

What does Incommensurability Tell us about Agency?ⁱ

Luke Elson (luke.elson@reading.ac.uk)

ORCID: 0000-0002-3013-8030

Abstract. Ruth Chang and Joseph Raz have both drawn far-reaching consequences for agency from the phenomenon of incommensurability. After criticising their arguments, I outline an alternative view: if incommensurability is vagueness, then there are no substantial implications for agency, except perhaps a limited form of naturalistic voluntarism if our reasons are provided by desires.

Terminology is a minefield here: one must choose between ‘incommensurate’, ‘incommensurable’, and ‘incomparable’, at least. The terminologies are incommensurate, but ‘incommensurate’ has stuck, so I use it. Two items or options A and B are *V-incommensurate* when with respect to some value V, A is not better than B, B is not better than A, and A and B are not equally good. *V-incommensurability* is the failure of all three *trichotomous* comparisons (better, worse, and equally good) with respect to V.ⁱⁱ

Examples of incommensurability almost always involve two options or items that each do better along different dimensions of the relevant standard or value. Cooking at home is healthier and cheaper, but ordering delivery is tastier and more convenient. In some classic cases—Mozart versus Michaelangelo in creative terms (Chang, 1997: 15) or Stonehenge versus St Peter’s in terms of impressiveness (Broome, 1997)—it’s not even clear how to specify the relevant dimensions in an informative way.

We could be epistemicists about incommensurability: maybe in all of these cases, one of the comparisons does apply, but we don’t know which. A version of this view concerning the Moorean good has been defended by (Regan, 1997) and deserves respect, but I set it aside to ask a conditional question: if none of the trichotomous comparisons determinately holds, then what are the implications for agency?

Very often, we don’t simply investigate which option is better with respect to some value—we also choose between options. When acting, it’s very rare that we have no choice at all (our options can include stalling for time, doing nothing, or doing something completely outrageous). Much of the time, though, the best option is so obvious that we don’t even waste any time consciously deliberating, or thinking it over. See (Arpaly, 2002) for a classic discussion of the prevalence—or otherwise—of deliberation in our actions.

I will assume that we are making a choice with respect to some value, but I will suppress mention of the value, which doesn’t matter for my purposes. The crucial point is the connection between the value and our reasons: we ought to choose the best option (if there is one), and if A is better than B, then we have stronger all-things-considered reason to choose A than B. Hence I will write interchangeably of values and reasons.ⁱⁱⁱ

When one of the trichotomous comparisons obtains, and given these stipulations, the situation is fairly clear. If A is better than every other option—if A is *strongly dominant*—then only A is permissible, and doing anything else instead would be a mistake. But sometimes, one option A is merely *weakly dominant*: at least as good as every option. This can be because every option is equally good, or because there are two or more options (including A) that are as good as each other, but better than all others. When we are choosing cans of beans of the same brand at the supermarket, there are likely to be many cans tied for best, and perhaps some that are worse (the dented ones). Buying the worse beans is impermissible—not morally so, but it is forbidden by our reasons. It ought not be done.

If the set of options is not infinite (and the betterness relation is not cyclic), then we know that there must be at least one option that is not strongly dominated: there is no option better than it. But under incommensurability, there may be no weakly-dominant option. There may instead be (at least) two options A and B, incommensurate with each other—and so not weakly dominant—but better than all other options.

It may be, for example, that cooking for yourself and ordering takeaway are each better than the other options—going hungry, foraging—but are incommensurate with each other. Or if you are engaging in a spot of post-pandemic travel, the ranking of your options may have a fast express train and a cheaper slow train incommensurate with each other, but better than all other options (walking, a fast express train in the wrong direction, ... most of our options are so silly that we don't even bother to think about them).

In such cases, two pieces of phenomenology seem to pull in opposite directions. First, either option seems permissible—certainly, neither is criticisable. If neither cooking nor ordering takeaway is better than the other, aren't both acceptable? And when quizzed about your choice by a nosy neighbour, “well cooking would have been healthier, but delivery was more convenient so I went for that” could be a perfectly reasonable justification. (Your neighbour might be suggesting that you should value your health more than you do, but that's a different point, and even more nosy.)

Second, except perhaps in the most trivial choices (such as what to have for dinner), we don't stop worrying when we reach the conclusion that two options are incommensurate. We typically continue to agonise and deliberate about a choice between them, at least in major choices. Even if we have to choose *right now*, and therefore pick arbitrarily, we might retrospectively think the choice over. If you are like most people, when househunting there are several options open to you that are incommensurate with each other but dominate all others: the smaller place closer to work, the bigger place further away, the fixer-upper, and so on. But with such a major financial commitment as a mortgage or lease, coming to think that your remaining options are incommensurate does not simply stop your deliberation. Most of us do not at that point simply roll a die. Instead we think it over. (We might roll a die in the food case, but not all of us would in the train case. How much you agonise about your travel options has a lot of individual variation; I would probably enjoy travel more were I to simply relax about such choices as which train to take.)

