THE JOURNAL OF PHILOSOPHY
VOLUME CXXI, NO. 3, MARCH 2024

THE ISAAC LEVI PRIZE 2023

OPTIMIZATION AND BEYOND*

What, all things considered, should I do? This is of course the central question of practical reasoning, and people who are otherwise of quite diverse persuasions think that a plausible answer to this question is the following: Do your best in the situation in which you find yourself!

Now, thanks to much work by scholars who study rational choice, the basic idea involved in this claim is familiar enough, and it is this. We bracket the question of what ultimately matters. We nevertheless say that rational agents ought to go for the best—they ought to optimize—when such a best, or optimal, alternative is defined (with respect to whatever ultimately matters).[1] This might also come across

* I start with some debts that I can only acknowledge and never repay. The greatest of them all is to my doctoral advisors, Martin van Hees and Roland Iwan Luttens. Then there is my friend and colleague Marina Uzunova. It would be a mug's game to try to specify how Martin, Roland, and Marina have shaped this paper. I would also like to thank Francesco Guala, an anonymous member of the committee for the Isaac Levi Prize, and especially Stefan Wintein, for very perceptive comments on earlier drafts. For helpful discussion of the ideas in this paper, I am grateful to Constanze Binder, Douglas Bernheim, Franz Dietrich, Conrad Heilmann, Lucie White, and audiences at the Lake Como Summer School (which is one of the best things there is in philosophy). Thanks are also due to Alyssa Timin for editorial assistance. Finally, I would be remiss if I did not place on record my gratitude to the late, great Isaac Levi, and the great Amartya Sen. While I have never had the opportunity to discuss the ideas contained in this paper with them, the work of these two scholars has influenced me in ways that I cannot even begin to articulate. I therefore find it apt that this paper is published as part of a series that memorializes Levi's life and work, and that Sen is one of my main interlocutors here.

[1] It might be useful to define the concept of a best or optimal alternative here. Let R be a reflexive binary preference or value relation defined on a finite set of alternatives X, that is, $R \subseteq X \times X$, and call this relation *at least as good as*. R stands for the overall ranking of alternatives that a deliberating agent holds, and R can be based on a formal or intuitive weighing of a vector $\langle R_1, R_2, R_3, \ldots, R_n \rangle$ whose elements are the various

0022-362X/24/2103/121–146
https://doi.org/10.5840/jphil2024121312

as a fairly non-controversial idea. However, philosophy being as it is, the claim that rational agents ought to optimize is not without controversy, and has been forcefully disputed. Two of the most prominent dissenting positions are worth discussing at the outset, in part so I may set them aside.

First, Michael Slote argues that a choice is justified if one "satisfices," or if one opts for what is "good enough."[2] Observe, however, that this position can be understood as an optimization exercise as well. This is because what one is optimizing for can be understood as the consideration "good enough." That is, what justifies the actual choice is not that the chosen alternative is good enough, simpliciter. Rather, we can understand the satisficer's claim as follows: what justifies an actual choice is that the chosen alternative is *at least as good* as its rivals with respect to what is "good enough." So the alternative is "most good enough." To be sure, satisficers might resist such an interpretation of their view as an optimization exercise. However, this resistance would, I think, collapse the satisficing view into a specific version of the second position that has most prominently been articulated and defended by Michael Stocker.[3]

This second position that would dispute the claim that rational agents ought to optimize basically argues that, contrary to what optimization prescribes, the justification of a choice need not appeal to the *relative* merits of an alternative that is chosen. Rather, the justification of a choice is based on the *absolute* merits of the chosen alternative. So, for example, a satisficer might say—and this is why I believe

---

considerations that we have to take into account to arrive at the overall ranking R. Further, let $\mathcal{P}$ denote the set of all non-empty subsets of X, and a menu A in $\mathcal{P}$, or a *choice situation* as I shall call it here, is a non-empty subset of X, that is, $A \in \mathcal{P}$ and $A \neq \varnothing$. It is from this set A that a choice or selection is to be made. Call the set of alternatives that is chosen or selected in A a *choice set* and denote it by $c(A)$. One example of a choice set is the optimal set. The set of optimal, or best, alternative(s) in a given choice situation is defined as follows: an alternative x in a situation A is a *best* alternative with respect to a binary relation R if, and only if, for all $y \in A$, xRy.

[2] See Michael Slote, *Beyond Optimizing: A Study of Rational Choice* (Cambridge, MA: Harvard University Press, 1989), pp. 17–18. This position draws on, but is ultimately distinct from, Herbert Simon's idea of *satisficing*. For Simon, satisficing is a positive theory that *explains* behavior, and to satisfice is to use a heuristic or rule of thumb that economizes on the limited reasoning powers of the human mind. Slote goes further and makes the more radical—and philosophically more interesting—claim that satisficing can be rational in an unqualified sense. That is, Slote claims that rational agents *ought to* satisfice.

[3] See, specifically, Michael Stocker, "Abstract and Concrete Value: Plurality, Conflict, and Maximization," in Ruth Chang, ed., *Incommensurability, Incomparability, and Practical Reason* (Cambridge, MA: Harvard University Press, 1997). See also Elizabeth Anderson, "Practical Reason and Incommensurable Goods," in Chang, ed., *Incommensurability, Incomparability, and Practical Reason*, *op. cit.*, pp. 90–109, for a defense of a similar position.

that their view can collapse into the second position—that one is justified in going for an alternative that is "good enough," simpliciter (rather than going for an alternative that would qualify as being the "most good enough").[4] Or, one is justified in going for *x* because *x* is "pleasurable," simpliciter (rather than being the "most pleasurable" among its rivals). Or, one is justified in going for *x* because *x* is "beautiful," simpliciter (rather than being the "most beautiful" among what could have been chosen instead). In a more humorous vein, the relevant contrast involved here is that when asked the question, "Are you happy with your partner?," proponents of this position are not going to respond with the economists' quip: "Relative to what?"

More seriously, however, I wish to note at the outset that to proponents of this view, there is not much in what follows that will be of interest. This is because I shall not be presenting a general argument for an approach to rational decision-making that is based on the relative—rather than absolute—merits of an alternative.[5] Indeed, this is why these views that dispute the idea that rational agents ought to go for the best will not make an appearance in the claims and arguments that will be made in this paper. Rather, the line of reasoning that is presented here can be read as part of a debate that is internal to approaches to rational decision-making that assume the truth of the claim that the relative merits of an alternative justify a choice. Specifically, this paper asks those of us who believe that the relative merits instantiated in an alternative justify a choice the following question: can a rationally justified choice be made in a choice situation where the relative merits—represented by a binary preference or value relation—fail to establish the existence of an optimal al-

[4] To be sure, a satisficer may even dispute this understanding of their position, which collapses it into the second position. They may say instead that their view is what might be called a "threshold view," where "good enough" is some threshold along a chain of alternatives that are ranked according to their goodness. In this sense, satisficing would contrast with the claim that the absolute merits instantiated in an alternative are what justify a choice, as this understanding of satisficing appeals to the relative merits of an alternative. It would also contrast with optimization—despite being a view that appeals to the relative merits of an alternative—because it is not about going for the alternative that instantiates the "most good." The problem with this is of course a familiar one. There is no non-arbitrary way to specify what would count as a threshold.

[5] Sen has presented such a general argument for the domain of justice where the main claim is this. What we want from a theory of justice is a ranking of social states of affairs in terms of justice; specifically, we want a theory to tell us if some social state of affairs is more just than other social states of affairs. Crucially, this "comparative view" is contrasted with those theories of justice that seek to find the social state of affairs that would qualify as being "just," simpliciter. See Amartya Sen, "What Do We Want from a Theory of Justice?," this JOURNAL, CIII, 5 (2006): 215–38; and Amartya Sen, *The Idea of Justice* (Cambridge, MA: Harvard University Press, 2009).

ternative? More specifically, the question of interest here is this. If, in a given situation, you cannot go for an alternative that is the best with respect to your overall assessment of the alternatives, then is there no rationally justified choice for you to make in that situation? Put differently, are you facing a hard choice when you cannot optimize in a given situation?[6]

I shall be concerned with addressing this question here, and in order to do so we need to take note of the following. When a best alternative does not exist in a given situation, the reflexive binary preference or value relation that stands for a deliberating agent's all-things-considered evaluation of alternatives in the given situation will involve a violation of either acyclicity or completeness of the relation (or both).[7] Consequently, addressing our question will require scrutiny of two choice situations: namely, the class of situations involving a violation of acyclicity, as well as the class of situations involving a violation of completeness.[8] In what follows, it will be argued that: (i) the latter—the class of situations where a best alternative does not exist because of incompleteness of the binary preference or value relation—constitutes a hard choice; whereas (ii) the former—the class of situations where one cannot optimize because one holds a cyclic relation—does not constitute a hard choice. My defense of this pair of claims will proceed in three distinct steps.

