Caspar Hare March 2017
Draft

The Great Spectrum Paradox

1. Spectra

Population axiologists hope to shed light on central questions in population ethics (How many people should we want there to be? How well off should we want them to be? What if these things are in tension?) by ranking populations that differ with respect to the number of people they contain, and with respect to how well off those people are. But the enterprise of population axiology has, for thirty five years, been overshadowed by certain paradoxes – collections of propositions that are *individually truthy* (each looks true, at first glance), but *jointly inconsistent* (they cannot all be true). Here is one of the simplest[1]:

<u>The Paradox of Repugnance</u>
Consider a hundred different populations. Within each population a number of equally well off people ever exist. The populations differ with respect to how many people there ever are and how good their lives are – ranging from wonderful, level 100 lives, to barely-worth-living, level 1 lives.

	Pp.1	**Pp.2**	**Pp.3**		**Pp.98**	**Pp.99**	**Pp.100**
The Number of People	10^{11}	10^{12}	10^{13}	…	10^{108}	10^{109}	10^{110}
How Good Their Lives Are	100	99	98	…	3	2	1

Three individually truthy, but jointly inconsistent propositions:

Ascent Each successive population in this series is better than its predecessor. (Ten times as many people exist, only marginally less well off!) So Population 2 is better than Population 1, Population 3 is better than Population 2... etc.

Drop Population 100 is not better than Population 1. In Population 1 a very large number of people lead wonderful lives. That's just great. In

[1] Made famous by Derek Parfit in Chapter 17 of Parfit (1984).

Population 100 an almost unfathomably large number of people lead lives barely worth living. That's not great at all.

Transitivity — *Better than*, restricted to populations, is a transitive relation. For any populations *a*, *b*, *c,* if *a* is better than *b* and *b* is better than *c*, then *a* is better than *c*.

We know[2] that, outside of population axiology, there are many other paradoxes with a very similar form. Here is a nearby paradox that involves just one person:

<u>The Paradox of Heavenly Holidays</u>

God offers you a package holiday in Heaven. Packages differ with respect to how long you spend in Heaven, and with respect to the degree of Heaven you get to experience – from degree 100 Heaven (ecstatic bliss at every moment, never diminishing or getting old), to degree 1 Heaven (just good enough at every moment for suicide not to be a rational choice). Consider some packages.

	Pk1	**Pk2**	**…**	**Pk98**	**Pk99**	**Pk100**
Days in Heaven	10^{11}	10^{12}	…	10^{108}	10^{109}	10^{110}
Degree of Heaven	100	99	…	3	2	1

Individually truthy, but collectively inconsistent propositions:

Ascent — Each successive package in this series is better for you than its predecessor. (You get to spend ten times as long in a heaven, and only marginally less good a heaven!)

Drop — Package 100 is not better for you than Package 1.

Transitivity — *Better for you* is a transitive relation.

That particular paradox many not seem to you to be very important. Realistically speaking, you probably do not expect to be choosing between package holidays in Heaven anytime soon. But, realistically speaking, you probably do expect sometime soon to be choosing between longer

[2] Mostly due to Larry Temkin in Temkin (1996), and Temkin (2012), but with important contributions from Alistair Norcross in Norcross (1997) and Stuart Rachels in Rachels (1998).

periods of lesser goods-to-you and shorter periods of greater goods-to-you, and you probably do expect sometime soon to be choosing between more certain lesser goods-to-you and less certain greater goods-to-you. The paradox points to deep difficulties for the project of weighing up packages of lesser and greater goods and deciding which, on the whole, benefits you more.

And a paradox of the same general form has played a very important role in another area of ethics, the ethics of charity:

<u>The Paradox of Sacrifice</u>

Marcia is in position to save innocent strangers at a cost to her. Consider some situations in which Marcia saves ever more innocent strangers, and makes ever-greater sacrifices to do so – ranging from 0.1% sacrifice, in which she gives up a very small amount of money, to 100% sacrifice, in which she gives up all her money, every moment of her leisure, her shot at love, happiness career-fulfillment etc.