These two pieces of phenomenology are not contradictory. The former is about the permissibility of actions, whereas the latter is about the decision procedure—and not even about which decision procedures are permissible or appropriate, but about which

ones we in fact employ. But the two are certainly in tension: if both options are permissible, then why worry and deliberate about the choice?

Joseph Raz and Ruth Chang have each drawn quite far-reaching conclusions about the nature of human agency from incommensurability, and in particular from something like these two pieces of phenomenology. After criticising their arguments, I'll argue that if we see incommensurability as vagueness—a view defended by (Andersson, 2017), (Elson, 2017), and most famously (Broome, 1997) and **Broome in this volume**—then its implications for agency are minimal. I don't pretend to offer a rigorous defence of my alternative, but simply to explore how I think about choice under incommensurability, and to sketch a reasonably plausible picture.

1. Raz's Classical Conception

Joseph Raz distinguishes two views of human agency. The first is what he calls *rationalism*: “Paradigmatic human action is taken because, of all the options open to the agent, it was, in the agent's view, supported by the strongest reason” (Raz, 2002: 47).

Rationalism doesn't say that in paradigmatic human action, we do the action that is *actually* supported by the strongest collection of reasons. Instead, the action that the agent believes to be most supported by reason. And as we are often mistaken or uncertain, these need not be the same. But as I am setting aside epistemicism, I will focus on cases where her beliefs are accurate. What does 'paradigmatic' mean, here? I take it to refer to the central, 'normal' case, when things are working as they should. There may be exceptions, but they are either relatively isolated oddities, or involve some kind of error:

Rationalism. Normally, rational actions are those where the agent chooses what she believes to be the strongly-dominant option, because she believes it to be strongly dominant.

The rationalist picture is that there are many reasons for and against various actions, and the action one ought to do is that one supported by the strongest collection of reasons, after all the weights are in. (I'm setting aside here 'enticing reasons', which would complicate this picture, and of which I am skeptical. See (Dancy, 2006) for a defence of such reasons.) A huge part of the metaethics of the past 25 years or so consists of arguments about what these reasons are: are they just desires (or grounded in desires), or facts about impartial welfare, or particularist facts, or ...? Rationalism says that there is little distinctive role for the will, other than acting in compliance with the reasons we have.

There can be some exceptions to Rationalism's 'normally'. Cases where there are multiple weakly-dominant options—such as Buridan's Ass—are not a major challenge because those options remain, in Raz's words, “supported by the strongest reason”. Weak dominance is not a serious threat to the rationalist view.

The main threat is incommensurability, of an uncompromising form:

Hard Incomparability. A and B are hard-incomparable when each of the three trichotomous comparisons determinately fails to apply—A is determinately not better than B, B is determinately not better than A, and A

and B are determinately not equally good—and in virtue of this, *no* comparison holds between A and B.

In hard incomparability, all three trichotomous comparisons *determinately fail to apply*, and no other comparison applies either. The position of ‘determinately’ is crucial here, because the broader definition of incommensurability says only that the three trichotomous comparisons each *fail to determinately apply*. This is the difference between ‘she is not determinately tall’ (which leaves room for it to be indeterminate whether she is tall) and ‘she is determinately not tall’ (which doesn’t).

(Raz, 1986) argues that incommensurability is hard incomparability. If this is so, and if incommensurability is commonplace, then Rationalism starts to look untenable because there will often be no option that’s even weakly dominant. As hard incomparability becomes more widespread, the picture of normal action as reason narrowing down the options to a few dominant ones looks increasingly strained. Some instances can be accepted—if they are, as (Raz, 2002: 48) puts it, ‘relatively rare anomalies’—but Rationalism is a distortion if there not being any weakly-dominant option is the *more* common situation.

The more we move in this direction, the more support is lent to Raz’s second picture of agency, which he calls the *classical conception*: “paradigmatic human action is one taken because, of all the options the agent considers rationally eligible, he chooses to perform it” (Raz, 2002: 47). Here is my gloss:

Classical Conception. Normally, rational actions are those where the agent chooses from amongst actions she believes to be rationally eligible.

To say that an action X is ‘rationally eligible’ is to say that X is not strongly dominated, and so to choose X would not go against our reasons: it’s not the case that we ought not X.

On the Classical Conception, even after the strongly-dominated options have been eliminated, there typically remain a wide range of rationally-eligible options. If there are five incommensurate takeaway food options in Reading, but each of these five is better than all other restaurants, then reason will eliminate all but those five options. They are rationally eligible.

The pictures have quite different underlying moral psychologies. Rationalism implies that in normal cases the will has no role beyond executing one (strongly-dominant) or perhaps one of several (weakly-dominant) options. Raz argues that pervasive incommensurability construed as hard incomparability makes this picture a distortion, because the will must often choose from several rationally-eligible but incommensurate (and thus not even weakly-dominant) options.

He argues that incommensurability is indeed widespread, because desires are the only feasible candidate reasons (or source of reasons) that could remove all—or nearly all—incommensurabilities, but desires don’t provide reasons. Here is how I understand the master argument of (Raz, 2002):

- (1) Desires don’t provide reasons.
- (2) But desires are the only plausible candidate for commensurating values.
- (3) So reasons normally leave a number of hard-incomparable options.