[6] This pair of questions can be treated as invariant for two reasons. First, a best alternative is an optimal one—that is, an alternative that is at least as good as every other alternative that could have been chosen instead. Indeed, this paper will use 'optimal' and 'best' interchangably. Second, a hard choice is understood, at the most general level, as a situation in which there is no rationally justified choice for an agent to make.

[7] Completeness requires that for every pair of distinct alternatives in the finite set of alternatives X, either one alternative is *at least as good as* the other, the reverse, or both. Incompleteness of the preference or value relation is just the negation of this. Acyclicity requires that for all sequences x_1 to x_n in X, if x_1 is strictly better than x_2 and x_2 is strictly better than x_3 and so on till the end of the chain where x_{n-1} is strictly better than x_n, it should follow, then, that it is not the case that $x_n P x_1$. When a binary relation violates this constraint, we say that it is *cyclic*. The claim that a best or optimal alternative will remain undefined whenever there is a violation of either acyclicity or completeness of the relation (or both) follows from a well-known result due to Sen that states: given a reflexive relation that is complete, acyclicity is necessary and sufficient to establish that a best alternative exists in a given situation. See Amartya Sen, *Collective Choice and Social Welfare* (Cambridge, MA: Harvard University Press, 2018), lemma 1*l.

[8] A clarification is important to make in this context. It is no part of the claims being made here to argue for the existence of an incomplete relation, or a cyclic one. The line of reasoning presented here will simply assume that a binary preference or value relation that forms the basis of an actual choice can preclude the possibility of optimizing, or going for the best, because the binary relation is incomplete or cyclic. The question being investigated here is whether this would present a problem for practical reasoning.

First, it will be shown that the basic problem that needs to be addressed when optimization fails is not that a choice set—or *action-guiding proposal*, as I shall call it here—is undefined.[9] Rather, the problem that needs to be addressed is whether the choice sets or action-guiding proposals that can be invoked when the optimal set is empty are *justified* as a basis of rational decision-making. The second step is the heart of the paper, and it defends the main claims being advanced here. More specifically, this step of the argument starts by presenting two ways, due to Sen, in which we can scrutinize distinct action-guiding proposals, to wit: *case-implication scrutiny* and *prior-principle scrutiny*. Then, on these two grounds, it is argued that an action-guiding proposal to deal with the problem of decision-making with a cyclic relation is justified, while the most prominent proposal to deal with decision-making with an incomplete relation is not justified. Third, and finally, I consider and respond to a pair of objections that may be presented against the claims being defended here. A final section concludes.

I. THE PROBLEM OF ACTION GUIDANCE

I start by noting that the general problem that needs to be addressed in the context of our discussion of optimization is this: can a cyclic or incomplete relation be *action-guiding*? That is, if I cannot optimize because I do not hold an acyclic and complete evaluation of alternatives, then is there a defensible answer to the central question of practical reasoning which, recall, is this: "What, all things considered, should I do?" When the problem is formulated in this way, we can see that there are in fact two questions associated with the problem of action guidance. First, is it *possible* to give an answer to the central question of practical reasoning when one cannot optimize? I interpret this to mean the following: is there an action-guiding proposal that is defined when a deliberating agent holds a cyclic or incomplete evaluation of alternatives? Second, if such an answer exists—or, if such an action-guiding proposal is indeed defined—can it also be *justified* as the basis of rational decision-making?

Notice that—because ought implies can—a justification (or an answer to the second question) would not get off its feet without answering the first question affirmatively. Happily, this part of the problem of action guidance is not particularly serious for the class of situations under scrutiny here, and the following two claims are decisive in estab-

[9] Recall from footnote 1 that a choice set is the set of alternatives that are chosen or selected by the agent in a given choice situation. Optimization yields one example of a choice set, the optimal set or the set of best elements.

lishing why. First, Sen has presented an influential action-guiding proposal that deals quite explicitly with the problem of rational decision-making with an incomplete ranking. Second, an action-guiding proposal can also be defined when an agent holds a cyclic evaluation of alternatives. In fact, quite a few choice sets have been proposed by scholars interested in providing analytic foundations for decision-making with a cyclic relation.[10] The balance of this section will introduce, in turn, Sen's action-guiding proposals as well as a specific proposal to deal with the problem of decision-making with a cyclic relation.

Before we proceed, however, it is worth pointing out that a canonical action-guiding proposal that can be invoked in situations where agents cannot optimize is due to Isaac Levi, and it is called V-admissibility.[11] Per this proposal, an agent is seen as having two or more cardinal functions, each assigning an index number to every option. Each of these cardinal functions represents a different "permissible" way of evaluating alternatives. Further, every linear combination of permissible functions would also qualify as a permissible evaluation of alternatives. Now, an option is V-admissible if it is optimal according to at least one of these functions. Levi argues that a rational agent ought to restrict their choice to the set of V-admissible alternatives, but within this set, any choice is rationally permissible.

I bring up Levi's proposal in order to clarify two points concerning the scope of my claims. First, the question of whether Levi's proposal can be justified is not an issue that I shall pursue here. Levi's proposal introduces more mathematical structure than the two action-guiding proposals that I shall scrutinize here. The normative evaluation of this structure—which includes cardinality and convexity—would risk overburdening my analysis, which is restricted to ordinal structures of preferences or values. Second, despite the different formal frameworks, I note some congruence between Levi's claims and the claims that I shall defend here, which is this. Sen's action-guiding proposal does not give sound counsel. Indeed, this is one of the most luminous implications of Levi's analysis of hard choices, and it is hoped that the analysis presented here will provide further support for this message. But what is Sen's proposal?

[10] For a broad overview, see John Duggan, "A Systematic Approach to the Construction of Non-empty Choice Sets," *Social Choice and Welfare*, xxviii, 3 (2007): 491–506. See also Athanasios Andrikopoulos, "On the Construction of Non-empty Choice Sets," *Social Choice and Welfare*, xxxviii, 2 (2012): 305–23.

[11] See chapters 5 and 6 of Isaac Levi, *Hard Choices: Decision Making under Unresolved Conflict* (Cambridge, UK: Cambridge University Press, 1986), for a more systematic presentation of the idea of V-admissibility as it applies to hard choices.

Sen's action-guiding proposal to deal with the problem of decision-making with an incomplete relation is based on a departure that he makes quite early on in his classic work, *Collective Choice and Social Welfare*. Sen departs from optimization by presenting a weakening of optimization that is called maximization. With characteristic force and eloquence, Sen has gone on to argue that the maximal set—as opposed to the optimal set—should be the grounds of rational decision-making.[12] That is, Sen has advanced and defended the claim that in any given situation, a deliberating agent should identify—and opt for—a maximal alternative. A maximal alternative is one that is not strictly worse than any other available alternative.[13] Crucially, a complete ranking or evaluation of alternatives is not necessary to establish the existence of a maximal alternative; a reflexive binary preference or value relation that is acyclic is both necessary and sufficient to establish that a maximal alternative exists.[14] This is why Sen believes that there is a rationally justified decision that a deliberating agent can make when she cannot optimize because her ranking of the alternatives is incomplete.[15]

It is worth noting that the maximization view of practical reasoning can be given two related but distinct interpretations. The first is the one we have just discussed, where qualifying as an alternative that is not strictly worse than any other is a sufficient basis for rational decision-making. A second interpretation that has been suggested by Sen—and writers inspired by Sen, like Gerald Gaus—involves treating every maximal alternative *as if* it is optimal.[16] The basic idea here is

[12] See, *inter alia*, Amartya Sen, "Maximization and the Act of Choice," *Econometrica: Journal of the Econometric Society*, LXV, 4 (1997): 745–79; Amartya Sen, "Incompleteness and Reasoned Choice," *Synthese*, CXL, 1 (2004): 43–59; Amartya Sen, "Reason and Justice: The Optimal and the Maximal," *Philosophy*, XCII, 1 (2017): 5–19; and Amartya Sen, "The Importance of Incompleteness," *International Journal of Economic Theory*, XIV, 1 (2018): 9–20.