	St.1	**St.2**	**St.3**	**…**	**St.998**	**St.999**	**St.1000**
Marcia's Sacrifice (%)	0.1	0.2	0.3	…	99.8	99.9	100
The number she saves	1	2	3	…	998	999	1000

Say that one situation *morally dominates* another if any decent person should prefer that the one situation, rather than the other, comes about. Three individually truthy, but collectively inconsistent propositions:

Ascent Each successive situation in this series morally dominates its predecessor (Marcia saves an extra life, at a marginal extra cost to herself – everybody, even Marcia, should, on pain of indecency, prefer that!)

Drop Situation 1000, in which Marcia gives up everything and saves 1000 people, does not morally dominate Situation 100, in which Marcia gives up a lot and saves 100 people. (Marcia may prefer, without being indecent, merely to make a big sacrifice, rather than to sacrifice everything.)

Transitivity *Morally dominates* is a transitive relation. For any situations *a*, *b*, *c*, if everyone should, on pain of indecency, prefer *a* to *b*, and everyone should, on pain of indecency, prefer *b* to *c*, then everyone should, on pain of indecency, prefer *a* to *c*.

A small minority of philosophers reject *Drop*. Peter Singer is right, they say. We are indecent for failing to give our all. The paradox is really a convincing argument from *Ascent* and *Transitivity* to the negation of *Drop*. The majority of us take *Drop* to be an obvious truth. But we squirm a little, because *Ascent* and *Transitivity* look like obvious truths too.

The three paradoxes we have seen so far involve normative relations – *better than*, *better for you than*, and *morally dominates*. But paradoxes of this general form need not involve normative relations. Here's one with *stronger than*:

The Paradox of Strength

Consider some materials that differ only with respect to compressive strength (the ability of a material to resist deformation under compression -- measurable in megapascals) and tensile strength (the ability of a material to resist deformation under tension – again measurable in megapascals).

	Mt.1	**Mt.2**	**Mt.3**	**…**	**Mt.98**	**Mt.99**	**Mt.100**
Compressive Strength (mp)	10	10^2	10^3	…	10^{98}	10^{99}	10^{100}
Tensile Strength (mp)	10	9.9	9.8	…	0.3	0.2	0.1

Individually truthy, but collectively inconsistent propositions:

Ascent Each successive material in the series is stronger than its predecessor (it is ten times stronger under compression, and only marginally less strong under tension!)

Drop The last material is not stronger than the first material. (The first material is strong under tension and compression. Give me a fist-sized lump of the material and I won't be able either to stretch it or to squeeze it. The last material, though almost perfectly incompressible, is laughably weak under tension. I can pull it apart like foam bubbles.)

Transitivity — *Stronger than* is a transitive relation.

And (tradition demands it) here's one with *balder than*[3]:

The Paradox of the Hirsute

Consider some guys with medium-long hair, who differ with respect to scalp hair mass (measurable in grams) and scalp coverage (measurable as a percentage of scalp covered by hair).

	G.1	**G.2**	**G.3**	**…**	**G.98**	**G.99**	**G.100**
Scalp-hair mass (g)	10	10^2	10^3	…	10^{98}	10^{99}	10^{100}
Scalp coverage (%)	50	49.5	49	…	2	1.5	1

Individually truthy, but collectively inconsistent propositions:

Ascent — Each guy in the series is balder than his successor (he has a tenth as much scalp hair, and only marginally greater coverage!)

Drop — The first guy is not balder than the last guy. (The first guy has a decent enough head of hair. The last guy is bald, but for a pony tail denser than a neutron star.)

Transitivity — *Balder than* is a transitive relation.

There are many, many paradoxes just like this.

2. Many Paradoxes or One Paradox?

The paradoxes obviously have something in common. Roughly put: There's a superiority relation, and whether it holds between things is determined by how those things score along two dimensions. The paradox arises because we are inclined to think that any tiny

[3] Larry Temkin describes an example like this in the appendix of Temkin (2012). The idea is from Ryan Wasserman.

decrease along one dimension can be made up for by a sufficiently large increase along the other dimension, but some large decreases along the one dimension cannot be made up for by any large increases along the other dimension, but the relation is transitive. These three thoughts cannot all be right.