(4) Which supports the Classical Conception.

I've defended the inference from (3) to (4) as plausible, but I've not engaged with the rest of the argument. Below, I'll argue that if we assume that incommensurability is vagueness—in particular, that it's often *indeterminate* which option is better—then widespread incommensurability can be reconciled with something close to Rationalism, because (3) is false.

2. Chang's Hierarchical Voluntarism

Raz draws from incommensurability the thought that the will can choose without (and after) reason. Ruth Chang draws a different lesson: agency may involve creating reasons where there were none.

Chang argues that the trichotomy doesn't exhaust the available comparisons:

Parity. Sometimes, options are trichotomously incomparable but nevertheless comparable, because a fourth comparative relation called 'parity' holds between them.

According to Parity, at least sometimes, incommensurate options are *on a par*. Though central to Chang's overall view, the parity claim is not crucial to the deliberation argument that is my topic.

I will here discuss her view as defended across a series of papers, including (Chang, 2013) and (Chang, 2017), but my main focus is the fascinating argument of (Chang, 2009). That argument for voluntarist creation rests on the angstiness and continued deliberation discussed above: coming to think that our best options are incommensurate with each other doesn't usually stop deliberation in its tracks, as coming to think that they are precisely equally good does. Importantly for Chang, we don't always think this continued deliberation under incommensurability *irrational*, as for example a waste of mental resources would be.

To explain the rationality of continued deliberation under incommensurability, Chang argues that deliberation has more functions than we thought. I join her in a broad conception of deliberation as including 'discovering, recognizing, investigating, appreciating, and engaging with the reasons there are', but she adds to this list the voluntarist *creation* of new reasons (Chang, 2009: 259). There are other voluntarists, but her voluntarism is distinctively hierarchical: there are 'given' (non-voluntary) reasons, but they sometimes leave incommensurabilities. Voluntarist reasons can only rationalise or justify actions that are not strongly dominated, considering only the given reasons.

Here is how I picture it: the given reasons mark out the fences of the playground within which we must act, but they don't tell us what to do within that playground. Whereas for Raz the will must now simply choose what to do, for Chang the faculty of reason can—through its capacity to deliberate—*create* or endorse reasons to do one thing rather than another, remaining in the playground. So the view is rationalist in a certain sense: paradigmatic action involves acting for what we take to be the strongest reasons, but it's (within limits) up to us which reasons are strongest. Voluntarist creation is not required, and instead one may simply 'drift'—(Chang, 2017: 19) writes that 'law schools are populated with drifters'.

Why be suspicious about voluntarism? Because it says that we can create reasons. Of course we *can* sometimes create reasons: I can make it so that I have reason to give you £10 where before I did not, by promising to give you £10. But this is changing the facts of the matter—I have created a promise where before there was none. Voluntarism typically involves the more radical claim that I can change or create the reason-giving force of the facts. But it's far less clear—and more objectionable to anti-voluntarists—is that I could create a reason to give you £10 without doing some action like promising. Promising is quite different to an act of 'willing [that] creates normativity by creating new reasons whose normativity derives from the very act of will' (Chang, 2009: 255).^{iv}

Chang's view has the advantage that because one can only create voluntarist reasons within the area marked out by the given reasons, there is no way to create voluntaristic reasons to commit murder or do other heinous things (presuming that they are forbidden by the given reasons). If the anti-voluntarist objection to normative powers is really a worry that such powers threaten to allow us to make all sorts of unfortunate reasons—such as to commit murder—then Hierarchical Voluntarism looks immune to it—see (Chang, 2009: 270). Murder is outside the playground. But if the worry is instead an expression of discomfort with us having the ability to create *any* reasons, outside of promising and the like, then the hierarchical view doesn't mitigate the worry, or not very much.

But my main focus is not Hierarchical Voluntarism on its own terms, but whether deliberation under incommensurability supports it.

2.1. The Deliberation Argument

Chang argues for Hierarchical Voluntarism by appeal to our second piece of phenomenology—the angstiness that attends choice under incommensurability, and continued deliberation. Here is my understanding of the argument of (Chang, 2009: 249ff):

- (1) Sometimes our reasons run out.
- (2) When they do, we often keep deliberating.
- (3) Unlike with continued deliberation under betterness or equal goodness, we (often) consider this continued deliberation under incommensurability to be rational.
- (4) The best explanation for (3) is that unlike under betterness or equal goodness, deliberation under incommensurability is not a waste of time.
- (5) The best explanation for (4) is that deliberation under incommensurability involves the voluntaristic creation of reasons.

Though it captures the logical structure of the argument, this presentation overstates the argument's intended strength, which is intended to be more suggestive than abductive.^v And Chang defends the hierarchical or hybrid voluntarist view elsewhere too, but I focus on this argument because though I am not convinced by it, it deserves investigation and is (along with Raz's Classical Conception) one of the few attempts I am aware of to seriously explore the connections between agency, deliberation, and incommensurability.

