

Title Page

Manuscript title:

Stochastic Dominance for Incomplete Preferences

Authors:

Tomi Francis and Johan E. Gustafsson

Keywords:

stochastic dominance; opaque sweetening; rational choice; statewise
maximality; preferential gaps; incomparability

Tomi Francis (corresponding author)

Address:

Department of Philosophy
University of Oxford; Oxford OX2 6GG; United Kingdom

Email:

tomi.francis@philosophy.ox.ac.uk

ORCID:

0000-0003-1442-5854

Johan E. Gustafsson

Address:

Population Wellbeing Initiative
University of Texas at Austin; Austin, TX 78712; United States

Email:

johan.eric.gustafsson@gmail.com

ORCID:

0000-0002-9618-577X

Website:

<https://johanegustafsson.net>

Phone:

+1 (512) 645-8316

Ethical statement:

The work on this paper did not involve research on human participants nor animals. The authors did not receive funding that would result in any bias. The authors have no relevant financial or non-financial interests to disclose.

ACCEPTED MANUSCRIPT

Stochastic Dominance for Incomplete Preferences

Tomi Francis and Johan E. Gustafsson

ABSTRACT. According to Stochastic Dominance, it is rationally obligatory to prefer one gamble to another if it gives you the same chances of getting final outcomes you prefer. According to Statewise Maximality, it is rationally permissible not to disprefer a gamble if it is guaranteed not to result in a final outcome you disprefer. These principles conflict in cases involving incomplete preferences, known as *Opaque Sweetening* cases. In this paper, we argue for Stochastic Dominance and against Statewise Maximality in Opaque Sweetening cases. First, we rebut two standard arguments for statewise maximality, which we call the Argument from Full Information and the Argument from the Primacy of Final Outcomes. We then provide an argument for the verdict of Stochastic Dominance in Opaque Sweetening cases. This argument appeals to the principle of Transitivity, the Sure-Thing Principle, and the claim that stochastic reasoning is appropriate at least for choices which do not involve incomparability.

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

	<i>Coin Flip</i>	
	Heads ($1/2$)	Tails ($1/2$)
SUGAR	\$4	\$2
NO SUGAR	\$1	\$3

Which gamble should you choose? Choosing SUGAR would give you a one-in-two chance of getting four dollars and a one-in-two chance of receiving two dollars. Choosing NO SUGAR would give you a one-in-two chance of getting three dollars and a one-in-two chance of getting one dollar. Hence choosing SUGAR would give you the same chances of getting more money. So, if you like money (and who doesn't?), you are rationally required to choose SUGAR. Or, at least, this is so given the following requirement of rationality:¹

¹ See Quirk and Saposnik 1962, p. 141.

Stochastic Dominance If G and G' are gambles such that

- (i) for any final outcome X of G or G' , it holds that G is at least as likely as G' to result in a final outcome which is at least as preferred as X and
- (ii) for some final outcome Y of G or G' , it holds that G is more likely than G' to result in a final outcome which is at least as preferred as Y ,

then G is preferred to G' .

Stochastic Dominance is supported by the idea that, when you evaluate gambles, you only need to look at the probabilities of the potential final outcomes and your preferences over those outcomes.² The way the probabilities are aligned with states of nature should be irrelevant to your preferences. Thus, in the *Coin Flip* case, you shouldn't care whether a given one-in-two chance 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. You should only care about your chances of getting the money.

Stochastic Dominance is quite compelling. Yet one might doubt its implications in *Opaque Sweetening* cases.³ Such cases involve incomplete preferences. That is, for some final outcomes A and B , you have a *preferential gap* between A and B : it holds that A is not at least as preferred as B and B is not at least as preferred as A . If preferential gaps are rationally permissible, it's also rationally permissible for these preferential gaps to be *insensitive to mild sweetening*. That is, not only do you have a preferential gap between A and B , but also between A and B^+ , and between B and A^+ .⁴ Here, A^+ and B^+ are just like A and B respectively except that they are superior in some respect you care about.

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

<i>Opaque Sweetening</i>		
	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'.