[13] A clarification is in order. There are at least three ways in which an alternative can qualify as being not strictly worse than any other available alternative. It can be because the alternative is strictly better than, equally as good as, or unranked with respect to any other alternative that could have been chosen instead. In the first two cases, the distinction between optimization and maximization disappears. Indeed, in these cases the set of maximal alternatives is isomorphic with the set of best alternatives. It is the last case that is relevant for the arguments presented here, to wit: choice situations where an alternative qualifies as maximal because it is unranked with respect to other available alternatives.

[14] See Sen, "Maximization and the Act of Choice," *op. cit.*, Theorem 5.2.

[15] For an argument inspired by Sen's views, see Nien-hê Hsieh, "Is Incomparability a Problem for Anyone?," *Economics and Philosophy*, XXIII, 1 (2007): 65–80.

[16] See, specifically, Jerry Gaus's use of this interpretation for debates in political philosophy (Gerald Gaus, *The Order of Public Reason: A Theory of Freedom and Morality in a Diverse and Bounded World* (New York: Cambridge University Press, 2010), pp. 307–08).

this: call R^+ the completed extension of the primitive binary preference or value relation R such that for any pair of unranked maximal alternatives x and y, xR^+y.[17] Thus, a maximization story can be interpreted to be a case of *as if* optimization with some suitably devised ranking R^+ that treats every unranked pair of maximal alternatives *as if* they are equally good as each other.[18]

In fact, a similar interpretation—of treating every alternative in a (top) cycle *as if* it is optimal—can be given for the action-guiding proposal being advanced here to deal with the problem of decision-making when the underlying relation is cyclic. To see this proposal, consider the triple $\{a, b, c\}$ and let the ranking of this triple be such that: a is strictly better than b, b is strictly better than c, and c is strictly better than a. This is to say that if you have to make a selection from either a, b, or c, then you are in a situation where no alternative can be considered a best alternative with respect to the overall preference or value relation that you hold. Nevertheless—and this is the proposed action-guiding proposal—every alternative a, b, and c can be treated symmetrically from the point of view of deciding what should rationally be done. That is, every alternative in the (top) cycle can be treated *as if* they are a best alternative. The proposal can be stated thus: when caught in a cycle, the deliberating agent treats every alternative in the cycle *as if* she holds them to be equally as good as every other. She can, therefore, arbitrarily pick any alternative in the (top) cycle. To put this in mildly mathematical terms, the proposed action-guiding proposal for cyclic relations is the optimal set defined with respect to the transitive closure R^* of the primitive preference or value relation R.[19]

The important point to get across here is this: There are no grounds to claim that an acyclic and complete relation is required for a choice to be *possible*, as some writers have claimed.[20] Put differently, the ques-

[17] See Sen, "Maximization and the Act of Choice," *op. cit.*, Theorem 5.4.

[18] The converse does not hold. That is, not every optimal choice can be interpreted as maximizing choice. To be more precise, for some binary preference or value relation R which generates a class of optimal choice sets, there may exist no binary relation R^+ such that the maximal choice set defined with respect to R^+ is equivalent to the optimal set defined with respect to R, for all choice situations. See *ibid.*, Theorem 5.5.

[19] In fact, when the binary relation is both cyclic and incomplete, we may use the maximal set defined with respect to the transitive closure in order to define an action-guiding proposal. So even when a binary preference or value relation instantiates these two violations together, an action-guiding proposal can still be the basis of a choice.

[20] See, for example, section 2.5 in Gustaf Arrhenius, Jesper Ryberg, and Torbjörn Tännsjö, "The Repugnant Conclusion," in Edward N. Zalta, ed., *The Stanford Encyclopedia of Philosophy* (Winter 2022 Edition), https://plato.stanford.edu/archives/win2022/entries/repugnant-conclusion/.

tion of whether you are facing a hard choice when you cannot optimize turns on the question of whether the action-guiding proposals discussed here are defensible or not. This is the basic problem that needs to be addressed when a rational agent cannot optimize, and I shall turn to address this now.

II. THE SCRUTINY OF ACTION-GUIDING PROPOSALS

II.1. A Methodological Preface. In light of the preceding discussion, the pair of claims that will be defended here can be restated as follows:

(i) The action-guiding proposal concerning decision-making with a cyclic relation can be defended (which is why the class of situations involving a cyclic evaluation of alternatives does not constitute a hard choice); and

(ii) Sen's proposal for decision-making with an incomplete relation cannot be defended (which is why the class of situations where one cannot optimize because of an incomplete evaluation of alternatives does indeed constitute a hard choice).

This, however, raises quite an obvious methodological question. How are we to scrutinize distinct choice sets, or distinct action-guiding proposals? That is, when such a pair of claims is made—that one action-guiding proposal is defensible but another is not—how are we to defend these claims?

Now, one way to proceed could involve presenting a formal characterization of the distinct action-guiding proposals that have been advanced to deal with the problem of decision-making when one cannot optimize, and subsequently evaluating the appeal of the conditions involved in such a characterization.[21] In fact, just such an argument for the action-guiding proposal that was sketched above—to deal with the

[21] The basic idea here is to show that an action-guiding proposal is isomorphic with certain "consistency" conditions that would *prima facie* be irrational to violate, like Sen's famous property α and property β. See Sen, *Collective Choice and Social Welfare, op. cit.*, section 1.6. This is why a characterization result constitutes an argument for an action-guiding proposal because it shows that following the proposal basically amounts to following certain properties that would appear irrational to violate. Note, however, that this argumentative strategy has been disputed by, among others, Robert Sugden and Amartya Sen. See, for example, Robert Sugden, "Why Be Consistent? A Critical Analysis of Consistency Requirements in Choice Theory," *Economica*, LII, 206 (1985): 167–83; and Amartya Sen, "Internal Consistency of Choice," *Econometrica: Journal of the Econometric Society*, LXIII, 3 (1993): 495–521. Their arguments, however, can be interpreted as defending the general claim being made here because—and Robert Sugden is actually explicit about this—their arguments can be interpreted as an argument against viewing constraints on a binary preference or value relation like acyclicity as rationally required.

problem of rational decision-making with a cyclic relation—has been presented in the literature.[22]

A second way to proceed could involve demonstrating that the strongest arguments to dispute the pair of claims being made here are not sound arguments. Just such a demonstration will be offered in the subsequent section—section III—of this paper.

Yet another distinct way—a third way—to proceed involves checking the implications of an action-guiding proposal by taking up particular cases in which the results of employing that proposal can be seen in a rather stark way, and then examining these implications against our intuitions or considered judgments. The basic test here involves checking whether a specific action-guiding proposal has counterintuitive implications. If it does, then the fact that it has such counterintuitive implications constitutes a reason to believe that the action-guiding proposal is unjustified (without further reasons to revise our intuitions). Sen calls this argumentative strategy *case-implication scrutiny*.[23]

Fourth and finally, we can also check the consistency of a proposal with another principle, or set of principles, that are—or should be—acknowledged to be even more fundamental. Sen calls this final argumentative strategy *prior-principle scrutiny*.[24]

The balance of this section will deploy Sen's two argumentative strategies in order to defend the claims that are being made here.

II.2. Case-Implication Scrutiny. I start with *case-implication scrutiny* of Sen's action-guiding proposal, to wit: maximization. To this end, consider a paradigmatic hard choice.