My objection will be to premise (3): I'll argue that we sometimes rationally continue deliberating under trichotomy, and the explanation for this fact is similar to why we

rationally continue deliberating under incommensurability, and doesn't involve normative creation. Deliberation is an activity with its own costs—typically, the opportunity cost of spending time on it, but also other psychological downsides—but also its own benefits, which include but go beyond a greater chance of choosing the best option.

Since part of deliberation's point is information-gathering, I won't assume when discussing the deliberation argument that the agent in question has full information about her choice. If the stakes are high enough, then deliberation can continue to be rational even when one option appears best. When I presented this paper at the conference in Stockholm, I was also planning a trip to Texas to visit family, and could only find extraordinarily expensive travel insurance. After a reasonable amount of time, I found what seemed to be the best insurance policy ('Policy X')—the cheapest one that met some quality thresholds. But I continued to deliberate: rather than simply buying that apparently-best policy, I continued to search price-comparison websites and so on, thinking about which policy would really be best.

Finally, I bought Policy X. I knew this was likely to be the outcome; in the past, I have not usually been able to beat price-comparison websites. Does this mean that my continued deliberation was irrational? I'm not sure. Certainly there was an opportunity cost to it. But deliberation served several functions. I *looked* for a cheaper policy, spent time wondering how comprehensive my insurance really needed to be (what my quality thresholds should be), and so on. Assuming that my continued hemming and hawing didn't simply distort my preferences or lead me into irrationality (a non-trivial assumption), the continued deliberation meant that I engaged more with the reasons for and against various policies, and searched for policies I had not already seen, increasing the likelihood that I would find the best policy. I believed that Policy X was the best policy, but I was not certain, and the deliberation raised my credence in that proposition.

I also became surer that I was not wasting money, that the policy I settled on was in fact best, and that I wasn't making a costly mistake. This reassurance is worth something, and may rationalise some continued deliberation. So perhaps I would add 'confirming' to Chang's list of the functions of deliberation (discovering, recognizing, investigating, appreciating, and engaging with the reasons there are, and—perhaps—creating new reasons).

Sometimes the costs of making the wrong choice are so high as to swamp the cost of continued deliberation. Neither costs nor benefits need be financial: some of us find comparing insurance policies oddly fascinating, whilst also deploring its necessity as a feature of gotcha capitalism. All told, sometimes it can make sense to stay and re-check one's calculation: there is a definite cost to this, not least opportunity cost, but also the prospect of a gain. Continued deliberation can bring net benefits, either by making it more likely that we choose the correct option, or because deliberation itself brings non-causal benefits.

On the other hand, sometimes even when it's clear that continued deliberation *would* produce a better choice, deliberation is not worth it and thus not rational. I buy my fruit and vegetables from the local Aldi. Aldi is a discounter, so—inevitably—there are some bruised or otherwise damaged bananas in the large tub that one must reach down into. I *could* achieve a better banana-outcome by deliberating for longer about which banana

to take. But Aldi has decent quality controls, so I know that the truly undesirable bananas are rare, and such deliberation takes time. So it would be (even given my relative frugality) irrational for me to deliberate about each banana, not to mention about every item in my weekly shop. The food outcome would be better, but the opportunity cost and mental strain of the deliberation would be large. A quick survey to eliminate any obviously-inferior bananas is the best strategy. Applied to nearly every item, this strategy will get me out of the store in a decent amount of time, with several bags of acceptable produce.

But only *nearly* every item. There are some goods in the shop that repay careful examination. The most obvious one is the avocado: a bruised or damaged avocado is near-unrecoverable, and so avocados need careful inspection. I'm labouring this point because it illustrates one thing: the difference between avocados and bananas is *not* that the stakes are higher when choosing the former (they are, but only a little). The difference is that things are more likely to go wrong with avocados, so even given roughly equal stakes, the *expected payoff* of deliberating about them is higher. Even under undoubted cases of trichotomy, sometimes deliberation is rational and sometimes not, and the stakes are not the only things which vary.

What about continued deliberation under equal goodness? It must be conceded that if we are certain that two options are precisely equally good, then it would be perverse to continue to deliberate between them. Buridan's Ass should simply plump for one of the meals.^{vi} But when it comes to complex choices, we almost never encounter cases of equal goodness. The cans of beans at the supermarket might be precisely equally good, but basically any complex choice with multiple dimensions can engender incommensurability. I challenge you to think of a multidimensional comparison where two different options—no tricks, no implausible stipulations—are precisely equally good, despite being better along different dimensions. (And even in the can of beans case, if we look closely at the shelf, we might notice they vary slightly in appearance...) All this is by way of arguing that if two options are precisely equally good then continued deliberation would nearly always be irrational, but that this is likely only to happen in cases where the options are qualitatively identical.

My argument has been basically epistemic: the main benefits of deliberation are more knowledge (or certainty) about our situation and our reasons, and the consequent ability to act better; sometimes when the stakes are high, these can justify continued deliberation even when the facts are apparently all in.

