Stochastic Dominance for Incomplete Preferences[*]

Tomi Francis and Johan E. Gustafsson[†]

ABSTRACT. If incomplete preferences are rationally permissible, there are cases where two compelling principles of rational choice come into conflict: Stochastic Dominance, which says that one is rationally obligated to choose options which offer the same probabilities of preferred outcomes, and Statewise Maximality, which says that one is always permitted to choose gambles which are guaranteed not to result in a dispreferred outcome. We defend Stochastic Dominance by arguing against Statewise Maximality in cases where the two principles conflict. We rebut a standard argument for Statewise Maximality, namely, the Argument from Full Information. We then provide a direct argument against the application of Statewise Maximality to cases involving incomplete preferences. This argument proceeds from three premises. The first is that preferences are rationally required to be transitive. The second is the Sure-Thing Principle. The third premise is that probabilistic reasoning is appropriate in some cases where Statewise Maximality does not apply.

Consider a choice between the following two gambles, SUGAR and NO SUGAR, whose pay-offs depend on a fair coin flip:

Coin Flip

	Heads (1/2)	Tails (1/2)
SUGAR	$4	$2
NO SUGAR	$1	$3

Which gamble should you choose? Choosing SUGAR will give you a one-in-two chance of receiving four dollars and a one-in-two chance of receiving two dollars. Choosing NO SUGAR will give you a one-in-two chance of three dollars and a one-in-two chance of one dollar. Hence choosing SUGAR gives you the same chances of getting more money. So, if you prefer more money to less, you are rationally required to choose SUGAR. Or, at least, this is so assuming the following requirement of rationality:[1]

[*] Forthcoming in *Analysis*.

[†] We would be grateful for any thoughts or comments on this paper. They can be sent to *johan.eric.gustafsson@gmail.com*.

[1] See Quirk and Saposnik 1962, p. 141.

Stochastic Dominance If G_1 and G_2 are gambles such that

(i) for any X which is a final outcome of G_1 or G_2, it holds that G_1 is at least as likely as G_2 to result in a final outcome which is at least as preferred as X and

(ii) for some final outcome X' of G_1 or G_2, it holds that G_1 is more likely than G_2 to result in a final outcome which is at least as preferred as X',

then G_1 is preferred to G_2.

Stochastic Dominance is supported by the idea that when we evaluate gambles, we only need to look at the probabilities of each outcome occurring and our preferences over the outcomes.[2] The way the probabilities are aligned with states of nature should be preferentially irrelevant. In Coin Flip, for instance, it shouldn't matter whether a given 1/2 probability of a certain outcome is aligned with a state of nature in which a flipped fair coin lands heads, or whether it is aligned with the state in which the coin lands tails. What you should care about, on this line of thought, is your chances of getting rich.

Stochastic Dominance is compelling at first sight. Yet one may be tempted to doubt it in cases of *Opaque Sweetening*.[3] Such cases involve incomplete preferences — that is, there are outcomes A and B such that A is not at least as preferred as B and B is not at least as preferred as A. If incomplete preferences are rationally permissible, it's also rationally permissible for these preferences to be *insensitive to mild sweetening*. That is, there may be outcomes A^+ and B^+ such that A^+ is preferred to A but neither of A^+ and B is at least as preferred as the other, and B^+ is preferred to B but neither of B^+ and A is at least as preferred to as the other.[4]

Assuming that you have this configuration of preferences, consider the following choice between two gambles that are resolved on the basis of a fair coin flip:

Opaque Sweetening

	Heads (1/2)	Tails (1/2)
NO SWEETENING	A	B
FLIPPED SWEETENING	B^+	A^+

[2] Hare 2010, p. 240–1.

[3] Hare 2010, pp. 239–40; 2013, p. 46.

[4] Raz (1986, pp. 325–6) calls this kind of insensitivity to improvements the 'mark of incommensurability'.

Since FLIPPED SWEETENING offers the same probabilities of preferred outcomes to NO SWEETENING, Stochastic Dominance entails FLIPPED SWEETENING is preferred to NO SWEETENING. Accordingly, if Stochastic Dominance is a requirement of rationality, you are rationally required to prefer FLIPPED SWEETENING is preferred to NO SWEETENING. This judgement conflicts with the following (at first sight) compelling principle:[5]

> *Statewise Maximality* If it is rationally permissible that, in every state of nature, the outcome of gamble X is not preferred to the outcome of gamble Y, then it is rationally permissible not to prefer X to Y.

