoenfield and Bales, Cohen & Handfield combine these two perspectives: As they present it, the problem arises for rational choice if the latter is guided by considerations of value. All these authors' contributions will be shortly commented upon in what follows, primarily in footnotes. First, however, I need to describe the problem as such and suggest how I think it should be solved. I will propose a solution, but then I will identify a residual problem that I don't know how to solve.

1. Seting up the problem

To set up the problem, in general terms, suppose we consider any four possible outcomes, *a*, *b*, *a*+ and *b*+, such that *a*+ is better than *a*, *b*+ is better than *b*, but no other betterness or equal goodness relationships obtain between the four outcomes in question. Thus, both *a* and *a*+ are incommensurable with both *b* and *b*+. As I use this term, to say that two items are incommensurable in value simply means that neither is better than the other nor are they equally good. Incommensurable items might be on a par, which is a relatively innocuous form of incommensurability, but they might also be more radically incomparable. In what follows, we only need to assume that *a*+ is better than *a* but is not better than *b*, while *b*+ is better than *b* but is not better than *a*. To give an example, suppose that *a* and *b* are two attractive holiday trips to very different locations, say, *a* is a trip to Galapagos and *b* is a trip to Peru. *a*+ and *b*+ are the same trips with a small discount. The discount suffices to make $a+$ better than a and $b+$ better than b, but it is not significant enough to make $a+$ better than b and $b+$ better than a.

Let *S1* and *S2* be two equiprobable states of the world ('states of nature') that in themselves are axiologically neutral (i.e. they are devoid of any value, either positive or negative) and, furthermore, do not contribute value (positive ot negative) to anything else. Suppose that, for example, a fair coin will be tossed and the states are the two possible results of the toss: *S1* stands for Heads and *S2* for Tails. We assume that Heads and Tails both have the same objective probability .5.

I will interpret actions in Savage's manner, as assignments of outcomes to states. This way of understanding actions seems unproblematic as long as states are causally and probabilistically independent of actions, which I am going to assume in what follows.

Consider two actions, *X* and *Y*:

	S1	*S2*
X	$a+$	$b+$
Y	b	a

As the diagram shows, *X* yields $a+$ under state *S1* and $b+$ under *S2*, while *Y* yields b under *S1* and a under *S2*.

Consider, to begin with, a principle that all parties to the debate should be willing to accept:

Dominance (V): (i) one action is better than another if it under every state yields a better outcome than the other action, and (ii) it is at least as good as another action if it under every state yields an outcome that is at least as good as that of the other action.

Here, (V) stands for *value*, in order to distinguish this principle from its analogue that applies to rational preferences.

Now, consider a principle that seems close in spirit to Dominance (V):

Complementary Dominance (V): One action is not better than another if it under every state yields an outcome than is not better than the outcome of the other action.

Or, to put is shorter: One action is not better than another if there is no state under which it has a better outcome than the other action.[1]

While the first clause of Dominance (V) relates betterness of outcomes to the betterness of actions, Complementary Dominance postulate the exactly analogous relationship between outcomes and actions with respect to the *complement* of betterness: the relation of not being better than.

How close this principle is to Dominance (V) can be appreciated if one considers that, in the absence of incommensurabilities, "not better" is equivalent to "at most as good", i. e., to the

---

[1] It is easy to see that Complementary Dominance entails the following principle:

<u>Neither Nor (V)</u>: One action is neither better nor worse than another if there is no state under which it has a better or worse outcome than the other action.

In what follows, I am going to argue against Complementary Dominance (V), but my argument will also apply to Neither Nor (V).

In Schoenfield's decision-theoretical variant of the problem, the principle that corresponds to the conjunction of Dominance (V) with Neither Nor (V) is formulated as follows:

> LINK: In cases in which considerations of value are the only ones that are relevant, if you are rationally certain that one option, A, will bring about greater value than the alternative option, B, you're required to choose A. If you are rationally certain that neither of the two options will bring about greater value than the other, it's not required that you choose A, and it's not required that you choose B. (Schoenfield, p. 267)

Schoenfield's Link, which she defends in her paper, connects rational choice with evaluation. But in a footnote, she adds: "if you prefer to think of rational decision making as determined fully by the agent's preferences rather than objective facts about value, simply substitute 'A will bring about greater value than the alternative B' with 'A will bring about an outcome that is preferred to B'." (ibid, p. 276)

Bales, Cohen & Handfield (2014) formulate and defend a principle that is almost identical to Link, but applies not merely to binary choices. They call it Competitiveness and consider it to be "at least as plausible" as the requirement to perform the dominating action rather than the one that is being dominated (ibid. p. 460). To introduce it, they first introduce the notion of a *competitive action*, which they define as an action such that "for every way the world could be [i.e.., for every state of nature], its consequences [i.e.., its outcome] are no worse than the consequences of all alternative actions." (ibid.) The principle of Competitiveness states that it is rationally permissible to perform a competitive action. (ibid.)

converse of “at least as good”. Consequently, in the absence of incommensurabilities, Complementary Dominance (V) would reduce to the second clause of Dominance (V). But, in the presence of incommensurabilities, as in the case we consider, Complemenaty Dominance (V) represents an independent constraint.

Now, Complementary -Dominance (V) entails that <u>*X* is not better than *Y*</u>. Given or assumptions about the case, there is no state under which the outcome of *X* is better than that of *Y*.[2,3]

However, there is also an argument for the opposite conclusion, namely, that <u>*X is* better than *Y*</u>. This is, by the way, the conclusion, that I myself accept. So I need to argue for it, which I will do in the rest of this section and the next one. But I also need to explain what is wrong with Complementary Dominance (P). This will be done in sections 4 and 5, after some necessary preparations in section 3.

Here then is how an argument for *X* being better than *Y* can be constructed. I will first start with an argument that rests on very strong assumptions and will then show how these assumptions can be considerably weakened.