Chang, of course, is not blind to the epistemic possibility. She argues that we can be *practically certain* that we are facing a hard choice, and that the deliberation argument can be run from this point: if it is rational to continue to deliberate beyond practical certainty, then there must be something else going on:

Although we may never be in a position to know, in some strong sense of "know," that our reasons have run out in any particular case, we can, however, be *practically certain* that they have. If you are practically certain that p, it is irrational for you to act on the assumption that not p. (Chang, 2009: 250–251)

But need this state of practical certainty rule out further deliberation as irrational? Deliberation about a choice involving p is not straightforwardly acting on the assumption that not p. And even if I'm *absolutely* certain that p, deliberating about what

to do (where this depends on whether p) need not be irrational. As I've tried to argue, whether it is depends on the costs and benefits of continued deliberation, and they are not just epistemic. Consider:

Prison Escape. I'm locked in a cell and must wait an hour to make my escape attempt. I'm certain that it'll be better to bash one guard over the head with a flowerpot rather than lock him in the bathroom, and these are my only two options, but the opportunity to do either will not arise for an hour.

Is it really irrational for me to continue to deliberate? I don't think so. Of course there is the opportunity cost of not composing a prison sonnet or proof, but there are benefits too, such as confirmation that I am taking the right course, and planning for various contingencies. (Some of this might stretch the boundaries of 'deliberation', but I think they fall within the broad scope of *deciding and planning how to act.*)

Prison Escape is extreme, but the point is general: granting Chang's definition of practical certainty, practical certainty *need not* rule out deliberation as irrational. So even if we are practically certain that we are dealing with incommensurability, the apparent rationality of continued deliberation doesn't show that deliberation includes more than fact-finding or other non-voluntarist activities. And they seem rational even under trichotomy.

That is half of my response to the deliberation argument for Hierarchical Voluntarism. The other half is much shorter: if the stakes are small enough—or the expected payoff of deliberation small enough, more generally—then even under incommensurability, continued deliberation is not rational. Think again about bananas. Suppose that you are going to buy one bunch of bananas, and each bunch costs £1. One bunch looks in just slightly better condition: a little riper, a little less bruised; the other bunch is bigger. Both are perfectly usable, but they are each better along different dimensions. Assuming that this is just one item on a regular weekly shop, should you continue to deliberate amongst these incommensurate bunches of bananas? Of course not: getting a notebook and food scale and estimating some kind of quality-adjusted weight (akin to quality-adjusted life years in healthcare) for each banana would clearly be an irrational waste of your time.

That example was silly, but the overall argument is that there is no general connection between the presence of incommensurability and the rationality of continued deliberation beyond the point of well-founded belief about how your options compare.

There are several ways Chang could go here, such as restating the argument with a tweaked definition of practical certainty that p, as implying explicitly that continued deliberation about whether p is irrational. But I think this is not a plausible route. As the Prison Escape case shows, this kind of 'deliberation-forbidding' practical certainty is far rarer. Moreover, one of the assumptions of the current dialectic is that deliberation-forbidding practical certainty does *not* hold during many hard cases: the apparent rationality of continued deliberation under incommensurability is what we are trying to explain.

3. Indeterminist Rationalism

I'll now assume that incommensurability is vagueness, and that vagueness is indeterminacy: in the standard case, it's neither true nor false that (ie, indeterminate whether) A is better than B or that B is better than A. (It may be either false or neither true nor false that A and B are equally good.)

But what upshots does this view have for agency? What I'll call *Indeterminist Rationalism* says that Rationalism is true, but it's often indeterminate what we have most reason to do:

Indeterminist Rationalism (IR). Rationalism is true, but it is often indeterminate which options (strongly and weakly) dominate others.

IR is somewhere between the rationalist and classical views. For Raz (and for hard incomparability in general) there is often determinately *no* best option, because the ranking is incomplete. Under the indeterminacy view, on the other hand, it is indeterminate *which* complete ranking is the correct one: there are multiple candidate rankings. Under the broadly supervaluationist view I favour, this has the upshot that 'there is a complete ranking of the options' is true, because it's a supertruth: it is true on every candidate ranking, though the rankings disagree about what that complete ranking is. (Compare: 'there is a precise minimum number of hairs to be non-bald' is supertrue according to supervaluationism, because true on every sharpening of 'bald'; the sharpenings disagree about that number, however.)

I'll argue that IR can neatly explain the two pieces of phenomenology identified above—permissibility and rational continued deliberation. First, the permissibility intuition. Any sensible account of rational action under indeterminacy will say that when it's indeterminate whether A or B is best (and determinate that no other option is best), Aing is *acceptable* and Bing is acceptable. Why the new terminology? Because there is some dispute about whether in such cases, Aing and Bing are each permissible, or merely borderline-permissible. The most liberal kinds of rules say that both are permissible, because E-admissible (that is to say, best on one sharpening); (Rinard, 2015) has defended a view where in such cases, Aing and Bing are each indeterminately-permissible. But even on her view, all non-E-admissible options are impermissible.