Note that *FLIPPED SWEETENING* offers sweetened outcomes with the same probabilities as the unsweetened outcomes in *NO SWEETENING*. And, since you prefer the sweetened outcomes, Stochastic Dominance entails that *FLIPPED SWEETENING* is preferred to *NO SWEETENING*. So, if Stochastic Dominance is a requirement of rationality, you are rationally required to prefer *FLIPPED SWEETENING* to *NO SWEETENING*. This judgement, however, conflicts with the following apparently compelling principle:⁵

Statewise Maximality If it is rationally permissible that, in every state of nature, the outcome of gamble G is not preferred to the outcome of gamble G' , then it is rationally permissible not to prefer G to G' .

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 heads or it landed tails. If it landed heads, then the outcome of *FLIPPED SWEETENING* (B^+) is not preferred to the outcome of *NO SWEETENING* (A). Likewise, if the coin landed tails, then the outcome of *FLIPPED SWEETENING* (A^+) is not preferred to the outcome of *NO SWEETENING* (B). So the outcome of *FLIPPED SWEETENING* is guaranteed not to be preferred to the outcome of *NO SWEETENING*. And so Statewise Maximality entails that you are not rationally required to prefer *FLIPPED SWEETENING* to *NO SWEETENING*.

Provided that it's rationally permissible to have incomplete preferences, Opaque Sweetening cases show that Statewise Maximality can only be true if Stochastic Dominance is not a requirement of rationality. So, insofar as Statewise Maximality is compelling given this provision, we have a challenge to Stochastic Dominance.

We will 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'll then provide an argument, based on the Sure-Thing Principle, for

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 It's far from clear that incomplete preferences *are* rationally permissible, but, if they're not, then the objection to Stochastic Dominance from Opaque Sweetening does not get off the ground.

the verdict of Stochastic Dominance in Opaque Sweetening cases, thus ruling out Statewise Maximality (§ 2).

1. The Argument from Full Information

Let's say a gamble G is *statewise maximal* with respect to G' if and only if, no matter which state of nature eventuates, the final outcome of G' is not preferred to the final outcome of G . We can now state the Argument from Full Information:⁷

The Argument from Full Information

- (1) If it is rationally permissible for G to be statewise maximal with respect to G' , then you are rationally permitted to be certain that, given full information, you would not prefer G' to G .
- (2) If you are rationally permitted to be certain that, given full information, you would not prefer G' to G , then you are rationally permitted not to prefer G' to G .
- (3) So, if it is rationally permissible for G to be statewise maximal with respect to G' , you are rationally permitted not to prefer G' to G . (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 and G' , premise (1) must be true. So, granting premise (1), the soundness of the Argument from Full Information turns on premise (2).

There are two prominent arguments for (2). The first argument, which we can call the *Argument from Deference*, appeals 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' to G , you are, on this line of thought, permitted to share this lack of preference *ex ante*.⁸

The second argument, which we can call the *Argument from the Primacy of Final Outcomes*, attempts to derive (2) from the purported fact

⁷ Hare 2010, pp. 241–2.

⁸ See Hare 2010, p. 242.

that one should ultimately be concerned with the satisfaction of one's preferences over final outcomes.⁹ Since you're certain that you won't prefer the final outcome of G' to the final outcome of G , a preference for G' over G is unwarranted — choosing G' won't help you get a final outcome you prefer.¹⁰

THE ARGUMENT FROM DEFERENCE

Let's 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 that our fully-informed selves will prefer the outcome of one gamble to the outcome of another. But the Argument from Deference posits a more general form of deference.¹¹ At its most general, we might interpret it as permitting deference for any *preferential relation*, where a preferential relation is any relation on gambles definable in terms of propositional logical connectives and the weak preference relation.¹² But this more general principle of deference is 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 :

<i>Chancy Sugar</i>		
	Heads ($1/2$)	Tails ($1/2$)
SUGAR ON HEADS	A^+	A
SUGAR ON TAILS	A	A^+

In Chancy Sugar, you know that your future self will hold the *not indifferent* preferential relation between SUGAR ON HEADS and SUGAR ON

⁹ Schoenfeld (2014, pp. 267–9) makes an argument along these lines, although she's concerned with considerations of actual value rather than preferences.

