

The Hardness of the Practical Might: Incommensurability and Deliberatively Hard Choices

SERGIO TENENBAUM

University of Toronto

Abstract: Incommensurability is often introduced with the small improvement argument. Options A and B are shown to be incommensurable when it is neither the case that A is preferred to (or better than) B nor that B is preferred to (or better than) A, but a slightly improved version of A (A+) is still not preferred to B. Since A+ is preferred to A, but not to B, we must also conclude that it is also true that A and B are not indifferent (or equally good). Such incommensurable options seem incompatible with orthodox decision theory (and various forms of value theory) but options that obey the pattern described by this argument seem ubiquitous: my choice between lemon tarts and eclairs at my favourite pastry shop might exhibit this pattern, but so could my choice between jobs or careers. In trying to accommodate incommensurable options within various frameworks, philosophers have argued that we must preserve certain central features of the phenomenon. Among them is the supposed “hardness” of at least some incommensurable options: even if perhaps one would need to be a rather anxious gourmet to describe the choice between lemon tarts and eclairs as hard, the choice among careers could potentially be agonizing. However, it is not clear in which way choices among incommensurable options are “hard,” nor how and whether such hardness poses problems for the various accounts of incommensurable choices. To complicate matters, the deontic verdicts for choices between incommensurable options seem to be relatively straightforward: one appealing view is that in such circumstances I am rationally permitted to choose any option that is not worse than another option. This paper aims to provide a sharper formulation of at least a version of the hardness problem, to argue that various theories of incommensurability fail to account for the hardness of some incommensurable choices, and to propose that the theory of instrumental rationality I develop in *Rational Powers in Action*, aided by a Kantian insight, promises to provide an adequate explanation of the hardness of choice among incommensurable options.

Keywords: Incommensurability, hard choices, decision theory, practical rationality.

JEL Classification: D81

AUTHOR'S NOTE: I would like to thank Chrisoula Andreou, Jennifer Nagel, Julia Nefsky, Johanna Thoma, audiences at the Chambers Conference on Decision Theory at the University of Nebraska-Lincoln, the Department of Philosophy at the University of Lisbon, and the Choice Group at the London School of Economics, as well as an anonymous referee for this journal for comments on earlier versions of this paper.

I. INTRODUCTION

The possibility of incommensurability is often introduced with the small improvement argument. Two options A and B are shown to be incommensurable when neither A is preferred to (or better than) B nor B is preferred to (or better than) A, but a slightly improved version of A (A+) is still not better than B. Since A+ is clearly better than (or preferred to) A, but by hypothesis, A+ is not better than (or preferred to) B, we must also conclude that it is also true that A is not exactly as good (or indifferent to) B.¹ Given that A and B seem not to stand in any possible preference relation, or any conventionally measurable relation, their relation seems to fit poorly with various assumptions of decision theory and various forms of value theory more generally. Yet, it seems that choices that follow the pattern described by the small improvement argument are ubiquitous: the pattern can apply to anything from the choice between lemon tart and eclairs at one's favourite pastry shop to the choice between jobs or careers. In trying to accommodate incommensurable options² to various frameworks, philosophers have argued that we must preserve certain central features of the phenomena. Among them is the supposed 'hardness' of at least some choices among incommensurable options;³ even if perhaps one would need to be a rather anxious gourmet to describe the choice between lemon tart and éclair as hard, the choice among jobs could be an agonizing one. However, it is not clear in which way choices among incommensurable options are hard and how and whether such hardness poses problems to the various accounts of incommensurable choices. To complicate matters, the deontic verdicts for a choice among incommensurable options seems to be relatively straightforward: it seems that in such circumstances I am rationally permitted to choose any option that is not worse than another option. So why is it hard to choose in these circumstances? My aim here is to, first, sharpen the hardness problem, then argue that various theories of incommensurable choices fail to account for the hardness of some incommensurable choices, and finally propose that an alternative theory of rationality, with the help of a Kantian insight, can provide such an account.

¹ I will alternate between 'preferred/indifferent' and 'better/equally good' from now on, or just use '>' and '=' notation.

² There is much variation in the literature in how 'incommensurable' and 'incomparable' are used. I will use 'incommensurable options' stipulatively to refer to any set of options such that neither $A > B$, nor $B > A$, and there is some A+ such that $A+ > A$ but not $A+ > B$. This usage follows Andreou (2024). I will assume in all choice situations among incommensurable options below that there is no C such that $C > A$ and $C > B$.

³ See Chang (2017); Elson (2022).

II. A BASIC FRAMEWORK

First let me present a sketch of the theory of instrumental rationality I put forward in Tenenbaum (2020), the Extended Theory of Rationality (ETR). The initial presentation of ETR will serve two functions: first, unlike decision theory, ETR allows for the possibility (and even the ubiquity) of incommensurable options in an intuitive manner. On the other hand, exactly because it makes the possibility of these options intuitive, the theory seems to leave unexplained how these choices could be hard. Ultimately, I will rely on ETR to provide an account of the relevant form of ‘hardness’ involved in settling among incommensurable options. But meanwhile it provides the materials to present the main issues I will be discussing.

For our purposes, it is simplest to present ETR by contrasting it with decision theory, and, more specifically by contrasting it with a version of decision theory that has a direct *agent-guiding role*.⁴ ETR is also a *fundamental* theory of rationality: a fundamental theory provides an account of the basic principles of instrumental rationality as well as the attitudes or states that these principles are supposed to govern. So, for instance, let us take the following putative basic principle of rationality:

Persistence of Intention: If t_1 is earlier than t_2 , rationality requires of N that, if N intends at t_1 that **p**, and no cancelling event occurs between t_1 and t_2 , then either N intends at t_2 that **p**, or N considers at t_2 whether **p**. (Broome, 2013, 178)

Since this is a fundamental principle, a theory that accepts it takes intentions to be what I call ‘fundamental attitudes’; that is, attitudes whose relations to each other and to an agent’s actions determine what counts as acting in an instrumentally (ir)rational way.

A theory is action guiding (in a nutshell) if its principles not only evaluate but *explain* the action of an agent, when she acts rationally. Roughly, we can say that if decision theory is a fundamental theory with an action-guiding role, then:

- (i) Preferences are the fundamental attitude of decision theory.

⁴ Decision theorists do not typically endorse this interpretation of decision theory; in fact, some theorists explicitly reject it (Joyce, 1999). But since ETR does have such a role, this will make the contrast clearer.

