Intrinsically Desiring the Vague

Jack Spencer

Draft of 23-Feb-2022

## §1/ Introduction

This paper is about whether it is rational to intrinsically desire the vague. A proposition is *inconsequential* if neither it, nor its negation is rational to intrinsically desire. The objects of intrinsic desire are propositions, and the contradictory of propositional vagueness is propositional precision. Every vague proposition is not precise, and every precise proposition is not vague. The question to be pursued thus can be posed as follows: is every consequential proposition precise?

Philosophical orthodoxy says that every proposition is precise. So, if philosophical orthodoxy holds, then of course every consequential proposition is precise. But, as we will see, there are powerful arguments in favor of vague propositions. Whether there are vague proposition is, at this point of inquiry, an open question. Exploring the connection between consequentiality and vagueness, under the supposition that there are vague propositions is, I think and hope you will come to agree, worth the while.

Bacon (2018) is the best-developed theory of vague propositions to date, and Bacon accepts all three of the following claims:[1]

[I] Intrinsically desiring an inconsistent proposition is not rational, and, according to Bacon, the one and only inconsistent proposition is precise. But since I assume that Bacon thinks that it is rational to intrinsically desire the necessary proposition, I assume that he deems the inconsistent proposition consequential, nevertheless. Little but bookkeeping turns on this assumption, however. If Bacon thinks

**V-to-I**. Every vague proposition is inconsequential.[2]

**I-to-V**. Every inconsequential proposition is vague.[3]

**VP**. Some proposition is vague.

I will call the conjunction of these three theses, *inconsequentialism*. According to inconsequentialism, two distinctions, the distinction between the precise and the vague and the distinction between the consequential and the inconsequential, divide propositions into the same nonempty classes.

Inconsequentialists reject the orthodoxy. They affirm vague propositions. But like proponents of the orthodoxy, they maintain that every consequential proposition is precise.

Much of this paper is devoted to arguing against inconsequentialism, under the supposition that there are vague propositions. I am suspicious of I-to-V; I suspect that some precise propositions are inconsequential, if some propositions are inconsequential. But I focus most of my critical attention on V-to-I, defending the conditional consequentiality of the vague: the claim that some vague proposition is consequential, if some proposition is vague.

---

that it is not rational to intrinsically desire the necessary, then instead of I-to-V, we have: every inconsequential proposition is inconsistent or vague.

[2] See "Indifference," Bacon (2018: 38; 195).

[3] See "Richness of Utilities," Bacon (2018: 243).

The argument against V-to-I brings into focus an neglected but important distinction between things that are consequential only if precise and things that are consequential even if vague. It is that distinction that I want to under more clearly.

**§2/ The Boolean Ground Rules**

Following Bacon, I am going to assume classical logic and a Boolean approach to propositions.

The most familiar Boolean approach to propositions is the possible worlds approach, which identifies propositions with disjunctions of possible worlds (or, if you prefer thinking set-theoretically, sets of possible worlds). Exactly one possible world is actual. A proposition is true if the actual world is a disjunct of it, and false if the actual world is not a disjunct of it. So bivalence holds: every proposition is either true or false.

The assumptions below can be thought of as what remains when we subtract the possible worlds from the possible worlds approach. I assume that propositions are disjunctions of indices, not saying what indices are:

(2.1) Every proposition is a disjunction of indices, and every disjunction of indices is a proposition.

I assume that exactly one index is the actual index:

(2.2) Exactly one index is actual. A proposition is true if the actual index is a disjunct of it, and false if the actual index is not a disjunct of it.

So bivalence holds: every proposition is either true or false.

I assume that precise propositions are disjunction of precise indices:

(2.3) Every precise proposition is a disjunction of indices, and every disjunction of precise indices is a precise proposition.

And I assume that the vague supervenes on the precise:

(2.4) Every proposition is necessarily equivalent to some precise proposition.

I call the conjunction of these assumptions, the *Boolean ground rules*.

We recover the possible worlds approach to propositions, if we add the following claim to the Boolean ground rules:

(2.5) Every index is a possible world, and every possible world is an index.

But we do not assume (2.5), for these assumptions are meant to be accepted both by the proponent and the opponent of vague propositions, and (2.5) is inconsistent with vague

propositions, given the Boolean ground rules.[4] A proponent of vague propositions who accepts the Boolean ground rules must insist that there are necessarily equivalent indices, and thus must reject (2.5).

One might suspect that the Boolean ground rules tilt the playing field in favor of the orthodoxy. But, in fact, as Bacon shows: not only are vague propositions consistent with the Boolean ground rules; it is hard to resist vague propositions once the Boolean ground rules are laid down.