¹⁰ A similar principle regarding moral value, called the Principle of Full Information, is endorsed by Fleurbaey and Voorhoeve (2013, p. 121), but they 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 we find it slightly unclear what an 'array of preferences' amounts to, we think that it is most naturally interpreted as a set of preferential relations.

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

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

TAILS. Nevertheless, *ex ante*, you should be indifferent between the two gambles.

For the Argument from Deference to work, it needs to appeal to an intermediate principle of deference which says that you may defer to your future self for a certain class of preferential relations — including the *preferred* and *not preferred* relations but excluding other preferential relations such as the *not indifferent* relation.

What are the possibilities? One is to say that we should defer when it comes to positive preferential relations (describing how the relata are related), but not when it comes to negative preferential relations (describing how the relata are not related).¹⁴ But, even if the distinction between positive and negative relations can be made precise, this proposal still won't work: the *Chancy Sugar* case also shows that you should not defer when it comes to the *either preferred or dispreferred* relation, which seems to be a positive relation.

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 overgeneralizes. Suppose 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 are rationally permitted to defer when it comes to this decisive preferential relation. But this can't be true in general: In some cases of this form, it is rationally required not to bet.

The Argument from Deference, then, is on shaky ground without an explanation as to why we should defer when it comes to the *not preferred* relation in particular. So let us turn to another argument for premise (2).

THE ARGUMENT FROM THE PRIMACY OF FINAL OUTCOMES

What is the point of following rational requirements on our choices over gambles? Here is one answer: The point is to help us satisfy our preferences over final outcomes. We should care about things like expectations only in service to this goal.

But note that when we are rationally certain that we wouldn't 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.

¹⁴ Chang 2002, p. 663. While the distinction between positive and negative relations has some intuitive pull, it is far from clear that the distinction makes sense formally.

A decision theory that nevertheless prohibits G_2 would be going further than is warranted by our concern for final outcomes. Or so goes the Argument from the Primacy of Final Outcomes.¹⁵

We grant that there is some sense in which that the prescriptions of decision theory 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 prefer — that is, we should maximize our chances of getting what we prefer. But this would of course be implausible: surely, it's sometimes 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 of preference satisfaction (not necessarily using real-number values) across all states of nature — taking states to be neutral if neither outcome is preferred.

On this approach, we should prefer gambles insofar as their outcomes are preferred in particular states of nature but we should be neutral between gambles insofar 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:¹⁶

Strict Statewise Maximality It is rationally required that: if, in every state of nature, the outcome of G' is not preferred to the outcome of G and, in some state of nature, the outcome of G is preferred to the outcome of G' , then G is preferred to G' .

Strict Statewise Maximality, however, 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

¹⁵ See Schoenfield 2014, p. 268.

¹⁶ 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 over G' . 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 over G' are also reasons to prefer G to G' .

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 between A^+ and B . Now, consider the following three gambles G_1 , G_2 , and G_3 :¹⁷

	<i>Preference Cycle</i>		
	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^+

Strict Statewise Maximality implies that the following preference cycle is rationally permissible: G_1 is preferred to G_2 , G_2 is preferred to G_3 , and G_3 is preferred to G_1 . But this is false: Cyclic preferences are irrational.¹⁸ We should therefore reject Strict Statewise Maximality.

So it's implausible to take states to be neutral when neither outcome is preferred. 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 that 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. Suppose you could become either a lawyer or a poet, but not both.¹⁹ You would be better paid in law. But your work as a poet would be more fulfilling. And, since you don't have in mind a precise way of trading off these two features, you have a preferential gap between being a lawyer and being a poet. Another way of describing the situation is that there's a

¹⁷ 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 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.

¹⁸ 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; forthcoming.

¹⁹ In practice, however, these careers needn't be mutually exclusive. See Liptak 2002, p. 41.

financial consideration in favour of becoming a lawyer and there's a fulfilment consideration in favour of becoming a poet. Notice that, although you have no all-things-considered preference between the two careers, there are nevertheless considerations counting in favour of becoming a lawyer and considerations counting in favour of becoming a poet — it's just that these considerations are indecisive.