Does this mean that they not different paradoxes, just different versions of the same paradox? I don't think that it matters very much how we individuate paradoxes. The interesting question is whether there is a satisfactory solution to one of the paradoxes that readily generalizes, so as to serve as a satisfactory solution to them all.

A solution to a paradox identifies the false proposition (or *a* false proposition, if there's more than one), explains why it is false, and explains why it is truthy – why it seems to be true. A satisfactory solution puts you in a position to say this: "I see which proposition is false. I see why it is false. And I see why I mistakenly thought it was true. I can now, without pause, retract my earlier judgment."

Some philosophers suggest highly non-generalizable solutions to some of the paradoxes. So, with respect to the Paradox of Repugnance, some philosophers say that *Ascent* is false. They say that it is false because *comparativism* is true, because it is a mistake to make positive evaluative comparisons between outcomes in which different people exist.[4] But it is truthy because our commendable practice of balancing off well-being within lives (eg. 'It is better for me overall that I suffer now, so that I prosper later') lulls us into thinking that we can correctly balance off well-being across lives (eg. 'It is better overall that Anna suffer, so that Beth may prosper.' or 'It is better overall that Anna exist than that Beth exist.')[5]

[4] How, then, can we do population ethics? For sustained efforts to do population ethics while avoiding inter-personal evaluative comparisons, see Narveson (1967) and Roberts (1998), (2002).

[5] The locus classicus for this thought is Nozick (1974) pp. 32-3. See also Thomson (1993) and Thomson (2008).

Whether or not this is right (I do not think it is right[6]) it certainly doesn't tell the whole story, because it doesn't tell us about what is going on in the nearby Paradox of Heavenly Holidays, or in any of the other paradoxes. None of these paradoxes involve comparing worlds in which different people ever exist, or making inter-personal evaluative comparisons of any kind. This is rather like a solution to the Paradox of the Liar that turns on some idiosyncratic feature of *Cretans*, or a solution to the Sorites Paradox that turns on some idiosyncratic feature of *heaps*. The solution may or may not be right so far as it goes, but we can easily replace 'Cretans' with 'Athenians', 'heaps' with 'piles', and we have redescended into paradox.

Other philosophers have suggested that in the Paradox of Repugnance and the Paradox of Heavenly Holidays, *Drop* is false. It is better that there be 10^{110} people with lives barely worth living than that there be 10^{10} people with magnificent lives. It is better for me that I spend 10^{110} days almost at the point of rational suicide than that I spend 10^{10} days in ecstatic bliss. But they say that *Drop* is truthy because, in order to assess the merits of 10^{110} people with lives barely worth living, or of 10^{110} days almost at the point of rational suicide, we need properly to imagine these things. We think we can properly imagine these things ("It's a *lot* of people. It's a *long* time.") but we can't. The numbers are just too big. If we could properly imagine these things then we would realize how wonderfully good they are.[7]

Again, whether or not this is right (Again I do not think it is right. What do I need to do in order properly to imagine being near-suicidal for 10^{110} days? Do I need to spend 10^{110} days imagining being nearly suicidal? It is true that I cannot do that, but I have low confidence that if I did then I would or should emerge on the other side saying "That was just *wonderful*! Far better for me, overall, than 10^{10} days of ecstatic bliss!") it cannot be the whole story, because it

[6] I do not think it is right, for reasons I explain in Part II of Hare (2013)).

[7] See, for example, John Broome in Broome (2004) pp. 55-59. The strategy is carefully discussed, and resisted, by Theron Pummer in Pummer (2013).

doesn't tell us anything about what is going on in the other paradoxes. I can well enough imagine a material that is perfectly resistant to compression, but that stretches easily under tension. I can well enough imagine Marcia working tirelessly throughout her normal-length life to save the lives of 1000 people.

Yet other philosophers have suggested that, in at least many cases like this, *Transitivity* is false. [8] This is a suitably general response, but it really scuppers the whole project of ranking things. Suppose the relation *f-er than* is intransitive. Now we cannot establish a straightforward connection between whether one thing is *f-er than* another, and *how f* each of the things are. We cannot establish a straightforward connection between whether one thing is *f-er than* another and which of three things is the *f-est*. And, if we care about getting things that are *f-er than* other things, we cannot in any straightforward way use the resources of decision theory to guide our choices.