- (ii) The axioms of decision-theory are the fundamental principles of a theory of instrumental rationality

From (ii) and action-guidance, we can also infer that:

- (iii) When an agent acts rationally, they manifest a preference ordering that (approximately) conforms to the axioms of decision theory due to their implicit awareness of their soundness.

ETR instead takes only intentional actions extended through time (characterized as ‘pursuing the end of φ -ing’) to be fundamental attitudes⁵ and takes the only fundamental principle to be the Principle of Instrumental Reasoning (PIR). PIR prescribes, roughly, that agents take some sufficient means to all the ends they pursue. When I am typing the sentence ‘I am typing this sentence’ because it is a means to my end of providing an example, I am manifesting my (instrumentally) rational powers, and the explanation of my action depends in part on my implicit awareness of PIR. A principle of coherence also follows from PIR: a rational agent does not pursue incompatible ends at the same time. After all, if they pursue incompatible ends, they would not be able to be guided by PIR for at least some of their ends. On the other hand, since there are no comparative attitudes at the basis of the theory, ETR seems rather stingy when it comes to advice to an agent who realizes that her ends are incompatible. It cannot prescribe much more than: revise your ends so they are compatible.

However, even though the theory does not have any comparative attitudes as *fundamental*, *basic* attitudes, preferences still play an important role in the theory. Let us start with the observation that our ends are (nearly always) *indeterminate*; their contents do not fully specify what counts as pursuing a certain end. In the above example, the end of avoiding too much risk does not specify precisely how much risk is too much. Is it ok to rush to cross the street when I am late for my appointment, trusting that drivers will stop in time? Or jaywalking when no car is visible? Although one might think that this kind of indeterminacy is peculiar to ends which, like this one, are specified in a rather vague manner, it is in fact omnipresent: I aim to spend quality time with my children, but in order to pursue this end, I need to determine further what counts as quality time, or how much time

⁵ I defend the view that intentional actions should count as attitudes in Tenenbaum (2020).

is enough time; I pursue the end of writing a book, but the representation of the end does not fully specify how good or how long the book should be, let alone the precise content of the book. But exactly because the content of our ends is not fully determinate, we can ask not only whether an end is realized, but also distinguish between *better* and *worse* realizations of our ends. Although more needs to be said about why our ends have this structure,⁶ the structure provides for a limited basis of comparisons and thus an end-relative preference ranking. Let us take, for instance, my end of cooking a meal for my cousin. There will be cases in which I clearly realized the end: namely, there is a nice meal for them at dinner time. There are cases in which I fail: most obviously if I have no meal ready when they come, but also if what I make is barely edible or not enough food to sate them, etc. But there are also better and worse realizations of my end. Perhaps had I added truffles to the pasta, it would have been a nicer meal. Relative to this end, a meal that includes truffles is preferable. However, it is worth noting that there are multiple preference orderings of this kind, and they might conflict. Relative to my end of saving money, the pasta with truffles will rank lower than the plainer pasta. Since the theory does not determine how to ‘weigh’ these rankings, it will not provide anything like a complete preference ordering. Still, in some contexts, it might be able to approximate a preference ranking that satisfies the axioms of decision theory. For instance, if I am making decisions that pertain solely to the pursuit of my end of making money, my preference ordering relative to this end might be one that obeys the classical axioms of decision theory, or at least some less demanding set of axioms such as the one proposed by Buchak (2014).

Additionally, I might also have what I call ‘reflective preferences’, preferences that give priority to some end or another. Take, for instance, two ends of mine: spending time with my children and following my football team. Although I spend considerable amounts of time pursuing each of them, I also take the former to be more important, and thus I also pursue the higher-order end of giving priority to my end of spending time with the children. How exactly these reflective preferences play out is a complicated matter, but the details are not relevant to our purposes. Instead, I will just briefly introduce a final form of preference that this kind of ‘end-first/action-first’ theory can accommodate: Pareto preferences. To put together some of the ends I mentioned here, suppose I am going to watch my football team play tonight, either by myself or with my child, and I have two options:

⁶ One possible explanation relies on ‘the guise of the good’: since agents necessarily pursue what they conceive as good, the structure is inherited from the putative value being pursued. But other explanations are compatible with the general framework.

I can either take advantage of a promotion to get more youth interested in football that offers prime seats to the two of us for \$30, or go to the game by myself and buy a regular priced ticket in the distant stands for \$50. The first option is ranked higher by each of my ends: saving money, spending time with my children, and following my football team. In such a case, I have a *Pareto Preference* for the first option over the second option.

This brief sketch is enough to provide an account of the source of incommensurable options and the rationality of various actions within ETR's framework.

III. ETR AND INCOMMENSURABILITY

If we accept ETR, it should come as no surprise that the phenomena captured by the small improvement argument is ubiquitous. On this view, incommensurable options would simply be the product of my pursuing two different ends that upon further determination turn out not to be jointly realizable or such that their joint realization requires a worse realization of at least one of the ends. Since our different ends compete for opportunities in a world of limited resources, it is no surprise that it is extremely common that we make choices among incommensurable options understood in this manner. Let us look at a couple of examples.

Further Determination of Ends Requires Abandoning at Least One of Them: My plan for tonight is dinner and a movie. I have not given much thought yet about what I will eat and which movies I could watch. I start by looking at the movies. There are two acceptable options: *Novecento* or *An Elephant Sitting Still*. Unfortunately, both are about four hours long and they are shown from 6PM to 10PM. All the reasonable restaurants in town close at 10PM, and none of them serve dinner before 6PM. Thus I must give up one of these ends, but according to ETR, *instrumental* (more on this qualification momentarily) rationality does not determine which end I must give up. Even though there are preference orderings internal to each end, they would not help in such a case. Perhaps I conclude that the violence in *Novecento* will upset me and thus it will be a better experience to watch *An Elephant Sitting Still*. This might be a significant improvement insofar as the end of watching a movie is concerned, but it still says nothing about the comparison with going to a restaurant since there are *no relevant comparative attitudes* that determine a preference

relation here. And this is obviously even more so if the difference is small; a slightly better print of *An Elephant Sitting Still* might require that I watch the better print if I watch a movie at all, but it will not determine that I must prefer watching the movie over going out for dinner.

Further Determination of Ends Does not Allow for the Best Realization of All the Relevant Ends: Let us modify the situation above slightly. I can either watch *An Elephant Sitting Still* or *Guardians of the Galaxy 3*, which will also start at 6 but end earlier and would thus allow me to go to one of my favourite restaurants. I do learn however that there is a mediocre (but acceptable) diner that is open till midnight, so I could go for dinner after *An Elephant Sitting Still*. On the other hand, I do not know if I will ever get another convenient chance to watch *An Elephant Sitting Still* on the big screen, while *Guardians of the Galaxy 3* will probably be showing for months. It's a better realization of my end of going out for dinner that I watch the latter movie but a better realization of my end of watching movies to watch the former. And similarly here, unless I have a reflective preference, instrumental rationality as conceived by ETR, permits both actions and minor improvements on either option will make no difference to this deontic verdict.