Moving beyond an appeal to authority, I will say a few words about why I find a liberal account of action under indeterminacy plausible. Firstly, looking at the structure of the case, if it's indeterminate whether A or B is best, but determinate that every other option is strongly dominated, then if you A or B, then it's indeterminate whether you've done the best thing. If you do anything else, it's determinate that you did not do the best thing. Since there is no option that is determinately best, if you A or B, there is no option to which a critic can point and say 'you ought to have done that instead', and say something determinately true. Whether you A or B, you have done the best you can in your situation, because the best you can do is what is indeterminately-best (to flirt with paradox).

Second, permissive judgements about actions seem clearly right in everyday cases of indeterminacy. Consider a sorites forced march: you are walked along a row of men, starting with the very shortest, and each taller than the next by 1mm. Eventually you

are walking past some very tall men indeed. You are asked to comment accurately on whether each man is tall or not. Let's assume that you do comment correctly on the clearly not-tall and the clearly tall men. If that is not the case, then you have determinately failed to complete your task correctly. Beyond that assumption, our issue concerns the penumbra, of borderline-tall men: your answer to 'is this man tall?' should turn from No to Yes at some point in that penumbra. I think it intuitively obvious that you are not criticisable if your verdict changes somewhere in the penumbra.

We see how aspects of the Classical Conception are retained. Under a liberal decision rule, there will very often be choices not determined by reasons between options that are indeterminately-weakly dominant, and thus plausibly described as rationally eligible.

Our second piece of phenomenology is the connection between incommensurability and continued deliberation. As I argued above, there is no *necessary* connection: sometimes deliberation is rational under trichotomy, and sometimes it's irrational under incommensurability. IR simply appeals to the thought that deliberation is an activity with its own costs and benefits.

But it would be churlish to deny that there does seem to be *some* connection between incommensurability and something like continued deliberation. Indeterminist Rationalism explains this by appeal to some peculiar features of indeterminacy.

First, deliberation is rational when its expected benefit exceeds its expected cost. But under indeterminacy, it's often non-obvious or even indeterminate whether this is so: the net expected payoff of each choice will itself be indeterminate, and thanks to the phenomenon of second-order vagueness it may not even be clear *whether* we are facing an instance of incommensurability. It may also often be indeterminate how much deliberation costs (how much do you value your time, as against one of the goods at stake in the choice you face?) and so often indeterminate whether continued deliberation will bring some expected benefit.

Second, what looks like continued deliberation may simply be puzzlement. Indeterminacy is a puzzling, paradoxical phenomenon (at least, paradoxical in those cases where it engenders a sorites) and it shouldn't be surprising that we might linger when confronted with a practical manifestation of it. We can see this by confronting avowed cases of choice under vagueness: suppose, for example, that I ask you to choose the shortest tall man in this crowded room. Assuming that there is some indeterminacy about who is that man, it may seem obvious that you should simply pick a borderline-tall man. And yet I think in normal circumstances, you will 'deliberate' for some time, perhaps leading yourself on some sorites forced marches (to make this vivid, suppose that each extra mm of height costs you £1). Eventually you will likely settle on an arbitrary man in the penumbra of 'tall man': it is indeterminate whether he is the tallest short man, and it is indeterminate whether the slightly taller man next to him is the tallest short man.^{vii} Your arbitrary choice is no better than the choice you could have made when initially confronted with the problem, and yet we would not judge you irrational for spending some time on it. The slogan view of Indeterminist Rationalism is that it makes us stop and think, for much the same reason as indeterminacy everywhere can make us stop and think, and puzzle. And as I'll argue in the next section, if Humean accounts of reasons are true, then a broadly voluntaristic explanation of continued deliberation is indeed sometimes plausible.

It might be objected that I'm simply assuming that questions of action under indeterminacy don't themselves create any puzzling questions about agency. There is some truth to this, but perhaps not when it comes to Rationalism: if the indeterminacy *in our reasons* requires some theorising about agency, it will nevertheless not take us away from the core rationalist claim that rational action is acting for what we take to be the strongest set of reasons.

4. Indeterminist Humeanism and Naturalistic Voluntarism

To finish, I will explore the upshot of Indeterminist Rationalism when paired with the view that desires provide reasons:

Indeterminist Humeanism (IH). Indeterminist Rationalism is true, and all reasons are provided by desires.

There are many versions of 'Humean' or 'internalist' accounts of reasons, but they centre on the idea that somehow, an agent has a reason to do A *iff* Aing would promote the satisfaction of one of that agent's desires. The classic of the genre is (Williams, 1981), but there are many internalists of different kinds.

I won't engage in detail with arguments against desires as reasons, except to concede that they need an answer, but I will mention one of them. A common objection to Humeanism is that it would mean that deliberation about what to do is at base consideration of what we desire most, insofar as we deliberate about our reasons. This is often felt to be implausible, and in response some versions of Humeanism say that though reasons depend on desires, it is not true that the desires *are* the reasons. See (Schroeder, 2007) for an influential and sophisticated Humean response to this worry.