My aim in this paper is to develop a solution that is satisfactory, general and non-scuppering.

3. A Change of Topic: Vague Terms in Conversation

Let's begin with what may seem like a complete change of topic. In this section I will make some observations about how one-place vague terms, terms like 'gray' and 'tall', function in conversations. These observations do not amount to anything close to a complete theory of

[8] This is Larry Temkin's view in Temkin (1996). Temkin is well aware of the costs that come with denying transitivity. In Temkin (2010) he is less insistent that transitivity must go – the paradoxes reveal a tension that must somehow be resolved.

vague language, but they are data points that any complete theory of vague language must account for.[9]

To be clear, through repetition, the observations do not have to do with properties like *being gray* and *being tall*. They have to do how we use terms like 'gray' and 'tall'. And, again to be clear through repetition, the observations do not have to do with how we use two-place terms, like 'grayer than' and 'taller than'. They have only to do with how we use one-place terms, like 'gray' and 'tall'.

I will take 'gray' as my primary example, because grayness is easy to represent on paper.

Vague One-Place Terms Make Distinctions

When multiple things are mutually salient to speaker and listener, 'gray' may serve to distinguish some things from others.

The Distinctions they Make are Sensitive to What is Mutually Salient to Speaker and Listener

Whether a thing is distinguished as 'gray' may depend on which other things are mutually salient to speaker and listener.

The Distinctions they Make are Coarse

Speakers do not use terms like 'gray' to distinguish between mutually salient things that are very close with respect to grayness.

To get a sense of what all this means, let's work through some examples:

[9] For a more complete theory that suggests these things see Agustin Rayo (2008). I will borrow the finding-my-house trope from him.

<u>Conversation 1</u>

I give you directions to my house: "Go to Mud Street and you will see it. It's a gray house". Arriving at Mud Street, you face these two houses:

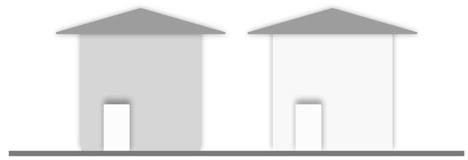

You knock on the door of the left house.

What has happened? I have used a vague term, 'gray', to distinguish the left house as mine.

<u>Conversation 2</u>

I give you directions to my house: "Go to Mud Street and you will see it. It's a gray house". Arriving at Mud Street, you face these two houses:

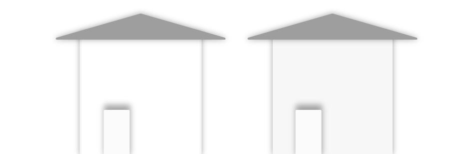

You knock on the door of the right house.

Again I have used a vague term, 'gray', to distinguish a house as mine. But this time it is the right house. The right house is distinguished as gray in this conversation, in which different houses are mutually salient to speaker and listener.

<u>Conversation 3</u>

I give you directions to my house: "Go to Mud Street and you will see it. It's the leftmost gray house". Arriving at Mud Street, you face just these six houses:

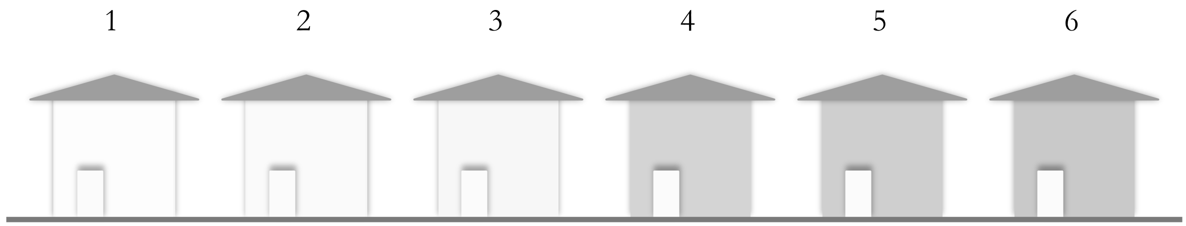

You knock on the door of house number 4.