ETR is a theory of *instrumental* rationality. Arguably, when I deliberate about incommensurable options, I am not just concerned with pursuing my ends well, but also with pursuing the right ends. So perhaps substantive normative reasons will constrain my choices in such cases in a way that is significantly more demanding than the deontic verdicts grounded on purely instrumental requirements; perhaps, it is substantively rational for me to give priority to the end of watching a movie over the end of dining, and this would settle the issue in favour of the second option in our example. However this is unlikely; the cases above did not seem to propose any action that would be substantively irrational and a theory that prohibited, say, watching *Guardians of the Galaxy 3*, would be thereby implausible. In fact, we could represent substantive constraints as prohibitions against engaging in certain actions or taking certain means to our ends or obligations to adopt certain ends (or engage in certain intentional actions). This would leave the general structure of the theory intact, and we have no reason to think that such prohibitions and requirements would eliminate the sources

of incommensurability mentioned above. Indeed, anyone who accepts Raz's popular Basic Belief according to which, in many situations, there will be "more than one response supported by reasons, with none of them supported to a higher degree than any of the others" (Raz, 1999, 111), is already committed to the view that the introduction of substantive normative reasons will not narrow down significantly the field of permissible ends that an agent might choose to pursue.

Moreover insofar as the ends that the theory of substantive rationality (or normative reasons) requires us to adopt are unlikely to be very robustly determinate, the concomitant pursuit of our ends will have to leave room for generating incommensurable options in trying to realize them jointly. Let us take, for instance, a proposal by Robert Noggle (Noggle, 2009), according to which, roughly, the duty of beneficence is a duty to have the well-being of others as one of our ultimate ends. An ultimate end is an end that has great importance in our lives. But, on this picture, benevolence is not our only ultimate end, and giving it priority does not mean always choosing it over another end. The end of benevolence is an end pursued through a rather extended period, arguably one's whole adult life, and various manifestations of it are acceptable realizations of this end. Moreover, even though you may not abandon an obligatory end, you could abandon various particular specifications of this end. Suppose I've been helping with our local soup kitchen every Wednesday, but I've joined a football league and they have just scheduled their games for Wednesday evenings. My commitment to the end of beneficence does not require that I realize it by working at this soup kitchen; it seems also permissible, for instance, that I continue with the football league and contribute to another soup kitchen or to another cause. In sum, the obligatoriness of the end does not rule out that we might have incommensurable options when deciding whether and how to pursue it in a particular occasion.

ETR turns out to be highly permissive with respect to incommensurable options, and this seems to speak in favour of it. Incommensurable options seem to be genuine cases in which vast permissions apply; this is exactly what lends plausibility to Raz's basic belief.⁷ But for this very reason, the theory seems to leave no room for the choices being *hard*; this very same permissiveness seems to tell us that choosing among incommensurable options would be rather trivial: any choice is permissible and there is nothing more to be said here. Since there are no basic comparative attitudes to draw on, as long as I am aware of all my pertinent ends in choosing between the

⁷But see (Portmore, 2019) for a more skeptical attitude towards this way of understanding the Basic Belief.

two careers, I can be rather confident that both my options are rationally permissible. It seems thus that an account of incommensurability based on ETR is very poorly positioned to explain why choice among incommensurable options is hard.

However, the difficulties faced by this kind of view are the direct result of accepting Raz's Basic Belief and will ultimately haunt any plausible account of incommensurability, or so I will argue below. I will then propose a constraint on rational agency (the Kantian condition), and argue that this constraint, together with the ETR-based account outlined above, provide a compelling explanation of why choices among incommensurable options are hard even when you know the relevant deontic verdicts of each option. But let us first try to understand better the different ways in which such choices might be hard.

IV. HOW HARD IS THE PRACTICAL MIGHT?

In what way can a decision be hard? Suppose my doctor tells me that I need an urgent operation (in the next week or so). The operation will result in the loss of my left arm, but it will save my life. She is ready to perform the operation at any moment; I just need to give her the go ahead, and she will get going. Or suppose I can save a child from an accident but I need to drive over a sympathetic racoon and kill it. These decisions are obviously hard in *some sense*. I need to say 'yes' to killing a racoon or to letting go of my arm, and these things seem to call for reactions of regret or sadness. If, say, in order to save my young children from the incoming Nazis I must leave behind my dying grandmother to suffer her last days alone, the choice might be clear and yet it is certainly devastating to choose to abandon one's dying relative. These choices are *affectively hard*. Affective hardness is not the result of a difficulty in ranking the options; it is clear in each of these cases what my choices should be. It is also not, or at least not clearly, the result of different kinds of values. Are the values involved in wanting to save my children and to provide a good life for my grandmother different kinds of value? Or when I have to decide to whom I will offer my spare room between two friends who have both fallen on hard times and are now homeless?⁸

Suppose that, in our first example above, my doctor has no adequate anesthesia. I need not withhold my consent as she operates on me in a rather painful way. My deciding to undertake the operation is, what I will call, '*volitionally hard*'; it requires great willpower. Some decisions can also be volitionally hard without being affectively hard. Pulling off the very stuck

⁸ See Hurka (1996) for a related point.

band-aid, drawing one's own blood, or jumping into the shower when there is no hot water can be like that. One might argue that all decisions that are affectively hard are also volitionally hard, but I do not think this is necessarily so. My decision to go ahead with the surgery when proper anesthesia was available or even to run over the raccoon seems to me *not* to be volitionally hard. Arguably, a distinctive mark of volitionally hard decisions is susceptibility to *akrasia*; such susceptibility does not seem present in these cases. At any rate, it does not matter for our purposes if the affectively hard decisions are a subset of the volitionally hard, as I will be mostly interested in the distinction between the affectively hard and, what I will call, 'the *deliberatively* hard'.