> *Sophie's Choice*: Sophie and her two children are prisoners in a Nazi concentration camp. A Nazi guard tells Sophie that *one*, and only one, of her two children will live, but Sophie will have to choose which one. So Sophie can save one of her children, but only by condemning the other to being killed. What, if anything, should Sophie do?

To paraphrase Lincoln on slavery, if this is not a hard choice, then nothing is a hard choice. However, Sen's action-guiding proposal—

[22] See, specifically, Martin van Hees, Akshath Jitendranath, and Roland Iwan Luttens, "Choice Functions and Hard Choices," *Journal of Mathematical Economics*, XCV (2021): 102479; see also Thomas Schwartz, "Rationality and the Myth of the Maximum," *Noûs*, VI, 2 (1972): 97–117; Thomas Schwartz, "Choice Functions, 'Rationality' Conditions, and Variations on the Weak Axiom of Revealed Preference," *Journal of Economic Theory*, XIII, 3 (1976): 414–27; and Rajat Deb, "On Schwartz's Rule," *Journal of Economic Theory*, XVI, 1 (1977): 103–10.

[23] See Amartya Sen, "Equality of What?," *The Tanner Lecture on Human Values*, I (1980): 197–220, at pp. 197–98.

[24] *Ibid.*

you can justifiably go for a maximal alternative—would imply that there is a rationally justified choice for Sophie to make. This, in turn, would imply that this situation does not constitute a hard choice. Sen has of course not explicitly articulated or defended this claim. It is nevertheless a conclusion that is implied by the maximization view of practical reasoning, and it is easy to see why. For intuitively, the alternatives in *Sophie's Choice* are unranked with respect to each other. Every alternative would thus qualify as being maximal. Therefore, if we take maximization to be a theory that offers a deliberating agent sound counsel, then it would follow that *Sophie's Choice* would not constitute a hard choice. The point is that this implication of employing Sen's action-guiding proposal is counterintuitive, and the fact that it is counterintuitive constitutes a reason to believe that Sen's action-guiding proposal fails.

To turn to decision-making with a cyclic relation, we can ask the same question. Does the action-guiding proposal that has been advanced to deal with this class of choice situations encounter such a problem of having to accept counterintuitive results? In order to address this question it will be useful to partition the cases that instantiate a cyclic relation into two broad classes. I discuss them in turn, and this discussion proceeds with the aim of showing that they do not encounter the problem of having to accept counterintuitive results.

The first class of cases that I wish to discuss is mostly found in the literature on the nature of value.[25] The aim here is to show the existence of a cyclic relation by appealing to our intuitive judgments about a specific case.[26] To explain, consider the example that Larry Temkin presents as the most convincing demonstration of the existence of a cyclical value relation. This is often called *Hangnails* for reasons that

[25] In this literature, Larry Temkin and Stuart Rachels have been among the most ardent and influential defenders of the claim that a value relation can violate acyclicity. See, *inter alia,* Larry S. Temkin, "Intransitivity and the Mere Addition Paradox," *Philosophy and Public Affairs,* XVI, 2 (1987): 138–87; Larry S. Temkin, "A Continuum Argument for Intransitivity," *Philosophy and Public Affairs,* XXV, 3 (1996): 175–210; Larry S. Temkin, *Rethinking the Good: Moral Ideals and the Nature of Practical Reasoning* (New York: Oxford University Press, 2014); Stuart Rachels, "Counterexamples to the Transitivity of *Better Than,*" *Australasian Journal of Philosophy,* LXXVI, 1 (1998): 71–83; and Stuart Rachels, "A Set of Solutions to Parfit's Problems," *Noûs,* XXXV, 2 (2001): 214–38.

[26] In this context it is worth noting that most of these examples aim to show a violation of what is called *quasi-transitivity*. That is, their explicit aim is to show that the "strictly better than" relation—or the asymmetric part P of the primitive binary relation R—is not transitive, rather than showing that the primitive relation violates acyclicity. Indeed, the distinction between quasi-transitivity and acyclicity is often ignored in this literature. But following Erik Carlson, I interpret these examples as violations of acyclicity for, as Carlson writes, "it seems implausible that betterness is sometimes intransitive but never cyclic." See Erik Carlson, "Intransitivity," in Hugh LaFollette, ed., *International Encyclopedia of Ethics* (Malden, MA: Wiley, 2013), p. 2.

will presently become clear, and it involves ranking a number of possible lives *A*, *B*, . . . , *Z*. Life *A* contains two years of excruciating torture, while *B* contains four years of slightly less intense pain. Similarly, *C* contains eight years of pain that is slightly less intense than that in *B*. And so on down the alphabet. Life *Z* contains millions of years of very mild pain, such as a hangnail. According to Larry Temkin, most would rank *A* strictly better than *B*, *B* strictly better than *C*, and so on down to *Z*. Acyclicity requires that *Z* is not strictly better than *A*. Temkin finds it clear, however, that *Z* is strictly better than *A*. Hence, acyclicity of value does not hold.[27] A graphical representation of this case is presented in Figure 1 below.

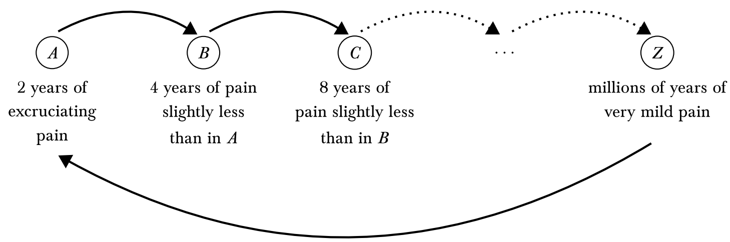

Figure 1. A graphical illustration of *Hangnails*.

There is, however, a problem with these cases. It is a matter of dispute as to whether such cases involve a violation of acyclicity. John Broome notes that such examples involve large numbers and argues that our intuitions about large numbers are unreliable. We may, for example, be unable to grasp what it would be like to have a hangnail for millions of years. Hence, the intuition that life *Z* in Hangnails is better than life *A* is unreliable.[28] Without questioning intuitions about large numbers, Ken Binmore and Alex Voorhoeve argue that such examples are based on a mathematical mistake.[29] It is also worth noting

[27] See Temkin, "A Continuum Argument for Intransitivity," *op. cit.*, p. 180.

[28] See John Broome, *Weighing Lives* (New York: Oxford University Press, 2004), p. 51.

[29] The mistake, according to Binmore and Voorhoeve, is that these purported cases of violations of quasi-transitivity (which requires that the asymmetric part of the primitive relation is transitive) or acyclicity are in fact cases of intransitive indifference, as in the famous Zeno's paradox. If they are correct, the upshot is that an optimal alternative is defined for these cases; they show this by demonstrating that a utility function—specifically, the familiar Cobb-Douglas utility function—is defined for these cases. See Kenneth G. Binmore and Alex Voorhoeve, "Defending Transitivity against Zeno's Paradox," *Philosophy and Public Affairs*, xxxi, 3 (2003): 272–79. To be sure, this argument has been challenged by Erik Carlson; see Erik Carlson, "Intransitivity without Zeno's Paradox," in Toni Rønnow-Rasmussen and Michael J. Zimmerman, eds., *Recent Work on Intrinsic Value* (Dordrecht, the Netherlands: Springer, 2005), pp. 273–77.

that such examples rely on controversial assumptions about how we must think of trade-offs between intensity of pain against duration of pain.[30] The relevant upshot for our discussion is this: Because it remains a matter of dispute as to whether these cases in fact involve a cyclic relation, it remains unclear as to whether these cases are apt to examine our intuitions about the action-guiding proposal that has been advanced to deal with the problem of decision-making with a cyclic relation.