---

[2] That *X* is not better than *Y* would also follow if we were to accept the following principle that relates incommensurability of outcomes to incommensurability of actions:

Incommensurability: Two actions are incommensurable in value if they yield incommensurable outcomes under every state.

This principle is very similar in form to the following implication of Dominance (V):

Equal Goodness: Two actions are equally good if they yield equally good outcomes under every state.

Incommensurability does not entail Complementary Dominance (V) nor is it entailed by the latter, but the two principles are closely related. (This is perhaps best seen if one compares Incommensurability with Neither Nor (V), which *is* entailed by Complementary Dominance.)

[3] Hare (2010), who does not discuss the value relations but instead focuses on rational preferences and rational choice, considers (but in the end rejects) a principle of *deference* to one’s “better informed self” (p. 242). In a comparison between two actions such as *X* and *Y*, my better informed self, who would know which state of nature obtains, would have no preference in favour of any of them, whatever the state: His preferences for actions would go by their actual outcomes and I know that he would not prefer the outcome of one action to the outcome of the other. But then – by the principle of deference – it is rationally permissible for me to have no preference between these actions either, even though, being less informed, I don’t know which state obtains. This would entail that, in the choice between two such actions, it would be permissible for me to choose either.

If states of nature are assigned objective probabilities, as we have assumed, we can compare actions in terms of the *lotteries* on outcomes they give rise to. To say that an action x gives rise to a lottery l, or, what amounts to the same, that l is determined by x, means that l is the probability distribution over outcomes that to each outcome assigns the sum of the probabilities of the states that deliver this outcome given x. We can refer to the lottery determined by x as l^x.

X gives rise to a lottery ($a+$,.5; $b+$,.5) while Y gives rise to a lottery (b,.5; a,.5), or, what amounts to the same, to a lottery (a,.5; b,.5). (The ordering in which the outcomes and their probabilities are specified in a description of a lottery doesn't matter for the lottery's identity.)

The von Neumann –Morgenstern expected utility theory assumes that preferences on lotteries satisfy the Independence axiom: replacing an outcome in a lottery by another one that is more preferred results in a more preferred lottery. A corresponding principle for value comparisons between lotteries would be as follows:

Independence (V): For all outcomes α, β, γ,
for all $p > 0$, α is better than β iff (α, p; γ, 1-p) is better than (β, p; γ, 1-p);
for all $p < 1$, α is better than β iff (γ, p; α, 1-p) is better than (γ, p; β, 1-p).

According to this principle, lottery prizes (outcomes) make independent value contributions to the value of a lottery: Improving a prize in a lottery always improves the lottery itself.

Since $a+$ is better than a and $b+$ is better than b, Independence (V), twice applied, implies that ($a+$,.5; $b+$,.5) is a better lottery than ($a+$,.5; b,.5) and the latter lottery is better than (a,.5; b,.5). Since betterness is transitive, it follows that ($a+$,.5; $b+$,.5) is better than (a,.5; b,.5).

Suppose we now assume:

Reduction to Lotteries (V): For all actions x and y, x is better than (equally as good as) y iff l^x is better than (equally as good as) l^y.

According to this principle, the value of an action goes by the value of the associated lottery.

It would then follow from Independence (V) and Reduction to Lotteries (V) that <u>action X is better than action Y</u>, contrary to what we have previously argued for.

An argument for *X* being better than *Y* can in fact be made in another way as well, without relying on such strong principles as Reduction to Lotteries and Independence. (As is well known, the latter principle leads to serious troubles, as exemplified by the famous Allais' Problem.) Instead of these principles, we could appeal to Dominance (V), which we want to accept anyway, together with the following weakening of Reduction to Lotteries:

Weak Reduction to Lotteries (V): If two actions give rise to the same lottery, they are equally good.

Thus, consider adding to *X* and *Y* a third action *Z*:

	S1	*S2*
X	*a+*	*b+*
Y	*b*	*a*
Z	*a*	*b*

Y and *Z* give rise to the same lottery, (*a*,.5; *b*,.5). (Remember that the ordering of outcomes in the description of a lottery doesn't matter for the lottery's identity.) Thus, *X* and *Y* are equally good by Weak Reduction to Lotteries (V). But, by (the first clause of) Dominance (V), *X* is better than *Z,* since under each state *X*'s outcome is better than *Z*'s. This means that *X* must also be better than *Y*, if betterness is transitive across equal goodness.[4]

[4] A somewhat different argument, which nevertheless also appeals to (a version of) Dominance, is provided by Hare (2010). He introduces the notion of *prospects* associated with an action, where a prospect is a possible outcome of that action paired with the agent's credence that the action will yield the outcome sin question. The set of all prospects associated with an action is thus very much like a lottery, but with objective probabilities replaced by credences. Hare then assumes that

Prospects Determine Permissibility: Facts about what it is rationally permissible for me to do are determined by facts about the prospects associated with the options available to me. (ibid. p. 240)

Clearly, this principle is similar to Reduction to Lotteries.
Hare now considers two choice situations: one in which we choose between *X* and *Y* (with objective probabilities replaced by credences) and the other in which the choice is between *X* and *Z*. Implicitly assuming (a preferential

Indeed, this detour via lotteries is not needed to establish the result we are after. We can instead base our argument on a principle that is even less demanding than Weak Reduction to Lotteries. Let π be any permutation on states. Consider any action x. An action y is the *image* of x under π iff for all states S, x assigns to S the outcome that y assigns to $\pi(S)$. To put it differently, y results from x by a reassignment of outcomes among the states. Now, consider the following principle:

Permutation (V): If π is a permutation on states that to every state assigns an equiprobable state, then for all actions x, x is equally as good as its image under π.

In other words, *a reassignment of outcomes among equiprobable states does not affect the value of an action*.