Many of the affectively (and volitionally) hard decisions above are *deliberatively* easy; they do not require much reflection and deliberation to arrive at the right choice. Devastating as it is that I have to run over a raccoon, it seems that little or no deliberation is needed to know what I ought to do in this case. However, some decisions can be deliberatively hard but affectively easy, or at least relatively, so. If the tax laws of my country are complex enough, the decision of whether I should file jointly with my spouse can be extremely challenging. And the difficulties might be not purely epistemic, at least not in a narrow sense of 'epistemic': I do not need just to gather the information, but know when to stop gathering it, or how much time is worth spending on this issue. But frustrating as it might be not knowing what to do, choosing one way to go rather than another is affectively easy, especially if I'm reasonably confident that I chose well. Similarly, choosing your NCAA bracket for the office pool is potentially deliberatively very complex, but at least for many of us, not particularly affectively challenging. Some cases might be affectively hard exactly because we do not know what the best action is, but even such cases might be deliberatively easy. Suppose I have two patients, one that will be saved by a drug and other for whom the drug will have no effect, but I do not know which is which: our computers have been hacked and I cannot access the relevant information. *Ceteris paribus*, it seems rather clear what I should do: I need to pick one patient, give them the drug, and hope I picked the right one. But it will certainly be an affectively hard choice (and it might be particularly hard *ex post* if I find out I picked the wrong one).

Finally, some decisions seem to be affectively hard due to the structure of human agency and the nature of the normative landscape. We need no empirical psychology to know that my choice to run over the raccoon to save the baby is (or at least should be) affectively hard; if I were cavalier about killing the raccoon while it stared trustingly at me, there would be

something wrong with me. And certainly more so if I experience no agent-regret in leaving my beloved grandmother behind. But some choices might be affectively hard due to unrelated features of our psychology. I am ashamed to confess that I find it extremely difficult to throw away unpaired socks even when they're cluttering the sock drawer. Among the affectively hard choices we could distinguish between the 'structurally hard' and the 'psychologically hard'. I think a similar distinction can be made within deliberately hard choices. The deliberately hard choices we considered above are structurally hard: they present intrinsically difficult questions (at the very least they are intrinsically difficult given some basic facts about human nature). The choice about what to have for dinner might be deliberately hard for me due to quirks about my psychology: for some reason, I cannot settle on a menu item. Some others might be hard due to specific failures of my mental capacities: because I have no sense of direction and cannot read maps, deciding which road to take might require some amazing inferential feats on my part.

Since it does not befit the armchair philosopher to discuss the psychologically hard, I will focus only on the structurally hard. I think there is little doubt that incommensurable options are often affectively hard. If I have to choose between visiting my ailing aunt or help my brother take care of his newborn baby, I will have serious regrets no matter what I choose.⁹ Some radical forms of normative hedonism might have difficulties explaining this fact, but for most theories of rationality the affectively hard does not present any special challenge.

But choice among incommensurable options are often deliberately hard and many, if not all, theories have difficulties explaining this fact. For how can it be deliberately hard to choose among incommensurable options if we already know the deontic verdicts for all the options? After all, if we know that A and B are incommensurable options, and that no other option is superior to either of them, then we already know that both A and B are permissible options. Philosophers have argued that on certain views of incommensurability, deliberation would be a waste of resources (Chang, 2017). But I think the real difficulty is that extant theories of incommensurability cannot explain why one is often *required* to deliberate among incommensurable options, or, at least, when it would not be appropriate not to deliberate; arguably, any theory can accommodate a *permission* to deliberate among incommensurable options.

⁹ For a more detailed discussion of similar cases, as well as a compelling argument that this kind of hardness is not distinctive to choice among incommensurable options, see Andreou (2024).

One could try to ground this permission by appealing to the fact that I can never be *certain* that I am in a situation in which I am choosing among incommensurable options. Suppose Mary is contemplating a choice between taking early retirement or working for a few more years. On the one hand, retiring now will ensure that she can stop going to Faculty meetings at this very moment, get a more flexible schedule, be able to spend more time with her grandchild who is about to be born, etc. On the other hand, postponing retirement will allow her to continue supervising students, have more money in her old age that she can use to take trips around the world, etc. Mary may conclude that these options are *probably* incommensurable, but how could she be *certain* that they are? So, one might argue that it is permissible to deliberate among incommensurable options because we are rarely, if ever, *certain* that our options are in fact incommensurable.

However, Chang correctly points out that the lack of certainty is not enough to make further deliberation appropriate. For, as Chang (2017) argues, in many cases we will have *practical certainty* that our options are, as Chang (2002) puts it, on a par. If I do not expect that further deliberation will improve my epistemic situation significantly, or at least not improve in such a way that I cannot expect that the benefits of further deliberating will outweigh the costs of engaging in it, then I have *practical certainty*. But it seems that in many cases of incommensurability we can rapidly achieve such practical certainty, and if even in these cases my choices are deliberatively hard, then the explanation of this fact must be sought elsewhere.

Chang's claim that practical certainty suffices to ground a good decision seems undeniable. But this would not be enough to show, on any model of incommensurability, that deliberation in cases of important decisions is always, or even often, inappropriate once I am practically certain that the options are incommensurable. It is true that sometimes deliberation can be very costly and excessive deliberation might be psychologically taxing. But deliberation is often nearly cost free; I can deliberate about my momentous life choices while I am running, in my commute to work, and so forth. It is not like in such situations, I have clearly in my mind an appetizing menu for better use of my intellectual abilities. Moreover, deliberation about a topic often comes to us unbidden. I might want to focus on writing my paper, but the question of what to have for dinner keeps coming to mind. The more important the decision is, the more likely that the low bar for non-wasteful deliberation in non-exigent circumstances have been cleared.¹⁰

¹⁰ See Elson (2022) for similar and further arguments for the permissibility of deliberation after one reaches practical certainty.

However, incommensurable options not only permit extended deliberation, but actually seem to *require* it. It would be exceedingly strange, and seemingly inappropriate or irrational, if someone decided among careers without deliberating the merits of each, by simply forming the rather plausible judgment that, say, a career in bricklaying and a career in carpentry are incommensurable. Let us go back to Mary, our academic friend who had to make a decision about whether to retire. Suppose Mary is a friendly colleague, very much appreciated for her judgment and deep knowledge of various topics. She has been working at Massive University for many years now, and many of her colleagues have relied on her advice on making decisions about early retirement. After hearing detailed descriptions of their situation from each of these colleagues, she invariably comes to the same conclusion: these options are, to use Chang's language, on a par. She helps each colleague figure out the financial consequences of each decision, gives them ideas for retirement projects that would be a good fit for their abilities and inclinations, but she always ends up refraining from recommending either option, as she always finds them on a par. It is now her turn to make a decision. She does not think she is significantly different from any of her colleagues or in any way unique. Thus she is practically certain that she is facing a choice among options that are on a par, or, as we have been describing, a choice between incommensurable options.