The second class of cases that I wish to discuss is mostly found in the literature in economic theory and decision theory. Unlike the first class of cases that we have just discussed, the instantiation of a cyclic preference or value relation is not a matter of dispute for this class of cases. However—and this is a crucial point to emphasize—these cases show that a cyclic relation can be an *implication* of accepting some "prior principles" that appear to be rationally justified, like majority rule, or Amos Tversky's lexicographic semi-order, among others.[31] Indeed, this is why the possibility of an agent holding a cyclic relation is not a matter of dispute for this class of cases.[32] The relevant lesson, then, from this class of cases is an important one, and it is this. They already constitute a "prior principle" case—albeit not a novel case—for viewing the action-guiding proposal being defended in the context of decision-making with a cyclic relation. This is because the proposal is consistent with some principles that appear to be rationally justified.

While we shall presently turn to a novel prior-principle case for the action-guiding proposal being defended here, the more general point that emerges from our discussion of the two classes of cases is an important one. They show us that there are no straightforward cases of situations involving cyclic relations against which we can examine our intuitions, and recognizing this point constitutes a reason to believe that the action-guiding proposal being defended here is justified. This is because there are no counterintuitive implications we have to accept with this proposal.

[30] In fact, Temkin acknowledges this, and chapter 5 of Temkin, *Rethinking the Good, op. cit.*, discusses this issue extensively.

[31] The classic references here are, of course, Kenneth O. May, "Intransitivity, Utility, and the Aggregation of Preference Patterns," *Econometrica: Journal of the Econometric Society* (1954): 1–13, for majority rule; and Amos Tversky, "Intransitivity of Preferences," *Psychological Review*, LXXVI, 1 (1969): 31–48, for the lexicographic semi-order.

[32] For a discussion of this class of cases, see the classic surveys of Paul Anand, "The Philosophy of Intransitive Preference," *The Economic Journal*, CIII, 417 (1993): 337–46; and Maya Bar-Hillel and Avishai Margalit, "How Vicious Are Cycles of Intransitive Choice?," *Theory and Decision*, XXIV, 2 (1988): 119–45. See also Philippe Mongin, "Does Optimization Imply Rationality?," *Synthese*, CXXIV, 1–2 (2000): 73–111, who presents a brief survey of cyclic relations before arguing for the very same position being defended here; that is, cyclic relations do not entail irrationality.

II.3. Prior-Principle Scrutiny. I turn now to prior-principle scrutiny of the action-guiding proposals that are under consideration here. So recall that prior-principle scrutiny of an action-guiding proposal involves checking the consistency of the proposal with another principle that is—or should be—acknowledged to be even more fundamental. Now, one such principle which I take to be more fundamental is the following condition that has been left as an unstated assumption throughout the exposition.

$\mathcal{E}$: Hard choices exist.[33]

I take it to be the case that $\mathcal{E}$ is true not just because of the intuitive plausibility of many examples, like *Sophie's Choice*. A theory of rational decision-making, I assume, should be able to show how it would analyze hard choices irrespective of whether they actually exist. Indeed, denying $\mathcal{E}$ by assumption, and without offering very decisive considerations for such an assumption, involves begging the question in its favor. But consider, now, in addition to $\mathcal{E}$, the following principle that (as far as I know) is new to the literature.

Existence of Reasons for Choice (*ERC*): Let A be a set of alternatives and y be an element that is not in A. If A does not constitute a hard choice, and if for any x that is selected or chosen from A (that is, any $x \in c(A)$) the choice between x and y is not a hard choice, then adding y to A (that is, $A \cup \{y\}$) is also not a hard choice.

A graphical representation of *ERC* is presented in Figure 2 below. It shows that if x is selected from A, and if the choice between x and some element y that is not in A does not constitute a hard choice, then adding y to A cannot make the new situation a hard choice.

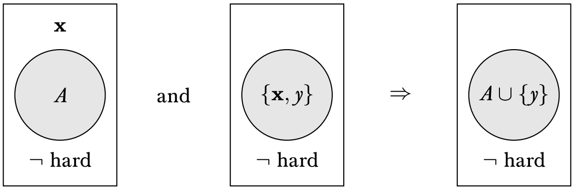

Figure 2. A graphical illustration of *ERC*.

But why, a skeptical reader may ask, should *ERC* also count as more fundamental along with $\mathcal{E}$? To answer this question, I start by taking the following claim by Ruth Chang—the most prominent contempo-

[33] To state this in mildly mathematical terms, for some $A \in \mathcal{P}$, $c(A) = \varnothing$.

rary scholar of hard choices—as a point of departure: "Hard choices are ones in which reasons 'run out': they fail, in some sense, to determine what you should do."[34] Now, if this is correct, an investigation of whether the class of situations that are under scrutiny here constitutes a hard choice can proceed by asking the more fundamental question of whether or not reasons "run out" in these situations. Adjudicating this question, however, requires a plausible condition that can adjudicate the question at hand, to wit: whether reasons have run out in a given situation. The claim being made here is that *ERC* is such a plausible condition, and it is made here on the following grounds.

If a choice situation *A* does not form a hard choice, then choosing among the elements of *A* is something that falls within "the domain of reasons." There are reasons for an agent to recognize and respond to in the situation *A*. Now consider a new element *y*. If there is a reason to make a choice between any of the chosen elements in *A* and *y*, then by adding *y* to *A* we do not leave the domain of reasons: an agent has not "run out" of reasons to recognize and respond to in the new situation in which they find themselves. Put differently, the addition of the new element cannot create a hard choice.

An example might also help getting across the intuitive appeal of *ERC*. Consider the situation where Dr. Anne and her nurse Bob are engaged with the issue of how they should diagnose a patient in their care—tubercular or not?[35] Assume, now, that after all the relevant evidence is in, and indeed, after careful scrutiny of the relevant evidence has taken place, Dr. Anne comes to the conclusion that the patient has tuberculosis. Suppose, however, that Bob asks Dr. Anne, "Have you considered whether the patient has COVID-19 instead?" *ERC* merely states that Dr. Anne cannot be facing a hard choice—nor indeed be in a situation where she has no rationally justified cure protocol with which to proceed—if (given the initial diagnosis of tuberculosis) she can in fact adjudicate the question of whether the present diagnosis should be tuberculosis or COVID-19.

Now, my claim is that Sen's action-guiding proposal is inconsistent with *ERC* and $\mathcal{E}$. In order to see this, consider the following principle—(M)—that formalizes Sen's action-guiding proposal.[36]

[34] See Ruth Chang, "Are Hard Choices Cases of Incomparability?," *Philosophical Issues*, XXII (2012): 106–26, at p. 112.

[35] This example is an adaptation from the classic R. Duncan Luce and Howard Raiffa, *Games and Decisions: Introduction and Critical Survey* (New York: Wiley, 1957), pp. 288–89.

[36] To be a little pedantic, this principle would follow from two premises. First, if a maximal alternative exists in a given situation, then there is a rationally justified choice that an agent facing the situation can make. This is of course the main claim involved

(M): If a maximal alternative exists in a given choice situation, then the choice situation does not constitute a hard choice.

I now present the proposition on which my criticism of Sen's action-guiding proposal relies.

Proposition 1. For any choice situation $A \in \mathcal{P}$, *ERC*, $\mathcal{E}$, and (M) cannot jointly be true.

The proof of this proposition can be found in appendix A. Here I note an implication of the proposition that is relevant for our scrutiny of Sen's action-guiding proposal. What this proposition shows is that *ERC* and (M) together violate $\mathcal{E}$—they rule out the existence of hard choices. Put differently, if one were to endorse *ERC* and $\mathcal{E}$, as I have argued that we should, then one would have to reject (M). This implication constitutes a reason to believe that Sen's action-guiding proposal fails. This is because it shows that the proposal is incompatible with principles—$\mathcal{E}$ and *ERC*—that should count as more fundamental.

We turn now to show why the choice situation involving decision-making with a cyclic relation would not constitute a hard choice. The argument for this will appeal to Sen's famous property γ, which, intuitively, requires that if one of the university's best teachers in nonclassical logic is also one of its best teachers in classical logic, then she is one of its best logic teachers.[37] More formally, it is defined as follows:

Property γ: For all $A, B \in \mathcal{P}$, $c(A) \cap c(B) \subseteq c(A \cup B)$.