If Statewise Maximality is true, then it's rationally permissible not to prefer FLIPPED SWEETENING to NO SWEETENING. This is because either the coin landed on heads or it landed on tails. If it landed on heads, then the outcome of FLIPPED SWEETENING is not preferred to the outcome of NO SWEETENING. Likewise, if the coin landed tails, then the outcome of FLIPPED SWEETENING is not preferred to the outcome of NO SWEETENING. Since the outcome of FLIPPED SWEETENING is guaranteed not to be preferred to the outcome of NO SWEETENING, you are not rationally required to prefer the former to the latter if Statewise Maximality is true.

Cases of Opaque Sweetening show that, provided it's rationally permissible to have incomplete preferences, Statewise Maximality can only be true if Stochastic Dominance is not a requirement of rationality. In so far as Statewise Maximality is compelling, this poses a challenge to Stochastic Dominance.

In this paper, we respond to this challenge by arguing that Statewise Maximality is false (if incomplete preferences are rationally permissible).[6] We will first rebut what we take to be the best available substantive argument for Statewise Maximality, namely, the Argument from Full Information (§ 1). We will then provide a direct argument, based on the Sure-Thing Principle, for the verdict of Stochastic Dominance in cases of Opaque Sweetening, thus ruling out Statewise Maximality (§ 2).

[5] Similar versions of this principle are considered in Hare 2010, p. 242 and defended in Schoenfield 2014, p. 267 and Bales et al. 2014, p. 460.

[6] We do not think it's obvious that incomplete preferences *are* rationally permissible, but if they are rationally impermissible, the objection to Stochastic Dominance from Opaque Sweetening does not get off the ground.

1. The Argument from Full Information

A standard argument for Statewise Maximality is the Argument from Full Information.[7] Say that a gamble G_1 is *statewise maximal* with respect to G_2 if and only if the final outcome of G_2 is not preferred to the final outcome of G_1 in any state of nature.

The Argument from Full Information

(1) If it is rationally permissible for G_1 to be statewise maximal with respect to G_2, then you are rationally permitted to be certain that, given full information, you would not prefer G_2 to G_1.

(2) If you are rationally permitted to be certain that, given full information, you would not prefer G_2 to G_1, then you are rationally permitted not to prefer G_2 to G_1.

(3) So, if it is rationally permissible for G_1 to be statewise maximal with respect to G_2, you are rationally permitted not to prefer G_2 to G_1. (That is, Statewise Maximality holds.)

Assuming that it's rationally permissible for the gathering of information not to alter your preferences over the final outcomes of G_1 and G_2, (1) must be true. Hence, granting this assumption, the soundness of the Argument from Full Information turns on (2).

There are two prominent arguments for (2) in the literature.

The first, the *Argument from Deference*, appeals directly to the claim that it's always permissible to defer to the preferences of fully-informed versions of yourself who share the same preferences over final outcomes: since your fully-informed self is bound not to prefer G_2 to G_1, this line of thought goes, you are permitted to share this lack of preference *ex ante*.[8]

The second, the *Argument from the Primacy of Final Outcomes*, attempts to derive (2) from the alleged fact that we should ultimately be concerned with the satisfaction of our preferences over final outcomes.[9] Since you are certain that the final outcome of G_2 won't be preferred to

[7] Hare 2010, pp. 241–2.

[8] See Hare 2010, p. 242.

[9] See Schoenfield 2014, pp. 267–9 for an argument along these lines, though concerned with considerations of actual value rather than preference relations.

the final outcome of G_1, a preference for G_2 over G_1 is unwarranted — since choosing G_2 won't help you get a final outcome you prefer.[10]

THE ARGUMENT FROM DEFERENCE

Let us consider the Argument from Deference in more detail. May we always defer to the preferences of our rational, fully-informed selves? In one sense, yes. We may do so when we know our fully-informed selves will strictly prefer the outcome of one gamble to the outcome of another. The Argument from Deference, however, posits that we may defer to "the preferences" of our fully-informed selves in a more general sense.[11] At its most general, we might interpret it as permitting deference for any *preferential relation*, where a preferential relation is any relation definable in terms of propositional logical connectives and the weak preference relation.[12] But this more general principle of deference turns out to be false. We may not always defer, for instance, when it comes to the "not equally preferred" relation.[13] To see this, consider the following two gambles, where A^+ is preferred to A:

Chancy Sugar

	Heads (1/2)	Tails (1/2)
SUGAR ON HEADS	A^+	A
SUGAR ON TAILS	A	A^+

You may rationally be certain that your future self will hold the "not equally preferred" preference relation between SUGAR ON HEADS and SUGAR ON TAILS. Nevertheless, you should equally prefer the two gambles *ex ante*.