It seems clear that it would be appropriate for Mary to engage in deliberation here. But this is a vast understatement. It would be positively bizarre, to say the least, if she decided that deliberation here was unnecessary; that is, if she decided, given that her inductive basis makes it practically certain that her choices are incommensurable, just to go for one of these options without any deliberation. Or suppose that Larry comes for advice and Mary immediately tells him, 'you know I've advised so many people by now, I can assure you your options are on a par', it would be especially inappropriate for *Larry* to make a decision solely on these grounds without any further deliberation. In sum, we need to account not only for a *permission* to engage in deliberation, but a *requirement* to do so. At the very least, in many circumstances, it seems a mistake, or inappropriate, to decide among certain incommensurable options without deliberation, even when I already know, or am already practically certain, that the options are incommensurable.

Let me make a further observation. Suppose I need to decide between two investment portfolios, and after deliberating for a long time, I am now fairly confident (and even practically certain) that one of them is somewhat better. But since I am not *certain*, I could deliberate even more. However,

we should expect that at this point, deliberation would take a somewhat different turn. Often this kind of deliberation will take a significantly larger focus on second-order considerations; if, after careful analysis of the evidence at my disposal, I am pretty confident that it's better to invest in US bonds than in bitcoin, if I continue to deliberate I'm likely to focus more on calculation mistakes I might have made, on whether my sources of information were indeed reliable, and so forth. Of course, this does not mean that first-order considerations would be left aside, but we'd expect that the more we are practically certain, the more second-order considerations become salient in challenging our certainty. Perhaps you disagree about this structure of deliberation in the face of ever increasing (practical) certainty. But what really matters for my argument in what follows is that this is certainly *not* the case in Mary's situation or for those who continue to deliberate in circumstances in which they can be confident that the choices are on a par. Deliberation in such situations seems *not* to take a different turn; we seem to focus on our reasons for each option in the same way. And I will argue next that extant theories of incommensurability have a difficult time explaining why Mary's choice is deliberatively hard in this way.

V. EXTANT THEORIES AND EXPLAINING HARDNESS

Here I will outline my reasons for dissatisfaction with extant theories of parity or incommensurability in explaining these concerns. First, I briefly examine the problems faced by the proposal that cases of incommensurability are cases of 'rough equality'¹¹ and the more specific idea that incommensurability is explained by the vagueness of the comparative evaluations (Broome, 1997, 2022; Elson, 2017). Then, I will examine in a bit more detail Chang's radically different proposal that combines her view about the existence of a fourth comparative value relation (parity) and limited voluntarism (Chang, 2009, 2017).

Let us start with rough equality. According to this account, in some choice situations, the values of some options are roughly equal (and there are no options that have values that are significantly higher than these options). In such a case, small improvements to one of them may still leave the options roughly equal, and this explains why we get a choice situation with the structure brought forth by the small improvements argument. It is unclear that this proposal can explain all cases of incommensurability, in part because it seems that often a 'large improvements argument' also

¹¹ Parfit (1984, 430–431) proposes an early version of this view.

applies to some central cases.¹² But leaving aside these concerns, we must conclude that, on this view, Mary already knows that her retirement options are roughly equal. So what would require her to deliberate further? Why would it be inappropriate for Mary not to deliberate at all? Under this account of incommensurability, she is by now quite confident that the options are roughly equal and thus that further deliberation would not improve her situation significantly. Would it not then be perfectly appropriate for Mary not to deliberate about her options?

According to the more specific proposal that incommensurability is vagueness, the structure of incommensurable options obtains when the relevant comparative evaluations are vague. In a nutshell, we can say that just as an extra hair will not turn a borderline bald person into a definitely non-bald one, a small improvement will not turn a borderline comparative evaluation into a definitely true one. Suppose Mary were confident that it was vague which retirement option would be better for her; that is, to use a supervaluational model, she knew that on some precisifications of 'better than' retiring now is better while on other precisifications, retiring later is better. This seems to be, if anything, a reason *not* to deliberate (but see Elson (2022)); after all, why would deliberation favour some precisifications over others? Moreover, at least on some views of vagueness, the content of deliberation would also change here; the question would switch, for instance, to questions about the linguistic practice of English speakers.¹³ Of course, on any proposal that incommensurability is some form of rough comparability, the most promising way to make room for deliberation among incommensurable options is to argue that we are often uncertain whether some options are really roughly equal. But this sends us back to the proposals we examined earlier.

Chang has famously defended a view in which there is a fourth value relation other than the classic $>$, $<$, and $=$. Seemingly incommensurable options are often cases in which neither option is better than the other, nor they are equal; they are instead, on this view, on a par. On its own, this view does not seem to explain why a choice among incommensurable options is deliberatively hard (though it might explain other ways in which it is hard). After all, it seems that if A and B are on a par, the deontic verdicts are still rather clear: again, they are both permitted. And even if it might not be fitting for some reason to flip a coin in some of these cases, it is not clear why deliberation is called for: we already know that either choice is

¹² See Andreou (2021, 2022) for this claim.

¹³ Schoenfield (2015) argues that all *moral* vagueness is ontic vagueness on the basis of a parallel concern.

permitted. However, more recently Chang has defended a 'two-stage' view of the source of reasons. Very roughly, according to Chang, evaluative facts rule out certain options. For instance, evaluative facts rule out the career of a contract killer as a possible career for any agent, and a career in neurosurgery or carpentry for someone with my level of manual dexterity. These reasons are *given* by the evaluative facts. However the given reasons still leave me with various options that are on a par. For those options that are on a par (and superior to all others), say a career in tax accounting or in landscaping, I must (or at least may) *create* the reasons to choose one or the other; although I have given reasons to make *some* choice, I must create my 'rational identity' as a landscaper or a tax accountant. Thus, on Chang's view, our deliberative job is not over when we realize that our options are on a par. I still need to (or at least may) choose a rational identity; I still need to (may) commit to landscaping or tax accounting by freely creating reasons to act one way or the other. So choice among incommensurable options can be hard because the agent's deliberative work does not end when the given reasons run out; we still need to create the reasons we need to act on.