Now, it is easy to see that action *Z* results from *Y* by such a reassignment. Thus, Permutation (V) implies that *Y* and *Z* are equally good and therefore again, by (the first clause of) Dominance (V) and the transitivity of betterness over equal goodness, it follows that *X* is better than *Y* since *X* is better than *Z*.[5,6]

---

version of) Dominance as a criterion of permissibility, he draws the conclusion that *Z* is impermissible in the latter choice problem. But since the prospects associated with *Y* and *Z* are the same and since in both problems the alternative action is the same (*X*), it follows that *Y* must be impermissible in the former choice problem if Prospects Determine Permissibility.

[5] Bales, Cohen & Handfield (2014) pose an objection to Hare's dominance-based argument for the impermissibility of *Y*. This objection can be adapted in order to question my argument as well. Here is how the adapted objection would look like: As Bales *et al*. point out, action *Z* would still be dominated by *X* even if states had unequal probabilities. However, if the states' probabilities were unequal, we could no longer use Permutation (V) (nor Weak Reduction t Lotteries) in order to establish that *Y* and *Z* are equally good. We would thus be unable to establish that *X* is better than *Y*.

I don't find this objection at all worrying. If the states' probabilities were unequal, then one should not expect *X* to be better than *Y*. It would under these conditions be perfectly possible that *X* and *Y* are incommensurable. Indeed, the modelling of value relations that I am going to develop below bears this out.

[6] Note that adherents of Dominance (V) will not be impressed by the above argument. They would of course reject one of its premises, Permutation (V). Since neither a is better than b nor vice versa, Complementary Dominance implies that neither Y is better than Z nor Z is better than Y. In other words, from this perspective, Permutation (V) must be an incorrect principle. (The same of course applies to strengthenings of Permutation (V), such as Reduction to Lotteries or Weak Reduction to Lotteries.) This shows that we are not home yet, as long as we haven't provided a

2. *The same outcomes?*

We need to consider a potential objection to the argument from Permutation (or from Weak Reduction to Lotteries for that matter). This argument assumes that the outcome of *Y* under *S1* is the same as that of *Z* under *S2*, and – analogously – that the outcome of *Y* under *S2* is the same as the outcome of *Z* under *S1*. But this assumption might be questioned.

One way in which one might question it is by pointing out that a more complete description of an outcome should incorporate the state under which the outcome is supposed to obtain. Thus, a more complete description of the outcome of, say, *Y* under *S1* is '*b* in state *S1*", while a more complete description of the outcome of *Z* under *S2* is '*b* in state *S2*". If states are incorporated into outcomes in this way, the two outcomes are not the same. Thus, *Y* is not obtained from *X* by a mere reassignment of outcomes. (And the lotteries determined by *X* and *Y*, respectively, are not identical.)

But this version of the objection is unsatisfactory. It is perfectly acceptable to exclude from the outcomes those components or features that are axiologically irrelevant. And we have supposed that the states in our example are irrelevant in this way. We have therefore no reason to incorporate them into the outcomes. Remember that *S1* and *S2* are just alternative results of a toss of a fair coin: Heads and Tails, respectively. While the result of the toss determines the outcome of a given action, it has been assumed to make no contribution to the value of the outcome it determines.

Another version of the objection distinguishes between seemingly identical outcomes by taking into consideration their *modal* features. In particular, when we evaluate the outcome of an action, one relevant consideration might be *what the action could have brought about if the state of nature had been different*. Think of an action, for example, which results in winning a large sum of money. For the evaluation of this outcome it might play a role what the action could have resulted in if the state of nature had been different. Could the agent have ended up with nothing or would he have won the money under every state? Such a guarantee might make the win more valuable for some (people who are risk-aversive) and less valuable for others (risk seekers). What

direct argument against Complemenary Dominance (V). I will try to do it later in this paper. I will then also provide reasons for accepting Permutation (V).

might have been the case under other states is thus potentially relevant to the value of the outcome in a given state; and, if it is relevant, then it should be included into the full description of the outcome in question.

But this kind of appeal to what might have been doesn't allow us to distinguish between outcomes in the problem we are considering. Thus, consider receiving *b* under *S1*, which is the outcome of action *Y* under that state. Under another state, the action in question would have resulted in *a*. The situation is exactly similar concerning *b* as the outcome of action *Z* under state *S2*: Under another state, the action would have resulted in *a*. So, in this respect, there is no difference between *b* as brought about by *Y* and *b* as brought about by *Z*. The same applies, mutatis mutandis, to outcome *a*, as brought about by *Y* or by *Z*.

The appeal to the modal features of outcomes might, however, also take another form: In evaluating the outcome of an action, it is one thing to consider what the action could have resulted in if the state had been different. But we might also consider what could have happened under the same state *if another action had been performed*. The latter kind of modal consideration looms large in Allais' problem. There, in one choice situation, if the agent chooses a risky action and loses (i.e., ends up with nothing), he would have won (gained a large sum of money) if he had acted otherwise: indeed, the alternative action is not risky. While in the other choice situation, both actions at the agent's disposal involve a risk of loss, though in slightly different degrees. There is a possible state in which the agent loses and would still have lost if he had acted otherwise. It might be argued that losing is less bad for the agent if it couldn't have been prevented by acting otherwise. Under such circumstances there is no reason to regret what one has done and the absence of such reasons might make the outcome more bearable.[7]

In the problem we face there does exist a modal difference of this kind between otherwise identical outcomes of *Y* and *Z* – a difference that *might* be value-relevant. Thus, the outcome *b* of *Y* under *S1* is not worse than the outcome at that state of any other action at the agent's disposal. Actions *X* and *Z* would have under *S1* resulted in *a+* and *a*, respectively, and none of these outcomes is better than *b*. By contrast, the outcome *b* of *Z* under *S2 is* worse than the outcome of

[7] For a discussion of different kinds of potential regret and in particular for the distinction between *outcome-regret* relating to what a given action could have resulted in instead and *action-regret* relating to what could have happened if one instead chose another action, see Bovens & Rabinowicz (2015).