## §3/ Two Argument for Vague Propositions

One argument for vague propositions goes by way of rational uncertainty. It seems that one can be rationally uncertain about something, without being uncertain about anything precise.[5]

Every credence function maps every proposition to some real number on the unit interval. Every rational credence function is a probability function. So, if the Boolean ground rules hold, then no rational credence function give middling credence to any precise proposition, conditional on any precise index. In other words, if the Boolean ground rules hold, then the following claim implies vague propositions:

[4] *Proof*: If (2.4) holds, then a proposition is precise if no distinct proposition is necessarily equivalent to it. If (2.1) and (2.5) hold, then every proposition is a disjunction of possible worlds. No disjunction of possible worlds is necessarily equivalent to any distinct distinction of possible worlds, so (2.1), (2.4), and (2.5) together imply that every proposition is precise; *cf*. Bacon (2018: 42).

[5] This argument differs from, but is closely related to, the argument in Bacon (2018: 69-95).

(3.1) For some proposition $p$, some precise index $y$, and some rational credence function $C$, $0 < C(p \mid y) < 1$.

Borderline cases evince the plausibility of (3.1). Let $y$ be a precise index throughout which Harry has exactly 30,000 hairs. The following three claims together imply (3.1):

(3.2) Harry is borderline bald, if he has exactly 30,000 hairs.

(3.3) If Harry is borderline bald, if he has exactly 30,000 hairs, then it is rational to be uncertain whether Harry is bald, conditional on $y$.

(3.4) If it is rational to be uncertain whether Harry is bald, conditional on $y$, then for some proposition $p$ and some rational credence function $C$, $0 < C(p \mid y) < 1$.

These claims are not indubitable; (3.3) and (3.4) merit further scrutiny. But all three enjoy considerable plausibility, and the three together imply vague propositions.[6]

A second argument for vague propositions goes by way of rational learning. It seems that one can learn something without learning anything precise.

[6] This argument parallels Bacon (2018: 96-123).

The argument can be adapted to different conceptions of rational learning. But suppose, to keep things simple, that the only rational response to learning is conditionalization:

(3.5) It is rational to transition from credence function $C_1$ to credence function $C_2$ upon learning *p* (and nothing more) only if $C_2 = C_1(- \mid p)$.

The following three claims imply that there are vague propositions.

(3.6) It is possible to learn that Harry is bald (and nothing more)

(3.7) If it is possible to learn that Harry is bald (and nothing more), then some proposition *p* is such that: it is rational to transition from credence function $C_1$ to credence function $C_2$ upon learning that Harry is bald (and nothing more) only if $C_2 = C_1(- \mid p)$.

(3.8) If it is possible to learn that Harry is bald (and nothing more), then no precise proposition *p* is such that: it is rational to transition from credence function $C_1$ to credence function $C_2$ upon learning that Harry is bald (and nothing more) only if $C_2 = C_1(- \mid p)$.

Any uniform substitution for 'Harry is bald (and nothing more)' that renders all three premises will do. For example, if one thinks that one always learns what one learns

when one learns rationally, then we could focus instead on the possibility of learning that one learns that Harry is bald (and nothing more).

Rejecting (3.7) is uncomfortable, if one accepts (3.5).

And (3.8) is true. Let $y_1$ be a precise index throughout which Harry has exactly one hair. Let $y_2$ be a precise index throughout which Harry has 30,000 hairs. It stands to reason that it is rational to partially shift one's credence toward $y_1$ upon learning that Harry is bald (and nothing more). That is to say, for some credence functions $C_1$ and $C_2$: $C_1(y_1)$, $C_1(y_2)$, $C_2(y_1)$, and $C_2(y_2)$ are positive; it is rational to transition from $C_1$ to $C_2$ upon learning that Harry is bald (and nothing more); and $C_1(y_1)/C_1(y_2) < C_2(y_1)/C_2(y_2)$. But no precise proposition is indecisive evidence for $y_1$ over $y_2$, relative to any rational credence function. If $p$ is precise, $C_1$ is a rational credence function, $C_2 = C_1(- \mid p)$, and $C_2(y_2)$ is positive, then $C_1(y_1)/C_1(y_2) = C_2(y_1)/C_2(y_2)$.

These argument underscore the strength of case for vague propositions. They also provide a sense of the role that vague proposition in rational psychology on the doxastic side of things, if there are vague propositions. Vague propositions are—oversimplifying slightly—the things that are potentially not scrutable from the precise: the propositions to which it is rational to give middling credence, conditional on a precise index.[7]

That prompts a question: what role do vague proposition play in rational psychology on the bouletic side of things, if there are vague propositions?

[7] See "The Principle of Plenitude" (Bacon 2018: 118), as one way, although not the only way, of fleshing this thought out.