I cannot do full justice here to this interesting approach. But I do want to raise some doubts about whether this can capture the fact that choice among incommensurable options is deliberatively hard. First, insofar as Chang is right that we deliberate about our rational identity, this deliberation cannot be by means of (given) reasons. Suppose that there were given reasons for me that would settle that I must choose the rational identity of being a tax accountant over the rational identity of a landscaper. In such a scenario, it would certainly be better for me to be a tax accountant (given the supposition that I have reasons to choose a rational identity), but this contradicts the hypothesis that the two options were on a par. And of course it is no help to appeal here to created reasons; this would just push the problem one step back. So if there is deliberation at all involved in creating reasons, it will be radically different from the deliberation in which we evaluate the merits of different options or that we examine whether our reasons for or against an action are valid, defeated, outweighed, and so forth. But this already violates one of our desiderata: deliberation among incommensurable options seems to be focused on the reasons for and against the various options and seems no different than deliberation among options in which you are still trying to settle which one is the better option. More generally, it seems far from clear what the content of the second stage of deliberation involves and why should it count as deliberation. On a certain conception of deliberation, we weigh the reasons for and against our

options; on another, we deliberate about whether potential reasons to choose an option are defeated by the reasons to choose some other option. But these forms of deliberation are ruled out since, *ex hypothesi*, we know that the given reasons are on a par. So how could we possibly deliberate about which rational identity we should create?¹⁴

Again, I do not mean these remarks to be conclusive but just to provide us with some incentive to look at other views about why we deliberate about incommensurable outcomes. In fact, as we will see I think there is some truth in both the uncertainty view and this proposed ‘existentialist’ view. But more on that below.

VI. THE KANTIAN INSIGHT

As we learn in the early days of our philosophical education, according to Kant only actions done from duty have moral worth; actions that are merely in accordance with duty, like the actions of his famous shopkeeper, do not have such merit. But Kant’s point is in fact much more general. For Kant, actions done from duty have this special place in his moral philosophy because they manifest our rational nature: they do not simply get things right by accident, but they also express our (effective) understanding of what makes them right. Rational action more generally is not just making the right choice but making the right choice for the right reasons (or via the grounds that make the choice the right choice)—I will call this condition on rational choice the ‘Kantian Condition’. We can see now that Mary’s non-deliberative choice would violate the Kantian condition. She would not be choosing to retire early because, say, this will allow her to spend more time with her family. She is deciding on the basis of second-order evidence that this is the correct choice. This is hard to model on a theory of instrumental rationality that is grounded on some form of decision theory. If what makes the right choice is the fact that it is the one that maximizes utility, then choosing something because it maximizes utility will suffice to comply with the Kantian condition. But this will not distinguish between, for instance, deciding on second-order evidence and first-order evidence, and thus it would not be enough to show that there is anything wrong with Mary’s choice. Moreover, those who accept that decision theory is not supposed to be an ‘action-guiding’ theory might allow that agents make decisions on the basis of heuristics whose grounds may bear little resemblance to the process of figuring out what maximizes utility. Of course, I am not trying here

¹⁴ Of course, this is not meant as a general criticism of the possibility of such created reasons, but only as manifesting skepticism that appeal to such reasons can explain the deliberative hardness of the decisions in question.

to refute normative decision theory; even if we can adequately defend the Kantian condition, nothing prevents the advocate of such a theory from finding another way to accommodate it. Perhaps, for instance, if the heuristics rely on a strong enough desire that is part of the explanation of why the choice maximizes utility, this will suffice to fulfill the Kantian condition. My only claim is that it is not immediately clear how such a theory can accommodate this condition.

But why accept the Kantian condition at all? If you are making the right choices, why would it matter that you are also making them for the right reasons? I'm not sure I can give a non-question begging and conclusive argument for this condition—one that is independent of any theory of rationality. My main aim will be to show that the Kantian condition is very plausible if we accept our ETR based account and together they explain why choices among incommensurable options are hard. But I do think that the condition is intuitively plausible, and that some other independent considerations favour it. Imagine someone who lives their life just trusting a planner, who explains very little to them about the grounds of the planner's decision. But the planner consistently tells our agent, call them Jeri, that they have nothing to worry; they are choosing the right thing. Jeri has learned from past experience that the planner is a much better reasoner than they are. So Jeri follows along the planner's recommendation: they go wherever the planner tells them to go, they marry whoever the planner tells them to marry (and divorces them when the planner tells them to divorce them), they choose whatever career the planner tells them to choose, etc. There seems to be something seriously amiss in Jeri conducting their life in this manner. One might argue that understanding the grounds of our choices is just one more good in our lives, and that once we add this value to our equation, we can dispense with the Kantian condition. But if it is just a good in our lives like any other, why could not our advisor be so excellent that it would compensate for the absence of such a good? And if we can accommodate the fact that it does not seem possible to compensate for such an absence, we might have just embraced the Kantian condition. One can also argue that given that many of the goods of a human life involve enjoyment, Jeri might not get these goods if the planner chooses for them. However, here too, it's not clear that the counterargument works. Why could Jeri not enjoy these goods that were chosen for them by the planner, even if they are not making the decision in light of such an enjoyment?

Another important consequence of the Kantian condition provides a further argument for it, and supplies a crucial piece of our account of the hardness of choice between incommensurable options. Let us say that an

agent, insofar as they are practically rational, choose only undefeated options. But if the Kantian condition holds, whether an option is defeated or not depends on my reasons for choosing it. Suppose I am at a restaurant and the menu has three options. Since the only relevant considerations for my decision are (i) how good each dish tastes, and, (ii) how healthy it is, we can represent my choice situation as follows:

	Tasty	Healthy
Asparagus soup	Yes	Yes
Raw broccoli	No	Very
Cheese souffle	Very	No

Table 1. Dishes available and relevant considerations.

Is the choice of asparagus soup permissible? If the Kantian condition holds, we cannot answer this question in an unqualified manner. There is *a way* of choosing the asparagus soup such that the choice is (arguably) undefeated: if I choose the asparagus soup *in order to eat something both healthy and tasty*. At least if we assume that this reason is defeated neither by the better taste of the soufflé nor by the superior health benefits of broccoli, the choice of asparagus soup *for this reason* is permissible. However if I choose the asparagus soup just because it is tasty or just because it is healthy, I have chosen a defeated, and thus impermissible, option, since the menu offers both a tastier and a healthier option. This aspect of the Kantian condition explains an otherwise surprising feature of choice between incommensurable options that Hare (2010) calls attention to. Suppose I have time to save from a fire either a beloved portrait of my grandmother or my son's graduation yearbook. Suppose also that these are incommensurable options for me. If I now realize that there is a \$100 bill attached to the box containing my grandmother's portrait, the extra cash does not rationally require me not to choose the yearbook. But the deontic statuses of my options seem to change if I do not know which option is improved—if I cannot remember whether the box with the \$100 bill contains the portrait or the yearbook. In such a situation, it seems that I must choose the box with \$100 (why would I choose the other box if I have no idea what it contains?), even though, had I known what was inside the other box, *irrespective of what it turned out to be*, I would be permitted to choose it. The Kantian condition provides an explanation of this difference between choosing under knowledge and under ignorance. If I know that the \$100 is attached to the

box with the portrait, I can choose the other box *in order to keep my son's yearbook*. But if I do not know this is the case, the best I could do would be to choose the cashless box *in order to have a (0.5) chance to keep my son's yearbook*. However, the other box also offers this chance of keeping the yearbook and, on top of it, it provides me with extra money. Thus, in such a case, choosing the cashless box is always a defeated option.¹⁵