I do not find it implausible that all deliberation is at least partly about what we desire. For example, I was unsure whether to travel by train or by plane to Stockholm for this paper's venue. I dislike flying, both for its hassle and its environmental damage, but this rail journey would take approximately two full days, and with a child at home I decided in the end to fly. I deliberated about this for some time, and I have no difficulty construing this deliberation in terms of desires: 'do I really want to be stuck on a train for two days, burning money on food out of boredom?' seems to me a wholly accurate description of my deliberative activity. It's not *obviously* false that we always deliberate about what we desire most.

But enough about the Humeanism. What about the indeterminacy? Indeterminist Humeanism will then locate the source of—at least some—incommensurability in indeterminacy in our beliefs and desires. If my desires are imprecise, then they may be sharpened in various ways, and on some sharpenings of my desires I prefer one house and have most reason to buy that one, and on other sharpenings I prefer the other and have most reason to buy *that* one. I want both a comfortable place to read and a short commute to work, but precisely how do these desires weigh against each other? How many commuting minutes am I willing to give up for an extra square metre of writing space? And so on. IH says that the answers to these questions are indeterminate, and so choices which depend on those answers will also manifest indeterminacy, in a version of Indeterminist Rationalism.

Indeterminist Humeanism engenders a *naturalistic quasi-voluntarism*, because what I desire is at least partly within my control—and if my desires change, then my reasons change.

It is voluntarism, in the sense that if I can change my desires, I can change the normative valence of some fact, without changing the facts of the matter, unlike in the promising case. If I can give myself a desire that some car is red, then I can give myself a reason to get out the spray paint. But it is naturalistic and only *quasi-voluntaristic*, because changing my reasons relies on causal, contingent mechanisms for changing my desires—there is no direct rational ‘endorsement’ of some feature or fact that makes it a reason. To give myself a reason to eat vegan food, I may undergo hypnosis, or choose to reflect deeply and vividly on the methods of industrial farming, believing that this will strengthen my desire to avoid many kinds of animal products. But the naturalistic nature of the view means that the process may not be predictable: perhaps I will simply be inured to the brutal realities. Animal farmers are not typically vegan, though there could be many explanations of that fact.

I have limited voluntary control over how my reasons evolve. As we have just seen, our desires may evolve in unpredictable ways—though ones that can seem to be retrospectively inevitable—we may move to the countryside in search of more space, and then either develop a taste for rural life, or come by grim experience to hate commuting. These developments are not only unpredictable but also chaotic and random: it may be that you hate commuting this year because of the pandemic, and because your next-door neighbour—who takes the same train as you—is obnoxious. But had you moved to the countryside next year, after the vaccine and after said neighbour moved away, you’d grow to appreciate the peace of a commute.

The point is that you can at least partly affect what you desire, and thus what you have reason to do. But doing so will not always be permissible. In particular, ‘strategic’ deliberation, with the aim of changing your desires and thus your reasons, will often be irrational. If at time t_0 you have most reason to A, which strongly dominates all your other options, then it is impermissible to engage in a path of deliberation that will make you less likely to A. If you have most reason to quit smoking, then it’s impermissible to engage in visualisations about how horrible withdrawal will be, visualisations which make it less likely that you will quit. This is because the deliberation *dis-promotes* the outcome that you quit smoking: it makes it less likely that you comply with your reasons.

But, of course, if your deliberation is successful, then your desires may have sufficiently changed that you now at t_1 have most reason to not-A: most reason to keep smoking. Yesterday, strategic deliberation in favour of smoking was impermissible; today, strategic deliberation in favour of *quitting* is impermissible. The deliberation has been a sort of transformative experience, similar to that described by (Paul, 2014). But this shouldn’t be too much of a surprise, because according to the Humean view I’m appealing to, our reasons may often change.

Indeterminacy permits a little more strategic deliberation. If at t_0 it is indeterminate whether I have most reason to A or to B, then assuming a permissive decision rule, not only is it permissible for me to A or to B, it’s also permissible for me to go to a therapist who will strengthen my desire to A, or to her colleague who strengthens B-desires. But doing so is optional: indeterminist ‘drifting’ is often respectable. And taking the wider

view, fixing my reasons is not always permissible—the cost of therapy may exceed any benefit of having determinate reasons.

Here we see a kind of hierarchy in the voluntarism: if it's indeterminate whether I have most reason to A or to B, I may only permissibly do things which make it more likely that I A or that I B. Typically, there will be a general causal connection between 'getting more reason to X' and 'becoming more likely to X', especially given a Humean view. So typically, if I ought not C, then I ought not engage in therapy that will foreseeably strengthen my reasons to C. And there will be exceptions to even this.

It's possible to break out of the hierarchy. At t_1 , if the set of permissible options has been changed, then—even if the process that began the change at t_0 was impermissible at t_0 —the change stands. Deliberating to create these reasons was impermissible, but we are where we are, and so acting upon them at t_1 is not. If all reasons are grounded in desires, then there is no hard boundary to the playground.

Of course, not only deliberation can change our desires. The actions I take now may well determine my reasons. But since the process is naturalistic and often unpredictable, we may reject the following argument due to Raz:

there is no reason for incommensurabilities among the options open to the agents, for when push comes to shove, the need to choose will concentrate the minds of the choosers, who will realize (or think that they do) that they want one of the options more than the others. (Raz, 2002: 49)

As a claim about the phenomenology of deliberation it seems false to me. Perhaps there are biases that reconstruct desires, but I often seem to plump without forming desires. Retrospectively, I can see that it would have been quite consistent with my wants to have gone the other way—but I had to choose. In principle, this is little different to choice under the weak dominance of Buridan's Ass.