## §4/ Intraprecision and Rational Utility Functions

According to inconsequentialism: no role at all.

The alleged inconsequentiality of the vague is easiest to appreciate if formulated using utility functions.

I will make two assumptions about the relation between rational utility functions and rational intrinsic desire. I will assume that every rational overall pattern of intrinsic desire is represented by some rational utility function, and I will assume that every rational utility function represents some rational overall pattern of intrinsic desire. The representation need not be unique in either directions.

These two assumptions allow us to formulate rational constraints on intrinsic desire as rational constraints on utility functions.

The notion that we need to cast inconsequentialism in terms of utility functions is intraprecision. A proposition is *consistent* just if at least one index is a disjunct of it. Two propositions, *p* and *q*, are *intraprecise*, just if, for some precise index *y*, both *y* & *p* and *y* & *q* are consistent. (If we think of propositions as sets of the indices that are disjuncts of them, then *p* and *q* are intraprecise just if each of them is a nonempty subset of some precise index.) Cast in terms of rational utility functions, V-to-I, I-to-V, and VP, respectively, become:

(4.1) A utility function is not rational if it maps intraprecise indices to different values.

(4.2) A utility function is rational if it does not map intraprecise indices to different values.

(4.3) Some utility functions are not rational.

Inconsequentialism—the conjunction of the three claims above—is a startling philosophical thesis. If there are vague propositions, then there is distinction to draw between precise and vague propositions. The Boolean ground rules afford us a circular way of drawing the distinction. A proposition is precise if and only if it is a disjunction of precise indices. But ideally we would break out of the circle, characterizing precise indices without appeal to precision or vagueness, and inconsequentialism affords us the ability to so. According to inconsequentialism, precise indices are the strongest objects of rational intrinsic desire: maximal disjunctions of indices that are equal in utility, relative to every rational utility function. Think of it this way: if inconsequentialism holds, and we draw a distinction between indices just when some rational utility function maps them to distinct values, then what we get at the end of our distinction drawing is the partition of logical space into precise indices. According to inconsequentialism, rational credence is finer-grained than rational utility.

If $p$ and $q$ are consistent, and every rational utility function maps every index that is a disjunction of $p \vee q$ to the same value, then it is not rational to prefer (i.e. strictly prefer) $p$ to $q$. Inconsequentialism thus predicts that, if $p$ and $q$ are intraprecise, it is not rational to prefer $p$ to $q$, and it is that prediction of inconsequentialism that I subject to

scrutiny. As I said above, I accept the conditional consequentiality of the vague: I believe that it is rational to have intraprecise preferences, if there are vague propositions.

**§5/ Pain**

Bacon does not offer any master argument for V-to-I. He takes it as a postulate, content to let it inherit support garnered by the overall theory.[8] But something about the postulate resounds. For many of us there is a deeply felt sense that it does not matter where in the borderline region the division between cases and non-cases lies. Since we are assuming bivalence, we are assuming that every borderline region has some such division with it. Some number of hairs is the most hairs one could pluck from my head without making me bald. Some nanosecond was the last nanosecond I was a child. But there remains a deeply felt sense that it does not matter where in the borderline region the division between cases and non-cases lies, and V-to-I is a natural regimentation of that suspected inconsequentiality.

And there's more: some things which are consequential, if precise, seem inconsequential, if vague. Pain is a good example.

The presence of pain is the proposition that there is pain, *o*. The absence of pain is its negation, ~*o*.

[8] Although see Bacon (2018: 195-201).

There could be borderline pains. So, if there are vague propositions, then the presence of pain is possibly borderline. Any proposition that could be borderline is vague, and the negation of any vague proposition is vague. So:

(5.1) If there are vague propositions, then the absence of pain is vague.

The truth of (5.1) is half of an argument against the V-to-I, under the supposition that there are vague propositions. The other half is:

(5.2) It is rational to intrinsically desire the absence of pain.

Everyone accepts (5.1), inconsequentialists notwithstanding. And, at first blush, (5.2) appears to be a Moorean fact. What could be more obviously consequential than pain?

But, in fact, (5.2) is doubtful. It is rational to intrinsically desire the absence of pain, if the absence of pain is precise, but it is not clear that it is rational to intrinsically desire the absence of pain, if the absence of pain is vague.

Suppose—simplifying things, to get the crux of the matter more clearly into view—that degrees of painfulness are precise and linearly ordered. If $d_x$ is some degree of painfulness, let $x$ be the proposition that the painfulness of something equals or exceeds $x$. Also suppose—again to the crux of the matter more clearly into view—that pain supervenes on degree of painfulness: that every index agrees that $o$ is necessarily equivalent to some $x$ or other.