What has happened? Though (e.g.) house 3 is noticeably grayer than house 2, and house 5 is noticeably grayer than house 4, they are close with respect to grayness. By convention we do not use ‘gray’ to distinguish between things that are close with respect to grayness. Houses 3 and 4 are not close with respect to grayness. Because they are not close, I am distinguishing 4 as gray, 3 as non-gray.

<u>Conversation 4</u>

I give you directions to my house: “Go to Mud Street and you will see it. It’s the leftmost gray house”. Arriving at Mud Street, you face just these six houses:

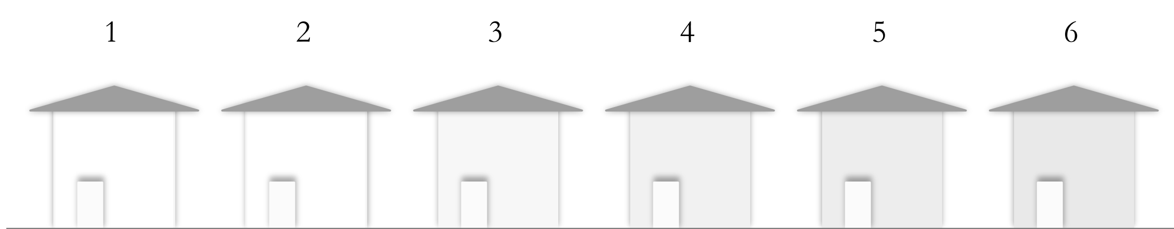

You knock on the door of house number 3.

Last time I distinguished house 3 as non-gray. This time I distinguish it as gray. Why? Well, this time, though house 4 is noticeably grayer than house 3, they are close with respect to grayness. This time houses 2 and 3 are not close with respect to grayness.

Conversation 5

I give you directions to my house: "Go to Mud Street and you will see it. It's the leftmost gray house". Arriving at Mud Street, you face these thirteen houses:

1 2 3 4 5 6 7 8 9 10 11 12 13

You curse my poor directions. You knock on the door of house 7. But that's a bit of a guess. On the basis of my directions you could just as well have knocked on the door of house 6, just as well have knocked on the door of house 8.

Why are my directions poor? Why are you *guessing*? By convention we do not use the term 'gray' to distinguish between mutually salient possibilities that are close with respect to grayness. But, also by convention, we do not count things like house 1 as gray. House 1 is unquestionably white. So I am in some way violating conventional norms here. And my violation is unhelpful because, on the basis of the conventional norms, plus what I said, plus what you see before you, you are not in a position to know which house is mine. (Note that epistemicists about vagueness[10] will say that I have managed to pick out a house, and maybe I picked out *my* house. No matter. Epistemicists and their opponents agree that it is not possible for you to know, on the basis of my directions, which house is mine. In that sense epistemicists and their opponents agree that these are bad directions.)

[10] Following Timothy Williamson in Williamson (1994).

Conversation 6

I give you directions to my house: “Get to Mud Street and you will see it. It’s the leftmost gray house”. Arriving at Mud Street, you face just these two houses:

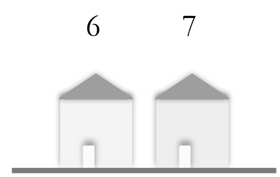

You knock on the door of house 6.

Last time, given my directions, you could just as well have knocked on the door of house 7 as the door of house 6. You would have been making no mistake by knocking on the door of house 7. This time, given the very same directions, you would be making a mistake by knocking on the door of house 7. Why? Because this time, though my use of the term ‘gray’ does not serve to distinguish one house from another (they are close with respect to grayness), my use of the term ‘leftmost’ does, and there is no convention that things like house 6 don’t count as gray.