In addition to property γ, my argument also relies on the following property that (as far as I am aware) is new to the literature:

Property κ: Let A be a set of alternatives, and let x and y be distinct alternatives. If x is selected or chosen from A (that is, $x \in c(A)$), but if y is selected or chosen in the situation involving only x and y (that is, $y \in c(\{x, y\})$), then y is selected or chosen in the situation where we add y to A (that is, $y \in c(A \cup y)$).

A graphical representation of κ is given in Figure 3 below.

Now, property γ is quite well known, and I shall not say more about why it should count as a more fundamental principle. Property κ,

in Sen's action-guiding proposal. Second, if there is a rationally justified choice that an agent facing a choice situation can make, then the choice situation does not constitute a hard choice. Together, these two premises entail the principle.

[37] See Sven Ove Hansson and Till Grüne-Yanoff, "Preferences," in Edward N. Zalta, ed., *The Stanford Encyclopedia of Philosophy* (Spring 2022 Edition), https://plato.stanford.edu/archives/spr2022/entries/preferences/, footnote 4.

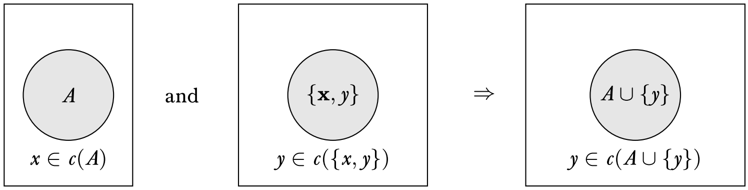

Figure 3. A graphical illustration of property κ.

however, is (as far as I am aware) new to the literature, and it is worth some defense before we proceed. So observe, first, that much of the defense for accepting *ERC* would also apply to κ. For both conditions help us adjudicate whether we have left the "domain of reasons." Indeed, κ requires that if you go for x from a given situation A (and are therefore within the domain of reasons), but if you can see yourself going for y in the choice between x and y, then you can see yourself going for y in the situation where we add y to A (and therefore do not leave the domain of reasons). Rather than restate this case for κ, I will try to bring out the intuitive appeal of this property with an example. Property κ states that if you can see yourself studying economics in the situation where you could have chosen to study any other social science instead, but you can see yourself studying philosophy when the choice is between studying economics or studying philosophy, then you can see yourself studying philosophy in the situation involving philosophy as well as all the social sciences you could have studied instead.

The final principle on which my argument relies is the following principle—(O)—with which we began the paper.

(O): If an optimal alternative exists in a given choice situation, then the choice situation does not constitute a hard choice.

Recall, now, that the action-guiding proposal that has been invoked here to deal with the problem of decision-making with a cyclic relation is the set of optimal or best elements defined with respect to the transitive closure R^* of the primitive relation R. If this action-guiding proposal were to be justified, then it would lead us to the following principle, (O*).

(O*): If an optimal alternative defined with respect to the transitive closure R^* of R exists in a given choice situation, then the choice situation does not constitute a hard choice.

The following proposition, then, is the ground on which my case for (O*) rests.

Proposition 2. If κ, γ, and (O) are true, then (O*) must be true.

The proof of this proposition can also be found in appendix A. Indeed, the case made here for (O*) is that κ, γ, and (O) are defensible principles. Thus, given this proposition, (O*) follows from accepting these "prior principles."

In sum, this section introduced distinct ways in which we may scrutinize action-guiding proposals, including *case-implication scrutiny* and *prior-principle scrutiny*. On the basis of these two ways in which an action-guiding proposal may be scrutinized, it has been shown that: (i) the action-guiding proposal concerning decision-making with a cyclic relation can be defended (which is why the class of situations involving a cyclic evaluation of alternatives does not constitute a hard choice); and (ii) Sen's proposal cannot be defended (which is why the class of situations where one cannot optimize because of an incomplete evaluation of the alternatives does indeed constitute a hard choice). I will now conclude my case for these claims by considering, and responding to, a pair of objections.

III. OBJECTIONS TO THE ARGUMENT

III.1. An Argument for Maximization and Why It Fails. Perhaps the most forceful argument to view maximization as a justified action-guiding proposal is Sen's rhetorically powerful illustration of the story of Buridan's ass. It is worth quoting this appeal to our intuitions in full.

> This is the tale of the donkey that dithered so long in deciding which of the two haystacks x or y was better, that it died of starvation z. There are two interpretations of the dilemma of Buridan's ass. The less interesting, but more common, interpretation is that the ass was indifferent between the two haystacks, and could not find any reason to choose one haystack over the other. But since there is no possibility of a loss from choosing either haystack in the case of indifference, there is no deep dilemma here either from the point of view of maximization or of that of optimization. The second—more interesting—interpretation is that the ass could not rank the two haystacks and had an incomplete preference over this pair. It did not, therefore, have any optimal alternative, but both x and y were maximal—neither known to be worse than any of the other alternatives. In fact, since each was also decidedly better for the donkey than its dying of starvation z, the case for a maximal choice is strong. Optimization being impossible here, I suppose we could 'sell' the choice act of maximization with two slogans: (i) maximization can save your life, and (ii) only an ass will wait for optimization.[38]

[38] See Sen, "Maximization and the Act of Choice," *op. cit.*, p. 765.

Now, I believe that a fair representation of the argument underlying the tale of Buridan's ass is as follows: The opportunity costs involved with not picking either stack of hay—death by starvation—is exorbitant. There are two related points to make here.

First, why are these costs not a reason for ranking the alternatives symmetrically, rather than evaluating them to be unranked? The point is that when one takes into account the opportunity costs in the tale of Buridan's ass, one recognizes that it is Sen's reinterpretation of this story that cannot be sustained. This is not a tale of facing unranked alternatives, but of having to decide between symmetrically ranked ones because of the opportunity costs associated with the alternatives in the given choice situation. This brings us to the following twist of Sen's conclusion: (i) correctly assessing alternatives can save your life, and (ii) only Buridan's ass will view this situation as involving two incomparable alternatives.

Second, and more generally, it is easy to see that one can present the opportunity cost argument as a consideration in favor of the second interpretation of the maximization story discussed above (as well as the action-guiding proposal discussed in the context of cyclic relations above). That is, when an agent recognizes the opportunity costs associated with unranked alternatives (or, for that matter, even alternatives ranked by a cyclic relation) the deliberating agent has a consideration for evaluating every unranked (or cyclically ranked) alternative *as if* each is equally as good as any other. There is an undeniable elegance and plausibility to this standard economic argument. It would also have the happy—if implausible—conclusion that there are no hard choices. However, before responding to it, and indeed, in order to respond to it, I wish to make the following clarification.

This argument from opportunity costs has moved us from an evaluative universe involving full information to a partial-information universe, and it is important to recognize this. For if we are in a full-information universe, then the opportunity costs associated with the alternatives in a choice situation are part of the overall evaluation of these alternatives that has brought us to an unranked (or cyclical) ranking of the alternatives. Indeed, this argument from opportunity costs seems to be the following. A recognition of the (as yet unconsidered) opportunity costs associated with every unranked (or cyclically ranked) alternative provides the deliberating agent with a consideration for evaluating these alternatives *as if* they are equally as good as each other. Consequently, the deliberating agent is justified in arbitrarily going for any maximal (or cyclically ranked) alternative.

Having made this clarification, there are two points that constitute my rejoinder to this argument. First, new information or further reflection, like recognizing opportunity costs associated with alternatives, may provide a consideration that allows us to make a rationally justified decision. This is not in dispute. What I wish to dispute is that the maximal alternatives remain unranked when this new information is available, or further reflection has taken place, and has provided a new consideration that allows a deliberating agent to make a warranted choice. The point is a generalization of the response to Sen's reinterpretation of the tale of Buridan's ass, and it is this. Behind what *prima facie* appears to be a situation involving unranked maximal alternatives in a partial information universe, the further information that has been introduced by taking cognizance of the opportunity costs has resolved the incompleteness associated with the ranking of the alternatives. We are justified in going for an alternative after taking the opportunity costs into account *because* the situation has been transformed into one where an optimal alternative does in fact exist. Second—and this is admittedly a more practical concern rather than an analytic one—in non-market environments, the opportunity costs associated with alternatives in a choice situation cannot always be identified—much less quantified—before a decision needs to be made. The point is that the argument does not operate over the entire domain of choice situations being investigated here. An example may help illustrate this.