Desiring that ~$x$ is desiring that $d_x$ exceed the painfulness of each thing, a desire that can be satisfied only if the degree of painfulness of each thing is sufficiently low. Each ~$x$ is rational to intrinsically desire, and ~$o$ is identical to some ~$x$, if every proposition is precise.

But the arguments for vague propositions are as strong here as anywhere. Even if ~$o$ and ~$x$ are necessarily equivalent, it seems rational to give middling credence to ~$o$, conditional on ~$x$, and vice versa. If there are vague propositions, then ~$o$ is distinct from each ~$x$. Each index agrees that some ~$x$ or other is necessarily equivalent to ~$o$, but they disagree about which. Some say that it is necessary that there is pain just if the painfulness of something equals or exceeds *this* degree, that $o$ is necessarily equivalent to *this* $x$. Others says that it is necessary that there is pain just if the painfulness of something equals or exceeds *that* degree, that $o$ is necessarily equivalent to *that* $x$. Intraprecise indices agree about how painful each thing is, and thus agree about the truth-value of each $x$. But they disagree about which borderline pains are pains—they disagree about which degree of painfulness is the necessary threshold for pain—and thus sometimes disagree about the truth-value of $o$.

Desiring that ~$o$, desiring the absence of pain, is desiring that the necessary threshold for pain exceed the painfulness of each thing. If ~$o$ is vague, if indices disagree about what the necessary threshold for pain is, then the desire that ~$o$ can be satisfied either by the degree of painfulness of each thing being sufficiently low or by the necessary threshold for pain being sufficiently high. Intrinsically desiring that ~$o$ is thus

bizarre, if ~*o* is vague; for someone who intrinsically desires that ~*o* then is intrinsically concerned with what the necessary threshold for pain is.

The point may be even clearer if put in terms of preference. Suppose that a sensation is a borderline pain if its degree of painfulness is $d_x$, and suppose that precise index *y* settles that the most painful thing is painful to degree $d_x$. If someone intrinsically desires that there be no pain, and has no other relevant intrinsic desires, then they will prefer *y* & ~*o* to *y* & *o*: they will prefer there being no pain and things being precisely thus-and-so to there being pain and things being precisely thus-and-so. Let *t* be the proposition that the necessary threshold for pain is less than or equal to $d_x$. While *o* and *t* are distinct, *y* & o and *y* & *t* are identical. To prefer *y* & ~*o* to *y* & *o* is thus to have a conditional preference about where the necessary threshold for pain is: it is to prefer *y* & ~*t* to *y* & *t*; it is to prefer things being precisely thus-and-so and the necessary threshold for pain being greater than $d_x$ to things being precisely thus-and-so and the necessary threshold being less than or equal to $d_x$. And that is a bizarre preference. Degrees of painfulness seem to screen off the consequentiality of pain: it is bizarre to prefer something that is not a pain and painful to degree $d_x$ to something that is a pain and painful to degree $d_x$.

I do not say that (5.2) is false. If the absence of pain is vague, then intrinsically desiring the absence of pain is bizarre, as is the pattern of preference it begets. But bizarre intrinsic desires may be rational. There is the Humean view:

(5.3) Every utility function is rational.

If the Humean view is true, then the irrationality of an intrinsic desire does not follow from its bizarreness.

This paper does not attempt to argue against the Humean view. The Humean view implies the conditional consequentiality of the vague, which I seek to defend.

It is inconsequentialists who must argue against the Humean view, and that is none too easy for them to do. Restrictive conceptions of rational intrinsic desire deem bizarre intrinsic desires irrational. But, like the Humean view, inconsequentialism is highly permissive. It deems rational many bizarre intrinsic desires. Inconsequentialists say that it is not rational to intrinsically desire that a glass of water be at least pretty full, but they say that it is rational to intrinsically desire that the glass of water be at least 65.72% full, even if the objects of those bizarre intrinsic desires are necessarily equivalent. Inconsequentialists are thus committed to a hyperintensional distinction in bizarreness: a distinction between the rationality-excluding bizarreness of intrinsically desiring that a glass of water be at least pretty full and the rationality-non-excluding bizarreness of intrinsically desiring that the glass be at least 65.72% full. It is not clear that the needed distinction is there to be drawn. I have not seen a convincing argument against the vague Humean view, which accepts both vague propositions and (5.3).

But I am interested in whether the conditional consequentiality of the vague can be defended without assuming a permissive conception of rational intrinsic desire.

The argument above might have seemed to fit the bill. The absence of pain would seem to be something that it is rational to intrinsically desire, even supposing a restrictive

conception of rational intrinsic desire. But if bizarre intrinsic desires not rational, then (5.2) is true only if the absence of pain is precise.