4. Context Slippage

Because the import of directions is sensitive in the above delicate ways to what the speaker takes to be salient to speaker and listener, if you are trying to follow my directions, you must keep careful track of the things that I took to be mutually salient to us when I gave my directions. It is very easy for you to confuse the possibilities that I took to be mutually salient to us with the possibilities that are presently salient to you, and thereby become confused about what best satisfies my directions. Here is a situation that really invites such confusion:

<u>Conversation 7</u>

I your boss, and you, my assistant, stand before a range of holiday houses, considering which one to rent for a week on my behalf. Here are the houses and their weekly rates:

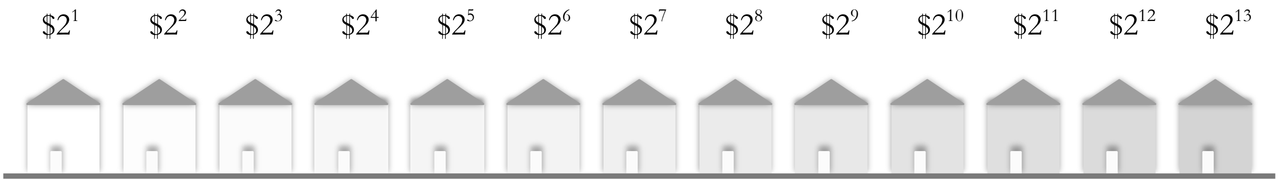

I say "First and foremost what I want is a gray house. Secondarily I want to spend less money. So please book me the cheapest gray house." and leave. You curse my poor instructions. Was it that I had a particular house in mind (the $\$2^9$ house, say) and I did a very crude job of identifying it to you, or was it just that I had certain values (Thumbs up to grayness! Thumbs up to inexpensiveness!) and I did a crude job of expressing them? You don't know. But either way, you do not take it of any house that booking it best satisfies my instructions. You decide to book the $\$2^7$ house, without taking it that the $\$2^7$ house satisfies my instructions better than the $\$2^8$ house or the $\$2^6$ house... But, before you have a chance to make that official, the rental agent tells you "Actually only two of the houses are available – the $\$2^6$ house and the $\$2^7$ house. The rest are all spoken for." So you focus hard on just those options:

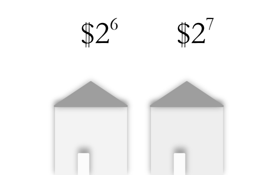

The $\$2^6$ house is only fractionally less gray. You really have to stare hard to see the difference in grayness. But it is half the price. Only a *grayness fanatic* would distinguish the $\$2^7$ house from the $\$2^6$ house on the grounds that the former is gray and the latter not. And nothing about me has suggested to you that I am a grayness fanatic. You conclude that the $\$2^6$ house satisfies my instructions better than the $\$2^7$ house, and book it.

Are you right? If so then, by similar reasoning, in a pairwise choice between any house and its predecessor, the predecessor better satisfies my directions. But, in a pairwise choice between the $\$2^7$ and $\$2^1$ houses, the $\$2^7$ house clearly better satisfies my directions – I said that first and foremost I wanted a gray house! So the relation *better satisfies my directions in pairwise choice* must be intransitive.

But you are not right. You have made a mistake. When comparing the $\$2^6$ and $\$2^7$ houses, you confused the collection of things that were at that time salient to you with the collection of things that I took to be mutually salient to us when I issued my instructions. Had I taken just those two houses to be mutually salient to us then the $\$2^6$ house would indeed have satisfied my instructions better. But, given that I took all thirteen houses to be mutually salient to us, it does not satisfy my instructions better. (Of course the mistake turned out to be harmless, because the $\$2^6$ house doesn't satisfy my instructions *worse* than the $\$2^7$ house. You might as well take the $\$2^6$ house. But still, it was a mistake.)

There don't need to be two people around for this confusion to arise. The 'conversation' can be between earlier and later incarnations of you:

<u>Conversation 8</u>

You are preparing to visit a rental agent and book yourself a holiday house. Not knowing exactly what options you will have, you resolve to book the cheapest gray house. To fix the resolution, you write in your diary: "Book the cheapest gray house!" The agent then shows you your two options:

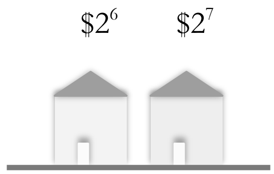

> You look at them, and then back to the words in your diary. You think to yourself: "The $\$2^6$ house is only marginally less gray. It would be *crazy* to distinguish the two houses on the grounds that one is gray and the other not. But the $\$2^6$ house is much cheaper. I can better satisfy my own instructions, and hence better abide by my resolution, by booking the $\$2^6$ house."