[10] A similar principle regarding moral value, called the Principle of Full Information, is endorsed by Fleurbaey and Voorhoeve 2013, p. 121. Fleurbaey and Voorhoeve, however, assume completeness.

[11] Hare (2010, p. 242) puts it like this: we may defer to any "array of preferences" which we know that our fully informed-self would hold. While it's somewhat unclear what an "array of preferences" amounts to, we think that it's most naturally interpreted as a set of preferential relations.

[12] Examples include the strict preference relation, the strict dispreference relation, and the *either preferred or dispreferred* relation. Here, 'disprefer' is not, as Fiske (2006, p. 119) claims, 'Idiotic for dislike'. Rather, it is a technical term defined as follows: X is *dispreferred* to $Y =_{\text{df}} Y$ is preferred to X.

[13] This point is also made in Rabinowicz 2022, p. 205.

For the Argument from Deference to work, then, an intermediate principle of deference needs to be true: we may defer when it comes to a certain class of preference relations, including the strict preference relation and the "not preferred" relation, but excluding other preferential relations such as "not equally preferred". What are the possibilities? One is to say that we should defer to positive preferential attitudes (those one *has*), but not to negative preferential attitudes (those one merely *lacks*). This can then be combined with the claim that, when you have a preferential gap between two gambles, you adopt towards them the positive preferential attitude of ambivalence. Even if this distinction can be made precise, however, it seems to us that it will not work: the *Chancy Sugar* case also shows that you should not defer when it comes to the "either preferred or dispreferred" relation, which is a preferential attitude you have, rather than one you lack.

Another possibility is to defer only when it comes to those preferential relations that are decisive regarding whether you ought, or are permitted, to choose an option. But this proposal also overgeneralises. Suppose that you may either bet that a coin lands heads, bet that it lands tails, or not bet at all. You know in advance that your fully informed self would disprefer not betting to one of the two betting options. On the present proposal, you should defer when it comes to this decisive preferential relation. This cannot be true in general, since it is sometimes rationally permissible not to bet in cases of this form.

The Argument from Deference, then, is on shaky footing without an explanation as to why we should defer when it comes to the *not preferred* relation in particular. There is, however, another argument for Statewise Maximality.

THE ARGUMENT FROM THE PRIMACY OF FINAL OUTCOMES

Next, consider the Argument from the Primacy of Final Outcomes. The idea here is that rationality is supposed to help us satisfy our preferences over final outcomes: we care about things like expectations only in service to this goal. When we are rationally certain that we would not prefer the final outcome of G_2 to the final outcome of G_1, we know, in advance, that choosing G_2 wouldn't result in us getting a final outcome we prefer. A decision theory which nevertheless rules out G_2 would be going further than is warranted by our concern for final outcomes (or so goes the argument).[14]

We grant that there is some sense in which it is true that the prescriptions of decision must help us to achieve our preferences over final outcomes. The question is how to make this platitude (the Primacy of Final Outcomes) precise in a plausible way.

One way of making it precise would be to say that we should choose a gamble whenever it will, more likely than not, lead to a final outcome we strictly prefer. But this would of course be implausible: it is surely rationally permissible to take a bet with a 40% chance of a large pay-off.

More promisingly, we could take the Primacy of Final Outcomes to require us to take into account information about the *extent* to which our preferences are better satisfied by our ending up with one final outcome rather than another. We might then compare gambles according to their probability-weighted sums (not necessarily using real-number values) across all states of nature — taking states to be neutral if neither outcome is better.

This approach suggests that we should prefer gambles in so far as their outcomes are preferred in particular states of nature, but we should be neutral between gambles in so far as we have no preference between their outcomes in other states of nature. Accordingly, it justifies not only Statewise Maximality but also the following principle:[15]

> *Strict Statewise Maximality* It is rationally required that, if, in every state of nature, the outcome of G_2 is not preferred to the outcome of G_1 and, in some state of nature, the outcome of G_1 is preferred to the outcome of G_2, then G_1 is preferred to G_2.