Now, consider a career-choice instance of Opaque Sweetening. You must choose between gambles involving the lawyer and poet careers and the same careers with a \$1 salary increase:

<i>Opaque Career Sweetening</i>		
	Heads ($1/2$)	Tails ($1/2$)
NO SWEETENING	Lawyer	Poet
FLIPPED SWEETENING	Poet + \$1	Lawyer + \$1

Considerations of fulfilment favour NO SWEETENING and SWEETENING equally overall. This is because, 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. This is because, while financial considerations favour NO SWEETENING on heads, stronger financial considerations favour FLIPPED SWEETENING on tails. Taken together, the two features of final outcomes which you care about favour FLIPPED SWEETENING, in line with Stochastic Dominance.²⁰

This way of understanding the primacy of final outcomes also undermines the following putative requirement of rationality:

Negative Dominance If gamble G is preferred to gamble G' , then at least one final outcome of G is preferred to at least one final outcome of G' .²¹

In the context of incomplete preferences, Negative Dominance exhibits an overall-ranking fetish (as does Statewise Maximality): It requires you to care about whether final outcomes are overall preferred or dispreferred as such — in a way that goes beyond caring about the considerations in virtue of which the outcomes are preferred or dispreferred.

²⁰ For a similar argument, see Hare 2010, p. 240 and Doody 2019, pp. 1087–9.

²¹ See Lederman forthcoming for discussion of this principle.

To see this, suppose that you have been offered a job, but you have a choice of compensation packages involving varying salary levels and vacation entitlement.²² You could pick the plain option:

(i) \$160,000 per year; 35 days vacation.

Alternatively, you could pick the “mystery” option, consisting of a fifty-fifty gamble between the following compensation packages:

(ii) \$100,000 per year; 55 days vacation.

(iii) \$200,000 per year; 5 days vacation.

If you value both money and vacation entitlement linearly, you would presumably prefer the plain option to the mystery option, since the plain option is expectedly better in both of the important respects you care about. And you might have this preference even if you do not prefer (i) to a sure thing of either (ii) or (iii). In that case, your preferences will violate Negative Dominance.²³

Accordingly, the primacy of final outcomes, properly understood, undermines rather than supports both Statewise Maximality and Negative Dominance. We will next argue against Statewise Maximality directly.

2. The Coin-Flip-Indifference Argument

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

²² For another counter-example, see the two-dimensional version of total utilitarianism in Gustafsson 2020, pp. 88–94, where a higher quality of life makes outcomes better but sufficient differences in population size can outweigh quality of life and make outcomes incomparable to each other. Let *A* and *B* be outcomes with a higher quality of life than outcome *C*. And suppose that *A* has a smaller population than *C* and *B* has a larger population than *C* so that overall *A* and *B* are both incomparable to *C*. The difference in population between *A* and *C* is the same as that between *C* and *B*. Now, let *G* be a fifty-fifty lottery between *A* and *B*. Then the expected population size is the same in *C* as in *G*. Since *C* and *G* are the same in expected population size, the guaranteed higher quality of life in *G* does make it better than *C*, which violates Negative Dominance (since neither outcome of *G* is better than *C*). But there’s no mystery where the extra value in *G* comes from — it comes from *G*’s having a higher quality of life than *C*.

²³ This case also serves as a counter-example to the Vagueness Sure Thing principle discussed by Manzini and Mariotti (2008, p. 308).

Transitivity If X is at least as preferred as Y and Y is at least as preferred as Z , then X is at least as preferred as Z .

Transitivity may be less compelling if preference gaps are rationally permissible than if they are not.²⁴ 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:²⁵

The Sure-Thing Principle Let G and G' be gambles over a set of states of nature U , and let V be a proper subset of U such that G and G' have the same outcome for each state S in V . Then G is at least as preferred as G' if and only if, conditional on none of the states in V obtaining, G is at least as preferred as G' .

The idea here is that since G and G' have the same outcome for all states in V , we can ignore those states when deciding which gamble should be preferred.