At the outset I mentioned the distinction between things that are consequential only if precise and consequential even if vague. A strong case can be made, I think, that pain is consequential only if precise.

**§6/ Morality**

Many cases pattern as pain does. Someone who prefers a relationship that is a friendship and precisely thus-and-so to a relationship that is not a friendship and precisely thus-and-so has a bizarre friendship fetish. One wants to say to such a someone, by way of rhetorical therapy: who cares if the relationship is a *friendship*, if it and the world around it are precisely thus-and-so? Someone who prefers appreciating a beautiful sunset that is precisely thus-and-so to appreciating a not beautiful sunset that is precisely thus-and-so has a bizarre beauty fetish. One wants to say, by way of rhetorical therapy: who cares if the sunset is *beautiful*, if it and the world around it are precisely thus-and-so? But not every case patterns as pain does. In the next section we consider, among other things, truth. It is not bizarre to prefer a belief that is true and precisely thus-and-so to a belief that is false and precisely thus-and-so. The rhetorical therapy, if tried—who cares if the belief is *true*, if the belief and the world around it are precisely thus and so?— backfires. This section considers the challenge that moral permissibility and moral value pose to V-to-I.

One example of borderline moral permissibility that has received attention in the literature is the following, presented and discussed in Schoenfield (2016: 262):

> Darryl is watching his two-year-old daughter play in the city park. It is permissible to divert his attention for one second. It is not permissible to divert his attention for five minutes. Is it permissible to divert his attention for 30 seconds? 31? 32? Plausibly, we can create a Sorites series, admitting of borderline cases of permissibility, out of a series of diversions whose lengths differ by a second.

This example is not rare or special. It may be borderline permissible for me to stand where I am standing, if you have not given me permission to be in your personal space, and the closest part of me to you is a borderline case of being in your personal space. Creating an additional someone may be borderline permissible, if the life the additional someone leads is borderline worth living. Aborting fetuses, instilling beliefs and values in children, the tradeoffs between duration and intensity of pain, liberty and security, flourishing and equality all supply examples of borderline moral permissibility. But it is helpful to focus on a particular case, and the case above suits the argumentative needs.

Let *d* be a possible diversion that is borderline morally permissible: a diversion lasting 30 seconds, say. Let P*d* be the proposition that *d* is morally permissible. If there are vague propositions, then some precise index *y* settles that *p*, but does not settle whether P*d*: the indices in *y* agree that *d* is true, but disagree about whether P*d* is true. If

each of *y* & P*d* and *y* & ~P*d* are consistent, it is possible to prefer *y* & P*d* to *y* & ~P*d*, and the following claim is inconsistent with inconsequentialism:

(6.1) If there are vague propositions, then it is rational to prefer *y* & P*d* to *y* & ~P*d*.

I want to suggest that we begin our evaluation of (6.1) in moral psychology, reflecting on what we expect psychologically from the decent people around us. Preferring *y* & P*d* to *y* & ~P*d* is not bizarre. It is commonplace. It is a preference that we expect Darryl to have, if he is decent.

Contrast pain. It is bizarre to prefer a sensation that is not a pain and painful to degree $d_x$ to a sensation that is a pain and painful to degree $d_x$, and that bizarre preference, if had, infects other parts of one's psychology. Someone with that preference, upon learning or supposing that there is a sensation that is painful to degree $d_x$, will worry that the sensation is a pain, and hope that it is not. Aand those are bizarre worries and hopes.

If Darryl prefers *y* & P*d* to *y* & ~P*d*, then that preference will affect his psychology in parallel ways. Upon learning or supposing *y*, he will worry that he acted impermissibly, and hope that he acted permissibly. But those worries and hopes are bizarre. They are expectable. In fact, their absence would be more bizarre than their presence. Upon learning or supposing *y*, Darryl is rationally uncertain whether he acted impermissibly. If he is not worried that he acted impermissibly, if whether he acted impermissibly is nothing to him, then we worry about him and his decency.

This point merits reiterating. Fill in the details: Darryl is screening on his smart phone, and then, after 30 seconds, suddenly remembers that he is at the city part with his daughter, and that he alone is tasked with watching after her. If degrees of negligence are precise, specify the degree to which the diversion is negligent. If degrees of self-indulgence are precise, specify the degree to which the diversion is self-indulgent. The diversion and the world around it are, in moral and non-moral respects, precisely thus-and-so. Impermissibility and mistreatment are penumbrally connected in this case: every index agrees that Darryl mistreated his daughter if and only if he diverted his attention impermissibly. Upon learning or supposing *y*, Darryl is rationally uncertain whether he mistreated his daughter. If he is not worried that he mistreated his daughter, if he does not hope that he did not mistreat his daughter, if whether he mistreated his daughter is nothing to him, then we worry about Darryl and his decency.