VII. ETR, THE KANTIAN CONDITION, AND INCOMPARABILITY

What exactly does Mary know when it comes the time to make her own decision about whether to retire early? Mary knows that there are *some* ways of retiring early that would be justified, and *some* ways of keeping working that would also be justified. But whether this is true in her case, no matter what her decision is, depends on why Mary decides to retire. As we saw above, given the Kantian condition, not every way she chooses one option is necessarily a correct one. If we accept ETR's contention that the Principle of Instrumental Reasoning is the only fundamental principle of instrumental rationality, at the very basic level, the agent satisfies the Kantian condition only if she is (knowingly) taking sufficient means to the end she is choosing to pursue.

On this view thus, if Mary just chooses to take early retirement solely on the basis of what she has learned inductively from her past experience, she is failing to satisfy the Kantian condition. In such a case she chooses to retire early on the basis that there is some way of choosing to retire early that is correct. But choosing on this basis does not make the choice correct; if she pursues at all the end of 'choosing an option such that there is no other option that is a better option than the one I choose', it is definitely not the end for the sake of which she is retiring early.

Of course, she could go on and decide to retire early to pursue her end of becoming a fencer. But note that this is no longer what she has evidence for. That is, her evidence establishes that there is *some* way of choosing to retire early that makes it a correct decision. However, her inductive evidence almost certainly does not settle that retiring early *in order to pursue the specific end becoming a fencer* is justified. And in taking in this next step, she might fail on various grounds. First suppose Mary wants to pursue the end of developing her fencing skills so that she will become an Olympic level fencer. Sadly, however, Mary does not have the natural ability to become such an outstanding fencer (and she is in a position to easily know

¹⁵ I discuss this point further and draw other consequences from it in Tenenbaum (2024).

this fact). So even though some way of choosing to retire early would make her choice justified, the choice she made was not one of them.

The failure need not be such an obvious failure of pursuing an end for which there are clearly no sufficient means available to the agent. Suppose that when making her final decision, Mary is focusing on two ends she is pursuing: the end of spending considerable time with her family in Minnesota (her job is elsewhere) and the end of having enough money for a comfortable life. She realizes that by retiring early, she could pursue both ends. Although she would be making less money, it would be enough for her needs. On the other hand, she would be able to move to Minnesota and spend significantly more time with her family. However, it turns out, that in thoughtlessly erasing all the emails from the administration, she overlooked how generous the post-pandemic work from home and flexible schedule policies were. Had she looked more into the issue, she would have realized that she could also move to Minnesota while continuing to work and thus be spending time of her family and yet have more money. In other words, she chose against a Pareto preference,¹⁶ and thus chose the option to retire on grounds that were defeated.¹⁷

This does not imply that there was no way that Mary could correctly choose the option to retire early. Had she retired early as a means to pursue her dream of being a novelist, she might have been making a choice on undefeated grounds; even the company's new flexible policies would not leave her enough time to write a good novel. This leads us to a further, compounding source of the demand to deliberate in such situations. The ends one pursues are always to some extent indeterminate.¹⁸ Suppose Mary decides to retire early so that she has more free time to pursue other projects. At this level of generality, this choice might not be defeated by any choice that Mary could have made had she chosen not to retire. Yet, if she ends up using her free time to spend more time with her family, it will be true that she had an option available to her that would have been better (that is, for which she had a Pareto preference); namely, the option of continuing to work and move to Minnesota. By itself, this does not mean that she made a bad (let alone an irrational) decision when she chose to retire early. Deliberation must come to an end, and no matter how much effort we put in specifying our ends, a great deal of indeterminacy will remain. When Mary decides she will spend more time with her family, she could go

¹⁶ See p. 1-aa above.

¹⁷ Assuming that there are no other considerations here that would permit choosing less preferred options.

¹⁸ See Tenenbaum (2020) for a more detailed defence of this claim.

on trying to make at least preliminary plans about with whom in her family she will be spending time, what she would do if various eventualities took place, the likelihood of each of them happening, and so forth. And it will always be true that an undefeated ground for choosing an option could be defeated by other options as one specifies in more detail the actions that one chooses to pursue (although becoming a novelist is not defeated by having enough money to buy a bigger house, *this way* of becoming a novelist is). These various intermediate plans could themselves be more or less specific. Granted, at some point Mary needs to stop deliberating. But with such momentous decisions that affects the pursuit of so many of our ends and that can be pursued in so many different ways, the no-deliberation point would be too cavalier a place to stop; Mary needs to at least have some good reason to think she is unlikely to end up in a defeated option as she pursues the ends for the sake of which she chose to retire early.

In sum, since our choices among incommensurable options are made on the basis only of a subset of the grounds that could determine to choose an option, it is not enough for me to know that *some* of my grounds to decide would make my choice the correct one; I need to ensure that the ones I chose are, as far as I can reasonably determine, undefeated. So knowledge that a permissible choice is possible either way does not take away the difficulty of deliberating: for, if, on the one hand the grounds on which I actually choose an option are relatively small in number, the potential defeaters are potentially unlimited. So a question remains: do the grounds on which I make my choice really suffice to justify what I do?

On this account, deliberation among incommensurable options does not change in character when Mary knows that her choices are incommensurable; she deliberates exactly in the same she would deliberate had she been uncertain about whether both options are permissible ones. This satisfies the desideratum we put forward above. But this account also explains the open-ended nature of deliberation about incommensurable options. There is no principled reason to stop deliberation at any specific point: we could always try to specify our pursuit of the relevant ends further, reassess the different conditions that might obtain and so forth.

Of course, this structure does not *guarantee* that in every such choice among incommensurable options, not even among important ones, deliberation would be required. If the grounds for choosing one option or another are clear, and if further specification does not change much the nature of such grounds, then our account cannot generate a demand to deliberate. But this is a welcome result; in such cases, indeed there is no such demand. My choice above between the portrait and yearbook does not seem to call

for deliberation. Even leaving aside the urgency of the matter, I could simply choose one of the options, even if preserving each item is of great value to me. Of course, a *permission* to deliberate may still apply (at least in a non-urgent version of the example), and our account can explain the content of such a deliberation: even here trying to specify more precisely what the pursuit of preserving these items would entail (would I look at my grandfather's portrait from time to time, or would I only be preserving it for future generation?), might engage in different ways with ends I already pursue or would be pursuing.