X under the same state: Under *S2*, *X* would have resulted in *b+*. Because of this modal difference then, it is arguable that outcome *b* of *Z* under *S2* is worse than outcome *b* of *Y* under *S1*. The same applies to *a*. Outcome *a* of *Z* under *S1* is arguably worse than outcome *a* of *Y* under *S2*: In *S1*, but not in *S2*, if the agent instead had chosen *X*, he would have brought about a better outcome, *a+*. This forms a basis for the argument that action *Z* is worse than action *Y*: Each of its outcomes is worse than the seemingly identical outcome of *Y*. Consequently, even though *X* by Dominance (V) is better than *Z*, it doesn't follow that *X* is better than *Y*.

The objection just presented is controversial, since it is debatable whether and to what extent the modal features of outcomes can be value-relevant. However, there is a way to finesse the objection instead of meeting it head on: we can simply change the example. Thus, let us suppose that there is a fourth action available to the agent, apart from *X*, *Y*, and *Z*. This fourth action; *U*, is the image of *X* under the same permutation of equiprobable states under which *Z* is the image of *Y*:

	S1	*S2*
X	*a+*	*b+*
Y	*b*	*a*
Z	*a*	*b*
U	*b+*	*a+*

By introducing *U* we make *Y* and *Z* perfectly symmetrical in their modal features, thereby also making their respective outcomes modally identical. Just as for *Z* it is under every state the case that the agent who performs *Z* would have brought about a better outcome by instead performing another action (action *X*), so it is for *Y* in our modified example: if the agent performs *Y*, he would under every state have brought about a better outcome by instead performing another action (action *U*). Thus, we can now no longer argue on modal grounds that *Y* and *Z* have different outcomes.

If we cannot argue for this on any other grounds either, then we can apply Permutation (V) to the comparison between *Z* and *Y* and conclude that these two actions are equally good.. In combination with Dominance (V) and the transitivity of betterness across equal goodness we can then derive that *X* is better than *Y* and thus get back the problem we have started with. For even in the modified example, with four available actions, Complementary Dominance (V) still entails that *X* is *not* better than *Y*.

Thus, if we want to hold on to Permutation (V), then we need to take a closer look at Complementary Dominance (V). Is this principle as compelling as it seems to be at first sight?

3. Value relations analysed

To approach this issue, it is helpful to ask how value relations should be analysed. In earlier work (Rabinowicz 2008), I proposed such an analysis in the spirit of the Fitting-Attitudes account of value. On this nowadays highly influential account, value is understood in terms of pro- and con-attitudes that it is fitting to hold towards a value bearer. The statement that an item *i* is valuable is understood as the claim that it is fitting to have a pro-attitude towards *i*. Or, to put it in stronger normative terms, *i* is valuable insofar as one ought to favour *i*.[8] Here, "favour" is a place-holder for a pro-attitude. Different values might call for different kinds of favouring. In case of value relations, such as betterness, a natural suggestion is that the relevant favouring consists in preference: an item *i* is better than an item *j* iff one ought to prefer *i* to *j*. Analogously, *i* and *j* are equally good iff one ought to equi-prefer *i* and *j*, i.e., to be indifferent between them. Consequently, *i* and *j* are incommensurable iff one neither ought to prefer one to the other nor ought to be indifferent. The kind of preferences and indifferences that are relevant in this context and not unstable, flighty attitudes. Betterness and equal goodness are analysed in terms of requirements on considered preferences and indifferences.

Suppose that **I** is the domain of items that are being compared and **K** is the class of all permissible preference orderings of that domain. We assume that in every ordering in **K**, weak preference (i.e. preference-or-indifference) is a reflexive and transitive relation. We can now define different value relations on **I** in terms of **K**. For any two items *i* and *j* in **I**, *i* is better than *j*

[8] This formulation ignores various qualifications and provisos we might want to impose, such as, for example, that the requirement to favour *i* only applies to agents who know what *i* is like or the qualification that the ought in question should be understood more in the spirit of a recommendation than as a strict demand. In what follows I disregard these complications.

iff in every ordering in **K**, *i* is preferred to *j*; which means that the preference for *i* over *j* is not merely permissible but required. *i* is equally as good as *j* iff in every ordering in **K**, *i* and *j* are equi-preferred; which means that the equi-preference of *i* and *j* is required. These definitions imply that the relation of betterness-or-equal goodness is transitive and reflexive (because weak preference has been assumed to be transitive and reflexive). Which, in its turn, entails (via the standard definitions of preference and indifference in terms of weak preference) that betterness is transitive and asymmetric, that equal goodness is an equivalence relation, and that betterness is transitive across equal goodness, i.e., that *i* is better than *k* if *i* is better than *j* and *j* and *k* are equally good. Just as we were assuming in our argument that *X* must be better than *Y*: this follows if *X* is better than *Z* and *Z* and *Y* are equally good.

Given this modelling of value relations in terms of class **K** of permissible preference orderings, *i* and *j* are incommensurable iff (i) there are orderings in **K** in which *i* is not preferred to *j*, (ii) there are orderings in **K** in which *j* is not preferred to *i*, and (iii) there are orderings in **K** in which *i* and *j* are not equi-preferred.
(i) – (iii) imply, respectively, that *i* is not better than *j*, that *j* is not better than *i*, and that *i* and *j* are not equally good.

Note that this necessary and sufficient condition of incommensurability is satisfied if (though not only if) the following holds:

K contains two preference orderings, such that *j* is preferred to *i* in one ordering and *i* is preferred to *j* in the other.