One can try the rhetorical therapy: who cares if the diversion is *permissible*, if the diversion and the world around it are precisely thus-and-so? Who cares if Darryl *mistreated* his daughter, if his treatment of her and the world around it are precisely thus-and-so? But the therapy does not hit home as it does in other cases.

Every preference among intraprecise propositions is a conditional preference about where in the borderline region the division between cases and non-cases lies, and the present example is no exception. Say that the threshold for permissible diversions is *no shorter than d* just if it a necessary truth that any diversion in circumstances like Darryl's that is not longer than *d* is permissible. If Darryl prefers *y* & P*d* to *y* & ~P*d*, then he prefers things being precisely thus-and-so and the threshold for permissible diversions

being no shorter than *d* to things being precisely thus-and-so and the threshold for permissible diversions being shorter than *d*. But as strange as that might seem from a purely point of view, that is an altogether ordinary preference to have. Parents often worry that they have diverted their attention for impermissibly long, and their worry is not entirely owed to uncertainty about the length or circumstances of the diversion. Conditional on the diversion and the world around it being precisely thus-and-so, they hope that they have diverted their attention permissibly, and worry that they have not. So, conditional on the diversion and the world around it being precisely thus-and-so, they hope that the threshold for permissible diversion is longer than their diversion, and worry that it is not.

An exhausted parent who loves their child might want to divert their attention for as long as they permissibly can. Where in the borderline region the division between permissible and impermissible diversions lies is not nothing to them—and they do not seem to be rendered irrational on that account.

There are other moral challenges to inconsequentialism.

Take consent. There are borderline cases of consent. So, if there are vague propositions, then some precise index *y* settles that you touched someone, but does not settle whether they consented to being touched by you. If it is rational to prefer, conditional on *y*, you touching them and them consenting to it to you touching them and them not consenting to it, then inconsequentialism fails.

Take rights. There are borderline cases of having a right. So, if there are vague propositions, then some precise index *y* settles that you prevented someone from using a

life-saving measure, but does not settle whether you had the right to prevent them. If it is rational to prefer, conditional on $y$, you preventing them and having a right to prevent them to you preventing them and not having a right to prevent them, then inconsequentialism fails.

But one particularly interesting moral argument against inconsequentialism focuses on comparative relations of moral value.

One precise index may be borderline morally better than another. In fact, each of two precise indices may be borderline morally better than the other.

Let $y$ be a precise index at which someone has a long, not very intense pain. Let $x_1$, …, $x_{1,000,000}$ be precise indices that are just like $y$, except that they have a short pain. The pain in $x_1$ is short and not very intense: $x_1$ is morally better than $y$. The pain in $x_2$ is short and just slightly more intense: $x_2$ is morally better than $y$. The pain in $x_{1,000,000}$ is very intense: $y$ is morally better than $x_{1,000,000}$. It seems that some $x$ is borderline morally better than $y$. Indeed, it seems, for some $x$, each of $x$ and $y$ are borderline morally better than the other. (Conditional on a precise index, it seems rational to give some credence to some $x$ being morally better than $y$ and also some credence to $y$ being morally better than that $x$.)

We have $x$, a precise index at which there are 99 happy people. We also have precise indices, $y_1$, …, $y_{1,000,000}$. Each $y$ has 99 happy people and one additional person. The additional person in $y_1$ is very happy indeed: $y_1$ is morally better than $x$. The additional person in $y_2$ is slightly less happy: $y_2$ is morally better than $x$. The additional person in $y_{1,000,000}$ is miserable: $x$ is morally better than $y_{1,000,000}$. If bivalence holds, then some $y$ is the first $y$ in the sequence that is morally worse than $x$: it is the best $y$ in the

sequence that *x* is morally better than. But the fact about which *y* in the sequence is the best *y* that *x* is morally better than seems epistemically inaccessible, in the same way that the fact about the greatest number of hairs one could pluck from my head without making me bald is. It seems that some *y* is borderline morally better than *x*. In fact, it seems that, for some *y*, each of *y* and *x* are borderline morally better than the other.

Abstracting away from any particular case, let M*pq* be the proposition that *p* is morally better than *q*: that it being the case that *p* is morally better than it being the case that *q*. If each of a pair of precise indices may be borderline morally better than the other, then the following claims holds:

(6.2) If there are vague propositions, then for some precise indices *x* and *y*, each of *x* and *y* is consistent with each of M*xy* and M*yx*.