But Strict Statewise Maximality should be rejected if it is rationally permitted to have preferential gaps that are insensitive to some mild sweetenings or sourings. This is because, under those conditions, Strict Statewise Maximality generates preference cycles. For instance, suppose that you prefer A^+ to A and that you have a preferential gap between A and B and

[14] See Schoenfield 2014, p. 268.

[15] Understood this way, the Primacy of Final Outcomes supports Strict Statewise Maximality rather than merely Doody's (2019, p. 1091) Principle of Predominance — which merely posits a rational permission to choose G_1 over G_2. While Doody accepts the Principle of Predominance but denies Strict Statewise Maximality, it seems to us that the reasons Doody offers in favour of it being permitted to choose G_1 over G_2 are also reasons to prefer G_1 over G_2.

between A^+ and B. Now, consider the following three gambles G_1, G_2, and G_3:[16]

Preference Cycle

	S_1 (1/3)	S_2 (1/3)	S_3 (1/3)
G_1	A^+	B	A
G_2	A	A^+	B
G_3	B	A	A^+

Suppose that one has the following preference cycle: G_1 is preferred to G_2, G_2 is preferred to G_3, and G_3 is preferred to G_1. Strict Statewise Maximality implies that this preference cycle is rationally permissible. We should therefore reject this principle, since cyclic preferences are irrational.[17]

We propose instead that the Primacy of Final Outcomes should be understood as follows: any consideration in favour of choosing one gamble over another must be grounded in considerations which favour the final outcomes of that gamble. By a *consideration*, we just mean some respect in which the outcome or gamble is preferred. But, understood this way, the Primacy of Final Outcomes turns out not to support the verdicts of Statewise Maximality in opaque sweetening cases after all; instead, it undermines them.

To see this, consider a typical sort of case involving a preferential gap. You might become a lawyer (A) or a clarinetist (B). You would be better-paid in law. But your work as a clarinettist would be more fulfilling. And, since you do not have in mind some precise way of trading off these two features, you have a preferential gap between A and B. Another way of describing the situation is that there is a financial consideration in favour of becoming a lawyer, and fulfilment consideration in favour of becoming a clarinettist. Notice that, although you have no all-things-considered preference between A and B, there are nevertheless considerations counting in favour of A, and considerations counting in favour of B; it's just that these considerations are indecisive.

[16] Bader (2018, p. 504) attempts a similar argument, but his example does not quite work. In his case, Strict Statewise Maximality does not entail that what he calls L_C is preferred to to what he calls L_B. So he does not get a cycle. In personal communication, Bader reports that the penultimate version of his paper had a working example.

[17] See the money-pump argument in Gustafsson and Rabinowicz 2020. It may be objected that the money-pump argument would prove too much in this context, since there are also money pumps for incomplete preferences. But note that the money-pump arguments against incomplete preferences need more assumptions than the best money pumps against cyclic preferences. See Gustafsson 2022, pp. 35–8.

Now, recall the opaque sweetening case from earlier. You must choose between gambles involving the lawyer and clarinettist careers, where A^+ and B^+ are, respectively, the lawyer and clarinettist careers plus some small salary increase:

Opaque Sweetening

	Heads (1/2)	Tails (1/2)
NO SWEETENING	A	B
FLIPPED SWEETENING	B^+	A^+

Considerations of fulfilment favour NO SWEETENING and SWEETENING equally overall: while considerations of fulfilment favour NO SWEETENING on tails, there are equal and opposite considerations of fulfilment favouring FLIPPED SWEETENING on heads. But financial considerations favour FLIPPED SWEETENING overall, whereas financial considerations favour NO SWEETENING on heads and stronger financial considerations favour FLIPPED SWEETENING on tails. Taken together, the two features of final outcomes which you care about seem to favour FLIPPED SWEETENING, in line with Stochastic Dominance. This is so, even though fulfilment and financial considerations do not precisely trade off against each other.[18]

Accordingly, in so far as it's plausible to understand the Primacy of Final Outcomes in the way we have suggested, the Argument from the Primacy of Final Outcomes fails to establish Statewise Maximality. But, since we haven't shown that the Primacy of Final Outcomes *must* be understood in this way, this does not give us a strong argument against Statewise Maximality. So we will next argue against Statewise Maximality directly.