You have again made a (harmless) mistake. If the instructions in your diary had been written by someone who took just those two options to be salient to herself and her addressee, then indeed those instructions would have been better satisfied by booking the $\$2^6$ house. But they weren't.

5. Back to the Spectrum Paradoxes

What do these observations about the way we use vague one-place terms, and about the confusion that can arise from misusing them, have to do with the spectrum paradoxes?

Our goal, remember, is to identify the false proposition in each of the spectrum paradoxes, explain why it is false, and explain why we (in particular: *you*) were initially inclined to think it true.

As a first step, note that, when we set about determining whether an ordering relation holds between things, there are sometimes procedures we can follow. Sometimes, these procedures can be described.

So, to take a bland example, if you are interested in determining which of two rectangles has *greater area*, then you will do well to follow these instructions:

Instructions for Comparing Rectangles

What matters is the product of length and width. So pick the one with the greatest product of length and width.

The instructions are short, accurate, and free of vague terms. We can give such instructions because *greater rectangular area than* is a very accessible and well-behaved relation. Whether one rectangle has greater rectangular area than another is fully determined by their lengths and widths – things that can be represented by numbers. We know just how the determination works: it is about which triangle has a greater product of length and width.

Generally, when there is a simple function f from pairs of numbers to numbers such that a thing with values w,x bears the relation to a thing with values y,z iff $f(w,x) > f(y,z)$, and we know just what the function is, we can give simple instructions for determining whether one thing bears the relation to another, free of vague terms, just by specifying the function.

But sometimes it is not so easy. For one thing, sometimes the ordering relations we are interested in are not so well-behaved. Sometimes the relation forms an *incomplete order*. The relation is *irreflexive* (not for some x, xRx), *antisymmetric* (not for some x,y xRy and yRx), *transitive* (for all x,y,z if xRy and yRz then xRz), but *negatively intransitive* (for some x,y,z, not xRy, not yRz, but xRz). In that case there is no function that determines the order in the manner above (though there will be a larger set of functions such that a thing with values w,x bears the relation to a thing with values y,z iff for every function f in the set, $f(w,x) > f(y,z)$).

For another thing, sometimes we don't know the precise extension of the ordering relation we are interested in – it depends in subtle, hard-to-know ways on our present interests and goals. (What further we want say about these cases will depend on what theory of vagueness we want to adopt. Epistemicists say that there is one relation we are interested in, with precise extension, but that it is hard or impossible for us to know what it is. Vagueness-in-the-world theorists say that there is one relation we are interested in, but that its extension is indeterminate. Supervaluationists say that there is no one relation we are interested in, just a family of relations. No matter. For present purposes all we need is that, in these cases, we don't know the precise extension of the ordering relation we are interested in.)

In some of these more difficult cases one-place vague terms naturally figure in our best descriptions of our best procedures for determining whether the ordering relations obtain.

Suppose, for example, that you and I are engineering the chassis of a car. Suppose that it is important that the materials we use not deform under the normal compressive and tensile stresses to which a car chassis is subject in acceleration, cornering etc. We don't want the chassis to break apart! Suppose I set about ordering materials by strength. My best description of a procedure for determining which of two materials is *overall stronger*, in the way that interests me, may then be something like this:

> *Instructions for Comparing Materials with Respect to Overall Strength*
> First and foremost what matters is that the material we use be both *strong under compression* and *strong under tension.* Secondarily what matters is the sum of compressive and tensile strength. So pick the material that is both strong under compression and strong under tension, with the greatest sum of compressive and tensile strength. Among materials that are not both strong under compression and strong under tension, pick the material with the greatest sum of compressive and tensile strength.