The more general point of this section has been that once we admit that desires can provide reasons—and especially if *only* desires can provide reasons—and that there can be indeterminacy therein, a kind of unpredictable naturalistic non-hierarchical voluntarism becomes an option.

5. Conclusion

Both of the authors I've considered draw conclusions about agency from incommensurability. Raz argues that agents can't create normativity, but the will can choose between incommensurate options, meaning a somewhat beefed-up role for the will, which does more than simply execute the verdicts of reason. Chang, by contrast, argues that agents can create normativity, but only within the space of rational freedom marked out by the given reasons. Here the faculty of reason and deliberation is beefed-up to cope with cases where that space is non-trivial.

I have argued, *pace* Raz and Chang, that incommensurability and related phenomena need not inspire radical views about agency. According to the two indeterminist views I've considered, incommensurability brings many quirks and oddities, but does not fundamentally affect the nature of deliberation.

6. References

- Andersson, Henrik. (2017). *How it all relates: Exploring the space of value comparisons*. Lund University.
- Arpaly, Nomy. (2002). *Unprincipled Virtue*.
<https://doi.org/10.1093/0195152042.001.0001>
- Broome, John. (1997). Is incommensurability vagueness? In Ruth Chang (Ed.), *Incommensurability, incomparability, and practical reason* (67–89). Harvard University Press.
- Chang, Ruth. (1997). Introduction. In Ruth Chang (Ed.), *Incommensurability, Incomparability, and Practical Reason*. Harvard.
- Chang, Ruth. (2009). Voluntarist reasons and the sources of normativity. In David Sobel & Steven Wall (Eds.), *Reasons for Action* (243–271). Cambridge University Press.
- Chang, Ruth. (2013). Grounding practical normativity: Going hybrid. *Philosophical Studies*, 164(1), 163–187. <https://doi.org/10.1007/s11098-013-0092-z>
- Chang, Ruth. (2017). Hard Choices. *Journal of the American Philosophical Association*, 3(1), 1–21. <https://doi.org/10.1017/apa.2017.7>
- Dancy, Jonathan. (2006). Enticing Reasons. In R Jay Wallace, Philip Pettit, Samuel Scheffler, & Michael Smith (Eds.), *Reason and Value: Themes from the Moral Philosophy of Joseph Raz* (91–118). Oxford University Press.
- Elson, Luke. (2017). Incommensurability as Vagueness: A Burden-Shifting Argument. *Theoria*, 83(4), 341–363. <https://doi.org/10.1111/theo.12129>
- Paul, L. A. (2014). *Transformative Experience*.
<https://doi.org/10.1093/acprof:oso/9780198717959.001.0001>
- Raz, Joseph. (1986). Value incommensurability: Some preliminaries. *Proceedings of the Aristotelian Society, New Series*, 86, 117–134.
- Raz, Joseph. (2002). Incommensurability and Agency. In *Engaging Reason*.
<https://doi.org/10.1093/0199248001.001.0001>
- Regan, Donald. (1997). Value, Comparability, and Choice. In *Incommensurability, Incomparability and Practical Reason*. Harvard University Press.
- Rinard, Susanna. (2015). A Decision Theory for Imprecise Probabilities. *Philosophers' Imprint*, 15(7), 1–16.
- Schroeder, Mark. (2007). *Slaves of the Passions*.
<https://doi.org/10.1093/acprof:oso/9780199299508.001.0001>
- Williams, Bernard. (1981). Internal and external reasons. In *Moral Luck* (101–113).

ⁱ For incisive written comments, I am indebted to Henrik Andersson, John Broome, Ruth Chang, Anders Herlitz, Brad Hooker, David Hull, Philip Stratton-Lake. I'm also grateful to participants at the 2019 'Incommensurability: Vagueness, Parity and other Non-Conventional Comparative Relations' conference in Stockholm, an audience at Oakland University, and participants in the Reading philosophy Graduate Class.

ⁱⁱ If the vagueness view I discuss below is correct, then 'incommensurate' may be a sensible usage, since at least some of the time incommensurability involves its being vague how options are to be compared on the same scale—how they are to be measured against each other.

ⁱⁱⁱ I'm grateful to John Broome and Brad Hooker for forcing me to be clearer here. Hooker points out that supererogation may complicate claims such as 'we ought to choose the best option'. But the arguments to follow can focus only on cases where that claim is uncomplicatedly true.

^{iv} I'm grateful to Brad Hooker and Philip Stratton-Lake for discussion of this point.

^v I'm grateful to Ruth Chang for discussion on this point.

^{vi} I'm grateful to Anders Herlitz for suggesting this connection.

^{vii} Unless we are at the edge of the penumbra, but that possibility is complicated by second-order vagueness, where 'borderline' itself has borderline-cases.