The second premise of the argument concerns the rationality of a certain pattern of preference:

(6.3) If for some precise indices *x* and *y*, each of *x* and *y* are consistent with each of M*xy* and M*yx*, then it is rational both to prefer *x* & M*xy* to *y* & M*xy* and to prefer *y* & M*yx* to *x* & M*yx*.

Think about conditional preference, first. If each of *x* and *y* are consistent with each of M*xy* and M*yx*, then it seems rational both to prefer *x* to *y*, conditional on M*xy*, and to

prefer *y* to *x*, conditional on M*yx*. And if that pattern of conditional preference is rational, then it is rational both to prefer *x* & M*xy* to *y* & M*xy* and to prefer *y* & M*yx* to *x* & M*yx*.

The third premise is also concerned with the rationality of a certain pattern of preference. If it is rational both to prefer *a* to *b* and to prefer *c* to *d*, then either it is rational to prefer *a* to *d* or it is rational to prefer *b* to *c*. Hence,

(6.4) If it is rational both to prefer *x* & M*xy* to *y* & M*xy* and to prefer *y* & M*yx* to *x* & M*yx*, then either it is rational to *x* & M*xy* to *x* & M*yx* or it is rational to prefer *y* & M*yx* to *y* & M*xy*.

But the conjunction of these three claims—(6.2), (6.3), and (6.4)—is inconsistent with inconsequentialism.

As stated, the argument requires that a pair of precise indices each be borderline morally better than the other, but a very similar argument requires only that some precise index be borderline morally better than the other:

(6.2*) If there are vague propositions, then for some precise indices *x* and *y*, each of *x* and *y* is consistent with each of M*yx* and ~M*yx*.

(6.3*) If for some precise indices *x* and *y*, each of *x* and *y* is consistent with each of M*yx* and ~M*yx*, then it is rational both to prefer *y* & M*yx* to *x* & M*yx* and be indifferent between *y* & ~M*yx* and *x* & ~M*yx*.

(6.4*) If it is rational both to prefer *y* & M*yx* to *x* & M*yx* and be indifferent between *y* & ~M*yx* and *x* & ~M*yx*, then it is rational to prefer *y* & M*yx* to *y* & ~M*yx*.

I do not see how either (6.4) or (6.4*) can be denied, so inconsequentialists have two options. Either they claim that moral value is precise—that no precise index is borderline morally better than any other—or they claim that it is not always rational to prefer *p* to *q*, conditional on *p* being morally better than *q*.

The former option is bolder. The hypothesis that no precise index is borderline morally better than another is very interesting. But it is also very implausible. When one compares the moral value of precise indies, one sees the epistemic signature of borderlineness. It is not hard to produce a precise index that is borderline morally better than another. In fact, it is not hard to produce a pair of precise indices that each borderline morally better than the other. The best option for an inconsequentialist is denying (6.3)/(6.3*).

If (6.3)/(6.3*) fail, then rationality sometimes requires amorality. Sometimes one must, on pain of irrationality, weakly prefer *q* to *p*, conditional on *p* being morally better than *q*. Could indifference to moral value of that sort really be a requirement of rationality?

**§7/ Epistemology**

Inconsequentialism also has some surprising epistemological consequences. I will start with two examples, where the consequences are more tolerable than they might initially appear, then turn to a third example, where the consequence is, I think, intolerable.

*First example*: Some things are borderline cases of believing that *p*. So, if there are vague propositions, then some precise index *y* settles that I believe that ~*p*, and settles that my beliefs are consistent if I do not believe that *p*, but does not settle whether I believe that *p*. If, conditional on *y*, it is rational to prefer having consistent beliefs to having inconsistent beliefs, then inconsequentialism fails.

Initially, the rationality of that preference might seem obvious. But bearing down on the example cast doubt on the initial impression.

If *y* does not settle whether I believe that *p*, then I am intimately related to *p*. Perhaps I affirm it under some guises, and deny it under others; perhaps I take it for granted in reasoning, but disavow it; perhaps my attitude is on the border between imagining and believing. Believing that ~*p* and being intimately related to *p* in the way that I am is epistemically bad in much the way that having inconsistent beliefs is, irrespective of whether the precise intimate relation I bear to *p* constitutes belief. The question is thus subtle. Is it rational to care about what I *believe*, over and above the precise relations I bear to propositions? Is it rational to prefer having consistent beliefs and being related to propositions precisely thus-and-so to having inconsistent beliefs and being related to propositions precisely thus-and-so? It is not obvious that it is. Belief may pattern as pain does.

*Second example*: There are borderline cases of having a credence function. So, if there are vague propositions, then some precise index *y* does not settle whether I have *C*, a rational (and probabilistic) credence function, or *D*, an irrational (and non-probabilistic) credence function. If, conditional on *y*, it is rational to prefer having *C* to having *D*, then inconsequentialism fails.