In Rabinowicz (2008), I suggested that the latter condition can be used to define a very common form of incommensurability, which Ruth Chang calls 'parity'. On this suggestion, then, two items are on a par iff it is permissible to prefer one of them to the other but it also is permissible to have the opposite preference. Cases like this arise when the comparison between two items is based on weighing several relevant dimensions, or aspects, against each other. One item can rate higher than another on some of the dimensions and lower on other dimensions. Different assignments of weights to dimensions can often be admissible in cases like this and different permissible weight assignments give rise to different permissible preference orderings of items. If one item comes higher up on one admissible assignment of weights, while the other on another admissible weight assignment, then it is permissible to prefer the former item to the latter, but it also is permissible to have the opposite preference. We thus have a case of parity.

4. *Deconstructing Complementary Dominance (V)*

Now, consider what this analysis of value relations in terms of permissible preferences implies for the problem at hand. If Dominance (V) and Permutation (V) are to be valid, then their preferential variants must hold for every ordering in class **K**:

Dominance (P): (i) One action is preferred to another if it under every state yields a more preferred outcome than the other action; and (ii) it is weakly preferred (i.e., either preferred or equi-preferred) if it under every state yields an outcome that is weakly preferred to that of the other action.[9]

Permutation (P): Every action is equipreferred with its image under permutation of equiprobable states.[10]

If Complementary Dominance (V) is to hold, **K** would have to satisfy the following, rather clumsy condition:

(*) If for every state *S* there is some ordering in **K** in which the outcome of action *x* in *S* is not preferred to the outcome of action *y* in *S*, then in some ordering in **K**, *x* is not preferred to *y*.

Note that this condition, unlike Dominance (P) and Permutation (P), is not a constraint on every preference ordering in **K**, but instead a condition on class **K** as a whole: It requires **K** to contain some orderings if it contains some other orderings.

---

[9] The first clause of Dominance (P) is both necessary and sufficient for the corresponding first clause of Dominance (V). However, the second clause of Dominance (P) is only necessary, but not sufficient, for the second clause of Dominance (V). Even if the outcome of one action in a given state is weakly preferred to the outcome of the other action in every ordering in K, it might not be preferred to it in every such ordering, nor equi-preferred with it in every such ordering. It might be preferred to it in some orderings and equi-preferred in others. But this means that it does not have to be at least a good, i.e., better or equally as good, as the outcome of the other action in the state in question .

[10] This condition is both necessary and sufficient for Permutation (V).

While (*) is both necessary and sufficient for Complementary Dominance (V), seems to lack intuitive support. To see this, consider a simple model that represents a case in which, as we have assumed, *a+* is better than *a*, *b+* is better than *b*, but *a+* is not better than *b*, nor is *b+* better than *a*. In this model, class **K** consists of just three preference orderings, *P1*, *P2* and *P3*, specified below. (Indeed, the inclusion of *P3*, although rather natural, is not necessary: it would suffice with *P1* and *P2*.)[11]

<u>*P1*</u>	<u>*P2*</u>	<u>*P3*</u>
a+	*b+*	*a+*, *b+*
a	*b*	*a*, *b*
b+	*a+*	
b	*a*	

While *P1* and *P2* linearly order the four items, *P3* contains two ties: *a+* and *b+* are equi-preferred and likewise *a* and *b*. Since *P1* - *P3* are the only orderings in **K**, it follows that *a+* is better than *a* (it is preferred to *a* in every ordering in **K**), *b+* is better than *b* (for the same reason), but *a+* is not better than *b* (*b* is preferred to it in *P2*) , nor is *b+* better than *a* (*a* is preferred to it in *P1*), just as we wanted.

The orderings in **K** are supposed to extend to all the items in the domain **I**. (This means, by the way, that *P1* - *P3* are not really single orderings, but rather representatives of sets of permissible orderings, since it is to be expected that other items in the domain can be permissibly ordered in many different ways. In what follows I ignore this complication, as it doesn't affect the argument.) If **I** also contains such items as actions, we need to say something about how *P1* - *P3* order actions *X*, *Y*, *Z*, etc. Permissible preferences satisfy Dominance (P), which means that all the three orderings rank *X* above *Z*. *If* permissible preferences also satisfy Permutation (P), then in all the three orderings, *Y* and *Z* are equi-preferred. By the transitivity of permissible preference

[11] The argument to follow would still hold if we added more orderings to **K**, provided that in all of them *a+* is preferred to *a* and *b+* is preferred to *b*. Otherwise, they might rank the four items differently from how these items are ranked in *P1* - *P3*. In the first part of the argument, it is also assumed that all the orderings in **K** are complete (i.e., contain no preference gaps). I am going to say more about incomplete preference orderings later.

over equi-preference, it therefore follows that *P1*, *P2* and *P3* all rank *X* above *Y*. Which means, on our analysis, that *X* is better than *Y*.

This conclusion comes into conflict with Complementary Dominance (V), which entails that *X* is not better than *Y*. But how compelling is that principle, or –what on our analysis amounts to the same – how compelling is its preferential version (*)?