2. The Coin-Flip-Indifference Argument

To argue against Statewise Maximality, we will make two assumptions. First, we assume that Transitivity is a requirement of rationality:

> *Transitivity* If G_1 is at least as preferred as G_2 and G_2 is at least as preferred as G_3, then G_1 is at least as preferred as G_3.

[18] For a similar argument, see Hare 2010, p. 240 and Doody 2019, pp. 1087–9.

Transitivity may be less compelling if preference gaps are rationally permissible than if they are not.[19] Still, even in the context of incomplete preferences, Transitivity is more compelling than Statewise Maximality.

Our second assumption is that the Sure-Thing Principle is a requirement of rationality:[20]

> *The Sure-Thing Principle* Suppose that X and Y are gambles over a set of states of nature U, that V is a subset of U, and that X and Y have the same outcome for each S in V. Then X is at least as preferred as Y if and only if, conditional on none of the states in V obtaining, the outcome of X is at least as preferred as the outcome of Y.

The idea is that, since X and Y are equivalent in the states in V, we can ignore those states. So, if X is at least as preferred as Y conditional on none of the ignored states obtaining, X should be at least as preferred as Y overall.

We now proceed with our argument against Statewise Maximality. We begin with some terminology. For any final outcomes X and Y, we can say that an agent is *coin-flip indifferent* between X and Y if and only if she has an equal preference between getting X on heads and Y on tails, or instead getting Y on heads and X on tails.

Our first observation is that, given our two assumptions, rationality requires that the coin-flip indifference relation is transitive. Suppose that an agent is coin-flip indifferent between X and Y and between Y and Z, and consider the following gambles:

	S_1 (1/3)	S_2 (1/3)	S_3 (1/3)
G_1	X	Z	Y
G_2	Y	Z	X
G_3	Z	Y	X
G_4	Z	X	Y

[19] For example, the money-pump argument that rational preferences are transitive in Gustafsson 2010 assumes that rational preferences are complete.

[20] Savage 1954, pp. 21–2 and Joyce 1999, p. 85. Although the Sure-Thing Principle is compelling, it has been challenged on the grounds of its incompatibility with the alleged rationality of Allais preferences; see Allais 1953, p. 527; 1979, p. 89. We are not persuaded by this argument, as we think that there are independent reasons to reject the rationality of Allais preferences. For instance, Gustafsson (2022, pp. 51-6) offers a money pump for Allais preferences.

First, compare G_1 and G_2. Conditional on the non-occurrence of state S_2 (in which both gambles yield the same final outcome), these gambles both yield equal chances of receiving X and Y. Since the agent is coin-flip indifferent between X and Y, she has an equal preference between these conditioned gambles. Hence, by the Sure-Thing Principle, she must have an equal preference between G_1 and G_2, without conditioning on the non-occurrence of S_2.

By repeating the same argument, it can be shown that a rational agent must have an equal preference between G_2 and G_3 and between G_3 and G_4. Transitivity thus requires her to have an equal preference between G_1 and G_4. Since G_1 and G_4 yield the same outcome in state S_3, the Sure-Thing Principle implies that the agent must equally prefer these gambles, conditioned on the non-occurrence of state S_3. That is a coin-flip between X and Z. The agent must therefore be coin-flip indifferent between X and Y.

Now consider again Opaque Sweetening, but with the addition of a third option, FLIPPED NO SWEETENING:

Extra-Flip Opaque Sweetening

	Heads (1/2)	Tails (1/2)
NO SWEETENING	A	B
FLIPPED NO SWEETENING	B	A
FLIPPED SWEETENING	B^+	A^+

It's easy to show, using the Sure-Thing Principle and Transitivity (or a statewise dominance principle), that an agent who prefers A^+ to A and B^+ to B is rationally required to prefer FLIPPED SWEETENING to FLIPPED NO SWEETENING, since the former is preferred to the latter in every state. Hence, if the agent is coin-flip indifferent between A and B, by Transitivity she must prefer FLIPPED SWEETENING to NO SWEETENING, in line with Stochastic Dominance and contrary to Statewise Maximality.

The key point of contention, of course, is whether the agent should indeed be coin-flip indifferent between A and B. The transitivity of coin-flip indifference supports this claim, if we also assume the following requirement of rationality:

> *Commensurable Coin-Flip Indifference* If final outcome X is at least as good as final outcome Y in every dimension the agent cares about, then the agent is coin-flip indifferent between X and Y.