Initially, the rational of the preference might seem obvious. But, again, bearing down on the example cast doubt on the initial impression.

If *y* does not settle whether I have *C* or *D*, then I am intimately related to both. Perhaps one better predicts my behavior by interpreting me as having *D*, but is more charitable to me by interpreting me as having *C*. Every precise relation I bear to every credence function is settled by *y*. What *y* does not settle, in this case, is what credence function I have. But is it rational to care intrinsically about *having*, over and above the precise relations I bear to credence functions? Is it rational to prefer having a rational credence functions and being precisely related to credence functions thus-and-so to having an irrational credence function and being precisely related to credence functions thus-and-so? It is not obvious that it is. Having may pattern as pain does.

*Third example*: There are borderline cases of truth. So, if there are vague propositions, then some precise index *y* settles that I believe that *p*, but does not settle whether *p* is true. If, conditional on *y*, it is rational to prefer truly believing *p* to falsely believing *p*, then inconsequentialism fails.

And, of course, the same threat is posed by other epistemic attitudes and activities. We instead could suppose that *y* settles that I give 0.6 credence to *p* and 0.4 credence to

~*p*, or that *y* settles that I guess that *p*. Whether something is a pain seems inconsequential, once we have specified how things are precisely, and, according to inconsequentialism, whether something is true is inconsequential, if its truth-value is not settles by how things are precisely, even if it is the content of our beliefs, guesses, and credences.

Here, again, the rationality of the preference initially might seem obvious, but in this case the initial impression is correct. Things being precisely thus-and-so settles many things about *p*. It settles that I believe that *p* (or guess that *p*, or give more credence to *p* than to its negation). But it does not settle whether *p* is true, and truth is consequential even if vague. Caring about what you believe, over and above how things are precisely, may be an irrational fetishizing of belief. Caring about credence function you have, over and above how things are precisely, may be an irrational fetishizing of having. But caring about whether epistemic attitudes and activities, over and above how things are precisely, is not an irrational fetishizing of truth. It is an appreciation of the truth-directedness of our epistemic attitudes and activities. It is an appreciation of the epistemic significance of accuracy.

If accuracy were not a matter of truth, things might be different. Say that a proposition is *determinately* true just if it is true and not borderline. Precise index *y* settles that my belief that *p* is not determinately true, we may suppose. If accuracy were a matter of determinate truth, then perhaps it would not be rational to prefer believing *p* truly to believing *p* falsely, conditional on thing being precisely thus-and-so.

But accuracy is a matter of truth. If bivalence holds, then, like necessary truth, determinate truth is truth plus something else. In the present context determinate truth is something like: being true and scrutable from the precise. Accuracy is not sensitive to what determinate truth adds to truth. If each of us guesses what number of hairs is the most hairs one could pluck from my head without making me bald, and every candidate number is guessed, then one of our guesses is accurate, even though none of our guesses are determinately true. It is not just the disjunction of what we have guessed that is accurate. Some guess is accurate, and every other guess is inaccurate. Accuracy is transparent to truth. If *p* is true and borderline, then a belief that *p* is accurate and borderline accurate.

Truth is as fine-grained as propositions are. For each index there is a belief that is more accurate at it than at any other index. Rational utility is thus as fine-grained as rational credence is, *pace* inconsequentialism.

References

Bacon, Andrew (2018): *Vagueness and Thought*, Oxford: Oxford University Press.

Constantinescu, Cristian (2014): "Moral Vagueness: A Dilemma for Non-naturalism," in Shafer-Landau (ed.) *Oxford Studies in Metaethics, vol. 9*, 152-85.

Dougherty, Tom (2014): "Vague Value," *Philosophy and Phenomenological Research* 89: 352-72.

Harman, Elizabeth (2015): "The Irrelevance of Moral Uncertainty," *Oxford Studies in Metaethics*: 53-78.

Hedden, Brian (2016), "Does MITE Make Right?: On Decision-Making Under Normative Uncertainty," in Shafer-Landau (ed.), *Oxford Studies in Metaethics, vol.11*, 102-28.

MacAskill, William, Krister Bykvist, and Toby Ord (2020): *Moral Uncertainty*. Oxford: Oxford University Press.

Schoenfield, Miriam (2016): "Moral Vagueness Is Ontic Vagueness," *Ethics* 126: 257-82.

Sud, Rohan (2019): "Moral Vagueness as Semantic Vagueness," *Ethics* 129: 684-705.

Weatherson, Brian (2019): *Normative Externalism*. Oxford: Oxford University Press.

------ (2014), "Running Risks Morally," *Philosophical Studies* 167: 141-63.