The answer is: Not compelling at all. The model described above can serve as an example. While it is true in this model that *X* under no state has a better outcome than *Y*, i.e., that there is for every state a preference ordering in **K** in which the outcome of *X* is not preferred to the outcome of *Y* (under *S1*, *P2* has this feature, while *P1* has it under *S2*), it still is true that in every preference ordering in **K** at least one of the possible outcomes of *X*, if not both, is preferred to the corresponding outcome of *Y*. (This applies in *P1* to the outcome of *X* under *S1*, in *P2* to *X*'s outcome under *S2*, and in *P*3 this holds for *X*'s outcomes under both states.) Thus, in each of these permissible orderings, the preference for *X* over *Y* is supported by the preference for at least one of *X*'s possible outcomes over the corresponding outcome for *Y*. Admittedly, in orderings *P1* and *P2* (though not in *P3*), the opposite preference for *Y* over *X* is also supported by the preference for some outcome of *Y* over the corresponding outcome of *X*. But this support for *Y* is weaker than the support for *X*. To see this, consider *P1* first. In this ordering, the preference for *X* is supported by the preference for *X*'s outcome under *S1*: $a+$ is preferred in *P1* to *b*. While the preference for *Y* is supported by the preference for *Y*'s outcome under *S2*: *a* is preferred in *P1* to $b+$. Now, it is easy to see that this support for the preference for *Y* is weaker - less pronounced – than the support for the preference for *X*: In a very intuitive sense, the preferential distance between $a+$ and *b* is in *P1* greater than the distance between *a* and $b+$. The former items occupy the first and the fourth place, respectively, while the latter occupy the second and the third place. Thus, in every possible assignment of numerical values to items that represents *P1*, the difference between the numerical values assigned to $a+$ and *b* must be larger than the corresponding difference between the numerical value assigned to a and $b+$, simply because of their placement in the ordering. In this sense, then, he 'ordinal distance' between $a+$ and *b* is in *P1* greater than that between *a* and $b+$. The argument with respect to *P2* is exactly analogous: In that ordering as well, the support that the preference for *Y* receives under one state is weaker (in terms of 'ordinal distance') than the support that the preference for *X* gets under the other state. But then, since the states are equiprobable, it should not come as a surprise that *X* is preferred to *Y* in both *P1* and *P2*, and indeed in every ordering in **K**. Indeed, this argument can be generalized for all **K** that consist of complete preference orderings. In any complete preference ordering in which (1) $a+$ is ranked above *a* and (2) $b+$ is ranked above *b*, the support for the preference for *X* is stronger (in terms of 'ordinal distance') than that for *Y*.[12,13] The feeling of mystery initially attaching to the claim that *X* is better than *Y* therefore disappears.[14, 15]

---

[12] Indeed, *P1* and *P2* are the only such orderings in which the preference for *Y* receives *any* support at all.

[13] Note that a similar argument can be used to defend Permutation (P), which – as we know – is a necessary and sufficient condition for Permutation (V). To see how it works consider, for simplicity's sake, how this works for the comparison between *Y* and *Z*. (This argument generalizes to all applications of Permutation (P).) In any preference ordering, either (i) outcomes *a* and *b* are tied and then the ordering will rank *Y* and *Z* equally, or (ii) one of these outcomes is ranked above the other. Suppose then that *a* is ranked above *b*. (The opposite ranking can be dealt with in the same way.) Then the outcome of *Y* is (1) preferred to that of *Z* under *S2* and (2) dispreferred to the outcome of *Z* under *S1*. But the support that the preference for *Y* over Z gets from (1) is exactly equal to the support the prefererence for *Z* over *X* gets from (2): the 'ordinal distance' from *a* to *b* is of course the same as that from *b* to *a*. Since *S1* and *S2* are equi-probable, it again follows that *Y* and *Z* will be equi-preferred in any such ordering, just as in orderings considered in case (i).

[14] Note, however, that the argument above depends on the equiprobability of states. If, say, *S1* were more probable than *S2*, the support for the preference for *Y*, while less pronounced in terms of ordinal distance, would gain *weight* in P2 and the support for the preference for X, while more pronounced, would lose weight. This could lead to *Y* being placed above *X* in *P2*. In *P1*, on the other hand, the preference for *X* would get more weight and thus would even more easily outbalance the less pronounced preference for *Y*. As a result, if the states had unequal probabilities, it could well turn out that actions *X* and *Y* are *on a par*, with each of them being preferred to the other in some permissible preference ordering.

[15] Schoenfield (2014) considers, but finally rejects, a somewhat similar argument in favour of *X* being better than *Y*. According to that argument, if an agent's value ordering is incomplete, but there is an action that would be unanimously recommended by every possible completion of her value ordering, then she ought to perform the action in question. She calls this principle *Unanimity*. Schoenfield assumes that every complete value ordering would have to obey axioms of expected utility theory, which is a very strong assumption, by the way. She therefore concludes that every complete value ordering would assign higher value to *X* than to *Y*. Consequently, Unanimity would imply that the agent ought to opt for *X*. While this argument, like the one I have proposed above, involves a move from an incomplete value ordering to a class of complete orderings, it interprets the latter as *value* orderings and not as permissible orderings of preference. The move to preference orderings is in my approach motivated by the analysis of value relations in terms of permissible preferences. By contrast, the move to complete value orderings in the argument considered by Schoenfield seems quite unmotivated (and she is right to reject it). Indeed, such a move

5. *Incomplete preference orderings*

There is, however, a natural line of defense for someone who wants to question this conclusion. In the argument above, we have been assuming that all the preference orderings in **K** are complete. But this might of course be questioned. Indeed, in Rabinowicz (2008), I do not restrict class **K** to complete orderings: I also allow orderings that contain preference gaps. Now, it might well be permissible to have an incomplete preference ordering of the four items $a+$, a, $b+$ and b – an ordering in which $a+$ is preferred to a, $b+$ is preferred to b, and no other preference relations obtain between these four items. In this incomplete preference ordering, none of the possible outcomes of X is preferred to the corresponding outcome of Y and thus the preference for X over Y is unsupported. A critic of my argument could therefore insist that if this permissible incomplete preference ordering is extended to actions, then such an ordering is not going to rank X above Y: Instead, it will contain a preference gap with respect to these two actions. Consequently, it will follow that X is not better than Y.