We now proceed with our argument against Statewise Maximality. We begin with some terminology. For any final outcomes X and Y , we say that an agent is *coin-flip indifferent* between X and Y if and only if the agent is indifferent between (i) a fifty-fifty gamble between getting X in state of nature S and getting Y in state of nature S' and (ii) a fifty-fifty gamble between getting Y in S and getting X in S' .

Our first observation is that, given our two assumptions, rationality requires that the coin-flip indifference relation is transitive. Suppose that you are coin-flip indifferent between X and Y and between Y and Z , and consider the following gambles:²⁶

²⁴ For example, the money-pump argument that rational preferences are transitive in Gustafsson 2010 assumes that rational preferences are complete.

²⁵ 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're not persuaded by this argument, as we think that there are independent reasons to reject the rationality of Allais preferences. For instance, see Gustafsson 2022, pp. 51–6 for a money pump for Allais preferences.

²⁶ Here, S_1 and S_2 are arbitrarily chosen states of nature. In order for our argument to go through, the construction of gambles G_1, G_2, G_3 , and G_4 needs to be logically possible. So we need to assume that S_1 and S_2 are not logically exhaustive.

	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

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 you are coin-flip indifferent between X and Y, you are indifferent between these conditioned gambles. Hence, by the Sure-Thing Principle, you must be indifferent 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 you must be indifferent between G_2 and G_3 and between G_3 and G_4 . Transitivity thus requires you to be indifferent between G_1 and G_4 . And, since G_1 and G_4 yield the same outcome in state S_3 , the Sure-Thing Principle implies that you must be indifferent between these gambles, conditioned on the non-occurrence of S_3 . Conditioned on the non-occurrence of S_3 , both G_1 and G_4 are fifty-fifty gambles between X and Z (but with X and Z obtaining in opposite states of nature). So, since we can run this argument for any states of nature in place of S_1 and S_2 , it follows that you must be coin-flip indifferent between X and Z.

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

<i>Extra-Flip Opaque Sweetening</i>		
	Heads ($1/2$)	Tails ($1/2$)
NO SWEETENING	A	B
FLIPPED NO SWEETENING	B	A
FLIPPED SWEETENING	B^+	A^+

Here, both Heads and Tails are states of nature.

It's easy to show, using the Sure-Thing Principle and Transitivity (or just 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. Moreover, if the agent is coin-flip indifferent between A and B, then the agent is indifferent between NO SWEETENING and FLIPPED NO SWEETENING. Then, by Transitivity, it follows that the agent 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* 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, eating poison. 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 underlying idea is that stochastic reasoning is appropriate when there is no possibility of incommensurability. To deny it, we would have to throw out stochastic reasoning almost entirely.²⁷

It may also be objected that if we accept Commensurable Coin-Flip Indifference, this must be 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*

²⁷ Since coin-flip indifference is transitive, any failures of Commensurable Coin-Flip Indifference cascade outwards: if the agent is coin-flip indifferent between *X* and *Y* but not between *X* and *Z*, then the agent must not be coin-flip indifferent between *Y* and *Z*. The upshot is that failures of coin-flip indifference are infectious.

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'école Américaine', *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.
- Chang, Ruth (2002) 'The Possibility of Parity', *Ethics* 112 (4): 659–688.
- 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.
- (2020) 'Population Axiology and the Possibility of a Fourth Category of Absolute Value', *Economics and Philosophy* 36 (1): 81–110.

- (2022) *Money-Pump Arguments*, Cambridge: Cambridge University Press.
- (forthcoming) ‘A Behavioural Money-Pump Argument for Completeness’, *Theory and Decision*.
- 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.
- Lederman, Harvey (forthcoming) ‘Of Marbles and Matchsticks’, in Tamar Szabó Gendler, John Hawthorne, Julianne Chung, and Alex Worsnip, eds., *Oxford Studies in Epistemology Volume 8*, Oxford: Oxford University Press.
- Liptak, Adam (2002) ‘Justices Call on Bench’s Bard to Limit His Lyricism’, *The New York Times* (December 15, 2002): 41.
- Manzini, Paola and Marco Mariotti (2008) ‘On the Representation of Incomplete Preferences over Risky Alternatives’, *Theory and Decision* 65 (4): 303–323.
- Quirk, James P. and Ruben 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.