Suppose, for instance, that outcome A is eating an apple and outcome B is eating an orange. We can assume that there is a third outcome C which is inferior in every dimension the agent cares about than each of A and B; for instance, being poisoned. Commensurable Coin-Flip Indifference entails that one is coin-flip indifferent between A and C and between C and B. By the transitivity of coin-flip indifference, it then follows that one is coin-flip indifferent between A and B.

It may be objected that appealing to Commensurable Coin-Flip Indifference assumes most of what is to be proved, since it is very similar to full Stochastic Dominance. But the challenge to Stochastic Dominance we are considering is precisely that its verdicts are questionable in the sorts of Opaque Sweetening cases considered in this paper. Commensurable Coin-Flip Indifference is not open to this challenge. Unlike standard Stochastic Dominance, it only concerns prospects where the outcomes are fully commensurable. The idea behind it is that probabilistic reasoning is appropriate when there is no possibility of incommensurability. To deny it, we would have to throw out probabilistic reasoning almost entirely.

It may also be objected that if we accept Commensurable Coin-Flip Indifference, this must be at least in part because the extent to which it would be preferable to get A rather than C on heads is the same as the extent to which it would be preferable to get A rather than C on tails; and similarly, of course, for B and C. This raises the worry that we could then compare A and B by comparing the extents to which each is better than C. It needn't be the case, however, that the extents to which final outcomes are preferable to others can always be placed on a unidimensional scale; indeed, we think that they had better not be if incomplete preferences are rationally permissible. In our case, A is preferable to C to the extent that apples are preferable to poison, and B is preferable to C to the extent that oranges are preferable to poison; but this does not imply that you can compare apples and oranges.

References

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.

— (1979) 'The Foundations of a Positive Theory of Choice involving Risk and a Criticism of the Postulates and Axioms of the American School', in M. Allais and O. Hagen, eds., *Expected Utility Hypothesis and the Allais Paradox: Contemporary Discussions of Decisions under Uncertainty with Allais' Rejoinder*, pp. 27–145, Dordrecht: Reidel.

Bader, Ralf M. (2018) 'Stochastic Dominance and Opaque Sweetening', *Australasian Journal of Philosophy* 96 (3): 498–507.

Bales, Adam, Daniel Cohen, and Toby Handfield (2014) 'Decision Theory for Agents with Incomplete Preferences', *Australasian Journal of Philosophy* 92 (3): 453–470.

Doody, Ryan (2019) 'Opaque Sweetening and Transitivity', *Australasian Journal of Philosophy* 97 (3): 559–571.

Fiske, Robert Hartwell (2006) *The Dictionary of Disagreeable English*, Cincinnati: Writer's Digest Books.

Fleurbaey, Marc and Alex Voorhoeve (2013) 'Decide as You Would with Full Information!: An Argument against phEx Ante Pareto', in Nir Eyal, Samia Hurst, Ole F. Norheim, and Daniel Wikler, eds., *Inequalities in Health: Concepts, Measures, and Ethics*, pp. 113–128, Oxford: Oxford University Press.

Gustafsson, Johan E. (2010) 'A Money-Pump for Acyclic Intransitive Preferences', *Dialectica* 64 (2): 251–257.

— (2022) *Money-Pump Arguments*, Cambridge: Cambridge University Press.

Gustafsson, Johan E. and Wlodek Rabinowicz (2020) 'A Simpler, More Compelling Money Pump with Foresight', *The Journal of Philosophy* 117 (10): 578–589.

Hare, Caspar (2010) 'Take the Sugar', *Analysis* 70 (2): 237–247.

— (2013) *The Limits of Kindness*, Oxford: Oxford University Press.

Joyce, James M. (1999) *The Foundations of Causal Decision Theory*, Cambridge: Cambridge University Press.

Quirk, James P. and Rubin Saposnik (1962) 'Admissibility and Measurable Utility Functions', *The Review of Economic Studies* 29 (2): 140–146.

Rabinowicz, Wlodek (2022) 'Incommensurability Meets Risk', in Henrik Andersson and Anders Herlitz, eds., *Value Incommensurability: Ethics,*

Risk, and Decision-Making, pp. 201–230, London: Routledge.
Raz, Joseph (1986) *The Morality of Freedom*, Oxford: Clarendon Press.
Savage, Leonard J. (1954) *The Foundations of Statistics*, New York: Wiley.
Schoenfield, Miriam (2014) 'Decision Making in the Face of Parity', *Philosophical Perspectives* 28 (1): 263–277.