This is certainly a weak point in my argument. But can I somehow respond to the critic? I think so. Note that gaps in preference can potentially be filled in. In other words, an incomplete preference ordering can – at least in principle – be completed, though of course it can have many different completions.[16] Now, consider a permissible incomplete preference ordering P and the set C_P of all its permissible completions. To be more precise, a preference ordering P^* belongs to C_P iff P^* (i) is complete, (ii) belongs to **K**, and (iii) includes P: for all i an j in **I**, if i is preferred to/equi-preferred with j in P, then it is also preferred to/equi-preferred with j in P^*. I will assume that, for any permissible P, C_P is a non-empty set: For a permissible preference ordering there is always a permissible completion.[17] Now, consider $\cap C_P$, the intersection of C_P. $\cap C_P$ is the

---

seems positively wrong-headed. Remember that the agent considers the items compared to be incommensurable in value. If this is correct, then filling in such gaps in value ordering is wrong. Judging that two items are incommensurable is obviously quite different from recognizing that one doesn't know how they compare in value. A knowledge gap allows of a completion, at least in principle. But incommensurability excludes completion.

[16] In this respect, preferential gaps importantly differ from judgments of incommensurability. See the preceding note.

[17] But what if this assumption is not satisfied? What if some items i and j in the domain **I** are radically incomparable, in value, in the sense that in all permissible preference orderings there is a gap with respect to i and j? To guarantee the non-emptiness of C_P we would then have to re-define this notion: to define it as the set of all *maximal permissible*

preference ordering that collects all preferences which are common to the permissible completions of *P*. I'd like to suggest that a preference ordering *P* is permissible only if it is identical to $\cap C_P$. I will say that *P* in such a case is *well-rounded*. What I want to suggest then is that, while K might well contain incomplete preference orderings, all the orderings in K are well-rounded.

If a preference ordering is not well-rounded, then there is something 'half-finished' about it: It contains a preference gap at a point at which no such gap should be present, given that all the permissible completions of the ordering in question would fill the gap in exactly the same way. Think, for example, of an ordering *P* that contains a preference for *i* over *j* and a preference for *j* over *k*, but no preference as regards *i* and *k*. It contains this gap even though all completions of *P* would contain a preference for *i* over *k*. Clearly, there is something inadequate about such an ordering. To the extent that permissible preference orderings should be rational, it is a fair requirement that they should all be well-rounded.[18]

Now, if we accept this constraint on **K**, then it no longer is a problem for my argument that K might contain incomplete orderings. Since, as I have argued, all permissible complete orderings contain a preference for *X* over *Y*, the same must hold for all permissible incomplete orderings as well: Being well-rounded, they inherit this property from their permissible completions. Which means that the critic's objection fails. Complementary Dominance (V) is not a valid principle even if incomplete preference orderings are allowed as members of **K**.

At this point, let me digress a little. Complementary Dominance (V) has an analogue for permissible preferences. According to this preferential analogue, the following must hold for all orderings in **K**:

Complementary Dominance (P): One action is not preferred to another if there is no state under which its outcome is preferred to that of the other action.

---

extensions of *P*. *P** is such an extension of *P* iff *P** belongs to **K**, includes *P*, and is not included in any other element of **K**. I will return to the issue of (radical) value incomparability below.

[18] This reminds a little about a condition on an ideal set of beliefs. While such a set might be incomplete (an agent might lack opinion on some matters), it should coincide with the intersection of all its permissible completions. This well-roundedness constraint guarantees that an ideal belief set will be closed under logical consequence and under other general requirements we might want to impose on all complete belief sets.

Unlike (*), this principle is neither necessary nor sufficient for Complementary Dominance (V).[19] But it appears to be similarly compelling. Are the appearances misleading even in this case? The answer is Yes, but it should be noted that Complementary Dominance (P) is perfectly innocuous as long as the preference ordering under consideration is *complete*. With respect to such an ordering, Complementary Dominance (P) immediately follows from Dominance (P). <u>Proof</u>: In a complete preference ordering, an item *i* is not preferred to an item *j* iff *j* is weakly preferred to *i*. Therefore, if there is no state under which the outcome of an action *x* is preferred to that of an action *y*, the outcome of *y* is weakly preferred to the outcome of *x* under every state. But then, by the second clause of Dominance (P), *y* is weakly preferred to *x*, which implies that *x* is not preferred to *y*.

If, however, preference orderings are allowed to be incomplete, then Complementary Dominance (P) no longer holds. As we have seen, in the example we are focusing on, in an incomplete preference ordering in which there is no state under which the outcome of *X* is preferred to that of *Y*, *X* still is preferred to *Y* if the ordering in question is well-rounded.

6. *Incomparability*

After this digression, let me come back to the main issue. Do we know now how incommensurability interacts with risk? No, of course not, there is much more that would require clarification regarding this general topic. But have I at least provided a satisfactory solution to the problem we have started with? Is the action such as *X* definitely better than its competitor *Y*? Well, perhaps not in all possible cases. The kind of case I have focused on is one in which outcomes *a+* and *a* are *on a par* with outcomes *b+* and *b*: It is permissible to prefer the former outcomes o the latter but it also is permissible to have the opposite preference. Now, parity is a

[19] It is not sufficient, since the antecedent of (*) is weaker than the antecedent of Complementary Dominance (P): it does not require the existence of any preference ordering *P* in **K** such that under no state the outcome of one action, *x*, is preferred to the outcome of the other action, *y*. Thus, Complementary Dominance (P) cannot be used to prove the consequent of (*) from its antecedent. Nor is Complementary Dominance (P) necessary for (*), since the consequent of (*) is relatively weak: It only requires that there exists some ordering in **K** in which action *x* is not preferred to action *y*. We could have a model in which some such ordering exists, even though **K** also contains *another* ordering that violates Complementary Dominance (P). Thus, (*) could hold even in the absence of Complementary Dominance (P). (As will be seen below, however, a model that violates Complementary Dominance (P) would have to allow incomplete preference orderings.)

form of incommensurability, the most common one, I would say. It is exemplified by such comparisons as the one between holiday trips, or by Caspar Hare's comparison between different restaurant options: Indian versus Chinese. Possibly, it also is exemplified by his other example, in which I have to choose which object should be saved from the fire: my Fabergé egg or my wedding album. (Cf. Hare 2010, pp. 237f)[20]