Let us come back briefly to the question of what justifies the Kantian condition. I hope the discussion above already has provided most of the answer, but a further consideration is also pertinent. As we take our basic attitudes to be intentions, desires, or preferences, it seems that the work of practical reason ends as soon as we form the relevant mental state. However, in choosing a retirement option, Mary still has to implement this choice, and her implementation is in no way trivial: it will need to further determine what a retirement life looks like. If writing her novel is not part of the grounds for which she chooses an option, this will inform how she implements her choices, and thus she cannot count on the fact that what she will actually be doing is undefeated by what she would otherwise be doing. And here is our element of truth in voluntarism: our choices do create reasons in a certain way. As we pursue our ends through extended actions, we further specify them and thus change the structure of how we continue to carry them out. Perhaps a bit far from Chang, or Korsgaard,¹⁹ style voluntarism, but possibly a more familiar form of reason creation.

Let us go back to our original puzzle: how can a decision be deliberatively hard, if we already know the deontic status of all options? It turns out that 'knowing the deontic status of all options' hides an ambiguity, once we accept the Kantian condition. It is only in a qualified sense that we know the deontic status of each option; for some of them (the non-defeated options), all that we know is that there is *some* way of choosing the option that is permissible. But this knowledge does not end deliberation. The importance of the Kantian condition is obscured when we think of practical reasoning ending at a momentary mental state that precedes the unfolding of the action, but once we note that practical reason does not stop short of the action, we realize that how we choose a certain option partly determines how we will be guided in rationally executing our decisions.

¹⁹ For Korsgaard's version, see, for instance, Korsgaard (1996).

Even if one agrees with my understanding of what makes incommensurable options deliberatively hard, it might seem that it can be easily co-opted by other views. An advocate of a normative version of decision theory might think that the same ambiguity can be captured as an ambiguity between a too sparse and a proper specification of the menu of options available. On this view, in limiting our choice set to ‘retiring early’ and ‘not retiring early’ we left our options massively underspecified. Our choice set should contain more, and more specific, options, such as ‘retiring while taking care of grandchildren’, ‘retiring while painting in the afternoons’ and so forth. But this would quickly create an unmanageable menu of options, and we’d be back having to understand the choice in terms of reliable heuristics. Whether such a reliance on heuristics can explain the deliberative hardness of our choice situation depends on the details of the account; here I can only express mild skepticism that any such view could find a way to do so that is compatible with the Kantian condition. Meanwhile we can tentatively accept that insofar as we focus on momentary mental states as the locus of practical rationality we cannot explain the deliberative difficulties involved in such choices, at least when the deontic statuses of all options are clear. I shall then (provisionally) conclude that in taking the view that our rational powers are primarily manifested in extended actions we also understand better the nature and rationality of (even momentary) decisions.

REFERENCES

- Andreou, Chrisoula. 2021. “Incomparability and the Huge-Improvement Arguments.” *American Philosophical Quarterly* 58(4): 307–318.
- Andreou, Chrisoula. 2022. “Parity without Imprecise Equality.” In *Value Incommensurability*, edited by Henrik Andersson and Anders Herlitz, 71–83. New York: Routledge.
- Andreou, Chrisoula. 2024. “Incommensurability and Hardness.” *Philosophical Studies*. <https://doi.org/10.1007/s11098-024-02123-9>.
- Broome, John. 1997. “Is Incommensurability Vagueness?” In *Incommensurability, Incomparability, and Practical Reason*, edited by Ruth Chang, 60–89. Cambridge, Mass: Harvard University Press.
- Broome, John. 2013. *Rationality Through Reasoning*. West Sussex, UK: Wiley-Blackwell.
- Broome, John. 2022. “Incommensurateness Is Vagueness.” In *Value Incommensurability*, edited by Henrik Andersson and Anders Herlitz, 29–49. New York: Routledge.
- Buchak, Lara. 2014. *Risk and Rationality*. New York: Oxford University Press.
- Chang, Ruth. 2002. “The Possibility of Parity.” *Ethics* 112(4): 659–688.
- Chang, Ruth. 2009. “Voluntarist Reasons and the Sources of Normativity.” In *Reasons for Actions*, edited by David Sobel and Steven Wall, 243–271. New York: Cambridge University Press.
- Chang, Ruth. 2017. “Hard Choices.” *Journal of the American Philosophical Association* 3(1): 1–21.

- Elson, Luke. 2017. "Incommensurability as Vagueness: A Burden-Shifting Argument." *Theoria* 83(4): 341-363.
- Elson, Luke. 2022. "What Does Incommensurability Tell Us About Agency?" In *Value Incommensurability*, edited by Henrik Andersson and Anders Herlitz, 181-198. New York: Routledge.
- Hare, Caspar. 2010. "Take the Sugar." *Analysis* 70(2): 237-247.
- Hurka, Thomas. 1996. "Monism, Pluralism, and Rational Regret." *Ethics* 106(3): 555-575.
- Joyce, James. 1999. *The Foundations of Causal Decision Theory*. New York: Cambridge University Press.
- Korsgaard, Christine. 1996. *Sources of Normativity*. New York: Cambridge University Press.
- Noggle, Robert. 2009. "Give till it Hurts? Beneficence, Imperfect Duties, and a Moderate Response to the Aid Question." *Journal of Social Philosophy* 40(1): 1-16.
- Parfit, Derek. 1984. *Reasons and Persons*. Oxford: Oxford University Press.
- Portmore, Douglas W. 2019. *Opting for the Best: Oughts and Options*. Oxford: Oxford University Press.
- Raz, Joseph. 1999. *Engaging Reason: On the Theory of Value and Action*. Oxford: Oxford University Press.
- Schoenfield, Miriam. 2015. "Moral Vagueness is Ontic Vagueness." *Ethics* 2: 257-282.
- Tenenbaum, Sergio. 2020. *Rational Powers in Action: Instrumental Rationality and Extended Agency*. Oxford: Oxford University Press.
- Tenenbaum, Sergio. 2024. "Can't Kant Count: Innumerate Views on Saving the Many over Saving the Few." In *Oxford Studies in Normative Ethics*, edited by Mark Timmons, 215-234. Oxford: Oxford University Press.

Sergio Tenenbaum is Professor of Philosophy and the Director of the Centre for Ethics at the University of Toronto.

Contact e-mail: <sergio.tenenbaum@utoronto.ca>