Imagine that you are a doctor having to treat patients infected with COVID-19, which has no known cure protocol. That is, there is no best alternative that works as a therapeutic drug. Nevertheless, you are aware of the fact that a cocktail of drugs, including the anti-malarial chloroquine and hydroxychloroquine, are being used as a therapeutic for treating patients with COVID-19. Further, you know that the efficacy of these drugs is being investigated in clinical trials. However, clinical data are still very limited and inconclusive, and the beneficial effects of these medicines as a cure protocol for COVID-19 have not been demonstrated. You also know that some studies have reported serious—and in some cases fatal—heart-rhythm problems with chloroquine or hydroxychloroquine, particularly when taken at high doses or in combination with the antibiotic azithromycin. In addition to side effects affecting the heart, they are known to potentially cause liver and kidney problems, nerve-cell damage that can lead to seizures, and hypoglycemia. Should you, as a doctor, prescribe such a cocktail of drugs that includes chloroquine or hydroxychloroquine? Assume that there is no better alternative, but you cannot rank prescribing such a cocktail of drugs against not prescribing such a cocktail as a therapeutic. These alternatives are maximal alternatives in a partial-information universe. But what are the opportunity costs associated with opting for either alternative?

The upshot of recognizing that opportunity costs are not immediately identifiable or measureable in non-market environments is this. Even when the incompleteness associated with unranked alternatives is resolvable with more information, like opportunity costs, or some other means of discovery, until such a resolution has taken place, we are not justified in treating unranked (or even cyclically ranked) alternatives *as if* they are equally as good as each other. To the contrary, when such a resolution of incompleteness is possible, the relevant lesson is that we need to pursue further inquiry, rather than ending inquiry and treating alternatives *as if* they are equally as good as each other. We may therefore set aside the opportunity-cost argument in favor of Sen's maximization view of practical reasoning.

III.2. An Argument against Cyclicity and Why It Fails. Consider, now, the canonical argument which gets presented to establish that acting on the basis of a cyclic relation is a profound error on the part of the deliberating agent—namely, the money pump. This argument relies on two premises.[39] First, an agent who acts on the basis of a relation that violates acyclicity is vulnerable to exploitation by a money pump. To illustrate using Amos Tversky's example:

> Suppose an individual prefers y to x, z to y, and x to z. It is reasonable to assume that he is willing to pay a sum of money to replace x by y. Similarly, he should be willing to pay some amount of money to replace y by z and still a third amount to replace z by x. Thus, he ends up with the alternative he started with but with less money.[40]

The conclusion that this is irrational is secured with the further premise which states that such vulnerability to being exploited by a money pump implies irrationality. Indeed, this second premise is the main premise of the argument, and it basically means that it is irrational to knowingly pay (in some currency that you care about) for what you could have kept for free. There is, of course, the following pair of influential objections to this line of reasoning: (i) that you could rationally avoid being money-pumped if you use foresight;[41] and

[39] There is, by now, an enormous literature on this famous argument, and the various positions involved in this literature are helpfully and critically presented in Johan E. Gustafsson, *Money-Pump Arguments* (New York: Cambridge University Press, 2022).

[40] Tversky, "Intransitivity of Preferences," *op. cit.*, p. 45.

[41] The basic idea here is that an agent who knows that she is being taken for a ride can avoid exploitation if she uses backward induction. The classic reference here is Frederic Schick, "Dutch Bookies and Money Pumps," this JOURNAL, LXXXIII, 2 (1986): 112–19. Note, however, that while foresight blocks the standard version of the money-

(ii) that you could rationally avoid being money-pumped if you are resolute and stick to a plan.[42]

Note, crucially, that both these objections are concerned with blocking the agents' vulnerability to exploitation by a money pump, rather than challenging the claim that it would be irrational if a deliberating agent is vulnerable to such exploitation. That is, they are directed at the first premise of the money-pump argument. I wish, however, to challenge the second premise here. That is, I wish to challenge the claim that vulnerability to such exploitation implies irrationality. Indeed, in the balance of this section I will argue that the implication involved in the main premise of the money pump fails for two distinct reasons.

First, as Wlodek Rabinowicz has noted, when we want to know whether a choice is rationally justified or not, we can consider two perspectives.[43] We can ask—and this is the first perspective—how successful the agent will be, by their own lights, if they act on the basis of the binary relation that they hold, and the situation in which they find themselves. This may be called the *pragmatic perspective*, and it is implicit in the claim that it is irrational to be vulnerable to exploitation (because the money pump sets out to show that violating some constraint would be to the agent's disadvantage by their own lights). This pragmatic perspective is, however, not the only one that is available to us to adjudicate the rationality of a binary preference or value relation that an agent holds in a given situation. A second perspective is this. We can also ask, given the situation, whether this binary relation is well-grounded in intuitive judgments and more general principles. Plainly, this perspective was implicit in the claim defended above that cyclical relations are rational as they are consistent with property κ and property γ. In this context it is worth stating that the argument given above follows arguments in the philosophy and psychology of

---

pump argument, there are other versions that work for agents who use foresight. See section 2.1 in Gustafsson, *Money-Pump Arguments, op. cit.*, and the enormous literature cited there for a critical *tour d'horizon* of this objection to the money-pump argument.

[42] To be resolute is to choose in accordance with plans one has adopted even if one would not actually choose in accordance with those plans if one had not adopted them. The classic reference here is Edward F. McClennen, "Prisoner's Dilemma and Resolute Choice," in Richmond Campbell and Lanning Sowden, eds., *Paradoxes of Rationality and Cooperation* (Vancouver: University of British Columbia Press, 1985), pp. 94–104. But crucially, see section 7 in Gustafsson, *Money-Pump Arguments, op. cit.*, for a discussion of why this objection does not succeed.

[43] The discussion here—and especially the distinction between the two perspectives that are introduced—is indebted to Wlodek Rabinowicz's discussion of the second premise of the money-pump argument. See section 6 of Wlodek Rabinowicz, "Money Pump with Foresight," in Michael J. Almeida, ed., *Imperceptible Harms and Benefits* (Dordrecht, the Netherlands: Springer, 2000), pp. 123–54.

decision theory which show that a binary preference or value relation which can be exploited by a money pump is well grounded (or consistent) with intuitive judgments as well as more general principles. Indeed, this literature abounds with examples of cyclical choices that appear to be rationally defensible, and therefore incompatible with the main premise of money-pump arguments.[44]

Second, and more importantly, the money-pump arguments prove too much: even a preference or value relation that would be sufficient for making what seems quite obviously to be a rationally justified choice could leave the agent who holds that binary relation vulnerable to exploitation by a money pump.[45] This is because the main premise of the money-pump argument conflicts with a far more plausible claim, to wit: optimization (or opting for the best) implies rationality. To see why, suppose that x, y, and z are tradable goods and that an agent prefers x to y but is indifferent between y and z, and also between x and z. Plainly, x and z qualify as best elements: they are at least as good as any other alternative that can be chosen instead. Nevertheless, the agent who holds such a binary preference or value relation is vulnerable to being exploited by a money pump. For starting out with y, our agent should then be willing to pay some amount of money, e, to switch to x. Having switched to x, and being indifferent between x and z, our agent should next be willing to switch to z, when offered a small monetary bonus, say $e/4$. Having switched to z, and being indifferent between z and y, the agent should then also be willing to switch to y if she is again paid $e/4$. But then she is back with y, having lost $e/2$, which makes the agent vulnerable to exploitation. We thus have a conflict between the following pair of claims: (i) optimization (or opting for the best) implies rationality, and (ii) vulnerability to exploitation by a money pump implies irrationality. Now, because I believe that it is a platitude to assert (i), we may conclude that (ii)

[44] See, for example, the surveys in Anand, "The Philosophy of Intransitive Preference," *op. cit.*; and Bar-Hillel and Margalit, "How Vicious Are Cycles of Intransitive Choice?," *op. cit.*

[45] The claim that money-pump arguments prove too much has been made most notably by Frank Arntzenius, Adam Elga, and John Hawthorne, "Bayesianism, Infinite Decisions, and Binding," *Mind*, CXIII, 450 (2004): 251–83. They defend this claim by showing that agents with preferences that are consistent with expected utility theory can nevertheless be vulnerable to being money-pumped. Indeed, this is why they say that money-pump arguments prove too much. Their setup, however, requires an infinite series of trades, or they concoct what have been called "infinite money pumps" to establish this point. See section 8 of Gustafsson, *Money-Pump Arguments*, *op. cit.*, for a critical overview of their arguments. In this context it is worth pointing out that my defense of the claim that money-pump arguments prove too much does not appeal to infinite money pumps.

is false. Put differently, the money-pump argument fails because its main premise cannot be sustained. We may therefore set aside this argument, which defends the view that acting on a cyclic relation is irrational.