But there might also, at least in principle, exist cases of incommensurability in which it is *not* permissible to prefer one item to the other or to be indifferent. Or, at least, this requirement can apply to considered preferential attitudes: Both considered preference and considered equi-preference might be disallowed. In such cases what is required is a preferential gap regarding the items compared. In Rabinowiczy (2008), I referred to such form of incommesurability as (radical) *incomparability* in value. Clearly, incomparability might well obtain between items that belong to different ontological categories. Thus, to take an example, persons are incomparable in value to, say, events or properties. But if the compared items are ontologically of the same kind, if they are, as in the cases we consider, different possible outcomes, then I think it is fair to say that incomparability is a much rarer phenomenon than parity. Still, at least in principle, its existence cannot be excluded.

A possible example in which we encounter incomparability is provided by Bales, Cohen and Handfield (2014): The outcomes *a* and *b* are in their example the death of your father and the death of your mother, respectively. Both deaths happen in terrible circumstances. *a+* and *b+* are slightly sweetened versions of *a* and *b* (though, admittedly, "sweetening" is not an appropriate term to use in such a tragic context): in the +-outcomes you receive a free grief counselling. You consider *a+* to be slightly better (slightly less bad) than *a* and *b+* to be slightly better than *b*, but *a+* and *a* are both incommensurable with *b+* and *b*. It might be argued that this incommensurability is of a radical sort: that it really is a case of incomparability in (personal) value. On this reading, it is not permissible for you to prefer the death of your father (whether sweetened or not) to the death of your mother, and vice versa. You ought not to have any

[20] Note, though, that Hare, who is interested in the preferential version of our problem, assumes that the agent lacks a considered preference in his examples. My interest is in the value relation and I thus allow as permissible both the absence of (considered) preference and its presence – but in latter case I take as permissible that I might prefer any of the items that are being compared.

(considered) preference in a case like this. Nor is it permissible for you to be indifferent between these outcomes: What is required on your part is a preference gap – the absence of a (considered) preferential attitude. Arguably, what characterizes genuine choice dilemmas. is the so-understood incomparability between the alternatives that you need to choose between.

7. *Residual paradox: Incomparability meets risk*

If gaps in a preference ordering are due to incomparabilities, then such gaps cannot permissibly be filled in. But this means that our argument for *X* being preferred to *Y* cannot in a case like this get off the ground: This argument was dependent on the permissibility of preferential completions. In cases of incomparability between outcomes, we cannot argue that action *X* must be preferred to action *Y* in any preference ordering on the grounds that it would be preferred to *Y* in any permissible completion of that ordering.[21]

So, where does this leave us? We have provided an argument showing that Complementary Dominance (V) is not compelling: An action might well be better than another even though there is no state under which its outcome is better. We have shown how this can be possible. But this argument does not apply to a weakening of Complementary Dominance (V), according to which an action cannot be better than another if there is no state under which their outcomes are comparable in value. Indeed, the following principle does seem compelling:

Incomparability (V): An action is incomparable with another action if their outcomes are incomparable under every state.[22]

This means, however, that we still confront a paradox. Incomparability (V) is intrinsically compelling but it is incompatible with Permutation (V). If the formers principle is correct, then Permutation (V) does not apply to cases in which the outcomes we reassign among equiprobable

[21] As was pointed out in a footnote above, in a model that contains incomparable items, the set of permissible completions of a preference ordering is empty. In such a model, well-roundedness must be re-defined in terms of maximal extensions: a preference ordering is well-rounded iff it is identical to the intersection of all its maximal permissible extensions. We can still impose the requirement of well-roundedness, so defined, on the orderings in K, even in the presence of incomparabilities. But this will not help us with the problem at hand.

[22] Note that if this principle is correct, then Permutation (V) does not apply to cases in which the outcomes we reassign among equiprobable states are mutually incomparable. Consequently, Permutation (P) does not apply to such cases either.

states are mutually incomparable. (Consequently, of course, Permutation (P) does not apply to such cases either.) The relationship between actions Y and Z exemplifies this point: If *a* and *b* are incomparable, then the outcomes of *Y* and *Z* are incomparable under every state. But then Incomparability (V) implies that actions *Y* and *Z* are incomparable and not equally good, as it would follow from Permutation (V).

Now, it might be claimed (I don't know if I would be prepared to make this claim) that Permutation (V) does not seem to be any less compelling in the case in which the outcomes of an action under different equi-probable states are mutually incomparable. Even in a case like this, it seems intuitive that a reassignment of outcomes among equiprobable states should not affect the value of an action. But then, despite of the progress we have made in solving the general problem we had started with – the problem of interaction between incommensurability and risk – we still are left with a more specific problem of how *incomparability* interacts with risk. Concerning this issue we seem to have at our hands something that looks like a genuine paradox: Two intrinsically compelling principles, Incomparability (V) and Permutation (V), are in conflict with each other.

TO BE CONTINUED …

References

Bales, Adam, Cohen, Daniel, and Handfield, Toby (2014) "Decision Theory for Agents with Incomplete Preferences.", *Australasian Journal of Philosophy 92*: 453-70.

Bovens, Luc, and Rabinowicz, Wlodek (2015), "The Meaning of 'Darn it!' in I. Hirose and A. Reisner (eds.), *Weighing and Reasoning: Themes from the Philosophy of John Broome*, Oxford University Press: Oxford, 129-139.

Hare, Caspar (2010) "Take the Sugar." *Analysis* 70: 237-47.

Rabinowicz, Wlodek (2008) "Value Relations." *Theoria* 74: 18-49.

Schoenfield, Miriam (2014) "Decision Making in the Face of Oarity." Ethical Perspectives 28: 263-77.