IV. CONCLUSION

To the best of my knowledge, philosophers and economists who are interested in whether a rationally justified choice can be made when optimization fails have focused on either one of the two classes of situations where an optimal alternative is undefined.[46] Indeed, they have focused either on the class of situations where an optimal alternative is undefined because of incompleteness of the underlying preference or value relation that forms the basis of a choice, or on the class of situations where one cannot optimize because one holds a cyclic preference or value relation. Here I have investigated the issue more generally and argued that the latter class of situations does not constitute a hard choice, while the former does. Specifically, I have argued that an action-guiding proposal to deal with the problem of decision-making with a cyclic relation is justified, while the most influential action-guiding proposal to deal with the problem of decision-making with an incomplete relation is unjustified.

APPENDIX A. PROOFS OF PROPOSITIONS 1 AND 2

Let R be a reflexive binary relation defined on a finite set of objects X. That is, $R \subseteq X \times X$, and let R^* be the transitive closure of R. That is, xR^*y iff $\forall x, y \in X$, there are $x_1, \ldots, x_k \in X$: $xRx_1R \ldots Rx_kRy$. We let P and I denote the asymmetric and symmetric parts of R, respectively. P stands for *strictly better than*, and I stands for *equally as good as*, and they are defined in the usual way. That is, xPy iff $(xRy \wedge \neg(yRx))$, and xIy iff $(xRy \wedge yRx)$. Now, let $\mathcal{P}$ denote the set of all the non-empty subsets of X, and call an element A in $\mathcal{P}$ a *choice situation*. For any choice situation A, and for any binary relation R, we let $B(A, R)$ denote the optimal set, that is, the set of best elements in A, and define it as follows.

$$B(A, R) = \{x | x \in A \wedge \forall y \in A : xRy\}$$

For any choice situation A, and for any binary relation R, $M(A, R)$ denotes the set of maximal elements in A, and we define it as follows.

$$M(A, R) = \{x | x \in A \wedge \nexists y \in A : yPx\}$$

[46] An honorable exception that establishes the point is Paul Anand, "Are the Preference Axioms Really Rational?," *Theory and Decision*, XXIII, 2 (1987): 189–214.

Finally, let $c(A)$ denote a justified choice set, or a justified action-guiding proposal. A hard choice, therefore, is any choice situation $A \in \mathcal{P}$ such that $c(A) = \varnothing$. For ease of exposition, we now restate the conditions—$\mathcal{E}$, *ERC*, κ, γ (M), (O), and (O*)—that were introduced in the paper.

Condition 1. $\mathcal{E}$: There exists some $A \in \mathcal{P}$ that is a hard choice (that is, for some $A \in \mathcal{P}$, $c(A) = \varnothing$).

Condition 2. ERC: For any $A \in \mathcal{P}$, and $y \in X \setminus A$. If $c(A) \neq \varnothing$ and if for any $x \in c(A)$ we have $c(\{x, y\}) \neq \varnothing$, then $c(A \cup \{y\}) \neq \varnothing$.

Condition 3. Property κ: For any $A \in \mathcal{P}$, if for some distinct x, y we have $x \in c(A)$ and $y \in c(\{x, y\})$, then $y \in c(A \cup \{y\})$.

Condition 4. Property γ: For all $A, B \in \mathcal{P}$, $c(A) \cap c(B) \subseteq c(A \cup B)$.

Condition 5. (M): $M(A, R) \subseteq c(A)$

Condition 6. (O): $B(A, R) \subseteq c(A)$

Condition 7. (O*): $B(A, R*) \subseteq c(A)$

Recall, now, that the first proposition presented above is as follows:[47]

Proposition 1. For any choice situation $A \in \mathcal{P}$, *ERC*, $\mathcal{E}$, and (M) cannot jointly be true.

Proof. We show that if *ERC* and (M) are true, $c(A) \neq \varnothing$ for any $A \in \mathcal{P}$ so that $\mathcal{E}$ is false. We do so by induction on the cardinality of A.
Induction base: If $|A| = 1$, one readily verifies that $M(A, R) \neq \varnothing$ so that by (M), $c(A) \neq \varnothing$.
Induction hypothesis: $c(A) \neq \varnothing$ for all A such that $|A| \leq n$.
Induction step: We have to show that $c(A) \neq \varnothing$ for all A such that $|A| = n + 1$.

Let A be such that $|A| = n+1$; let $y \in A$; and let $B = A \setminus \{y\}$. As $|B| = n$, it follows from the induction hypothesis that $c(B) \neq \varnothing$. Let $x \in B$, and observe that for any pair $\{x, y\}$, $M(\{x, y\}, R) \neq \varnothing$ so that it follows from (M) that $c(\{x, y\}) \neq \varnothing$. But then as $c(B) \neq \varnothing$ and as $c(\{x, y\}) \neq \varnothing$ for any $x \in B$, it follows from ERC that $c(B \cup \{y\}) = c(A) \neq \varnothing$, which is what we need to show. □

The second proposition presented above is as follows:

Proposition 2. If κ, (O), and γ are satisfied, then condition (O*) is satisfied.

Proof. Let κ, γ, and (O) be satisfied. Take some arbitrary A and some arbitrary $x \in A$. Suppose that $x \in B(A, R^*)$. We need to show that $x \in c(A)$. From $x \in B(A, R^*)$ it follows that, for every $z \in A$, xR^*z and hence,

[47] I would like to thank Dr. Stefan Wintein for this proof.

by definition, there is a set $A_z = \{x_0, x_1, \ldots, x_k\}$ where $k \geq 1$ such that $x_0 = x$ and $x_k = z$, and $x_0Rx_1, x_1Rx_2, \ldots, x_{k-1}Rx_k$. We distinguish, now, between two possible cases: $k = 1$ and $k > 1$.

Case 1, or when $k = 1$, we have $A_z = \{x, z\}$. Then xRz implies that $x \in B(A_z, R)$. By (O) we then get $x \in c(A_z)$.

Case 2, or when $k > 1$. Consider, now, $\{x_{k-1}, x_k\}$. Given $x_{k-1}Rx_k$, we have $x_{k-1} \in B(\{x_{k-1}, x_k\}, R)$. By (O), we have $x_{k-1} \in c(\{x_{k-1}, x_k\})$. Now consider x_{k-2}. Since $x_{k-2}Rx_{k-1}$, we have $x_{k-2} \in B(\{x_{k-2}, x_{k-1}\}, R)$, and therefore by (O) we have $x_{k-2} \in c(\{x_{k-2}, x_{k-1}\})$. We then get by κ that $x_{k-2} \in c(\{x_{k-2}, x_{k-1}, x_k\})$. Proceeding in this way we eventually arrive at $x \in c(\{x, x_1, \ldots, x_k, z\})$, or $x \in c(A_z)$.

Thus, in both cases we get $x \in c(A_z)$. Since z was chosen arbitrarily, we have for any $y \in A$ some A_y with $x \in c(A_y)$. Notice that, by the construction of the set A_y, we must have $\bigcup_{y \in A} A_y = A$.

This, in conjunction with repeated application of γ, implies that $x \in c(A)$. $\square$

AKSHATH JITENDRANATH

Paris School of Economics