Now, if you wish to follow instructions like this, the words alone may underdetermine how best to do it. Sometimes, to know how to follow instructions like this, you may need to know something about the context in which the instructions were issued. So, for example, you may need to be reminded of what is at stake. Which materials count as 'strong under compression' and 'strong under tension'? We are engineering a car chassis. That sets the bar for 'strong under compression' and 'strong under tension' at something around 8.5 megapascals. If we were engineering a children's toy then the bar would be much lower. If we were engineering an aircraft wing then the bar would be much higher.

And, as we saw in sections 4 and 5, you may need to know what materials I took to be mutually salient on issuing the instructions. Suppose, for example, that you are comparing materials in the series that figures in the Paradox of Strength:

	Mt.1	**Mt.2**	**Mt.3**	**…**	**Mt.98**	**Mt.99**	**Mt.100**
Compressive Strength (mp)	10	10^2	10^3	…	10^{98}	10^{99}	10^{100}
Tensile Strength (mp)	10	9.9	9.8	…	0.3	0.2	0.1

Suppose, to precisify the example, that you are comparing materials 15 and 16 (both on the edge of deforming under normal-for-car-driving tensile stresses). If the *Strength Instructions* were issued in a context in which all possible materials were mutually salient, then neither material 15 nor material 16 would satisfy them better. The instructions would be less than fully helpful. You would not be making a mistake by picking material 15. You would not be making a mistake by picking material 16. If, on the other hand, the instructions were issued in a context in which only materials 15 and 16 were mutually salient, then material 16 would satisfy them better than material 15. 'Strong under tension' is a vague one-place term. By convention we do not use it to distinguish between materials that are very close with respect to tensile strength. But material 16 most definitely has a greater sum of compressive and tensile strength than material 15. And the same goes for any adjacent pair of materials in the series – if the *Strength Instructions* were issued in a context in which just those two materials were mutually salient, then the higher-numbered material would satisfy those instructions better.

So, if you are comparing just two materials, with the *Strength Instructions* in mind, there are two (two of many really, but these are the natural ones) heuristics that you might wish to adopt:

The Global Heuristic
Pick the material that best satisfies the *Strength Instructions*, issued in a context in which all possible materials are mutually salient.

The Local Heuristic
Pick the material that best satisfies the *Strength Instructions*, issued in a context in which just the two materials you are focusing upon are mutually salient.

Which of these two heuristics is better? That depends on the nature of the relation you are interested in. The *Local Heuristic* is an accurate guide to an intransitive relation – each material in the series better satisfies the Strength Instructions than its predecessor (issued in a context in which just those two materials are mutually salient), but the first better satisfies the Strength Instructions than the last (issued in a context in which just those two materials are mutually salient). If you are not interested in an intransitive relation, if you want to *rank* the materials, then you are interested in a relation that forms at least a partial order, and the *Global Heuristic* is better.

This gives us a solution to the Paradox of Strength:

Which proposition is false? *Ascent* is false.

Why is *Ascent* false? It is false because there is something in the series that is no stronger than its predecessor. As we move along the series, the materials get stronger for a while. We are in the realm of materials strong under compression and tension. Each thing satisfies the *Strength Instructions* (issued in a context in which all possible materials are mutually salient) better than its predecessor. Then they no longer get stronger. We are in the realm of materials that are on the edge of being strong under tension. Material 16 satisfies the *Strength Instructions* (issued in a context in which all possible materials are mutually salient) no better than material 15. Then they get stronger again. We are in the realm of materials that are not strong under tension. Material, e.g., 100 satisfies the *Strength Instructions* (issued in a context in which all possible materials are mutually salient) better than material 99.

Why is *Ascent* truthy, though false? It is truthy because, when we compare adjacent materials in the spectrum we apply the local heuristic, where we should be applying the global heuristic. We confuse the materials presently salient to us (just these two) with the materials we should imagine to be mutually salient in the context in which the instructions we are trying to follow were issued (all possible materials.) This is a very easy mistake to make, as I hope the examples in section 5 show. It takes a great deal of mental effort to avoid it.

I suggest that, broadly speaking, this solution generalizes. Focus back on the Paradox of Repugnance, and the spectrum of populations that it involves.

	Pp.1	**Pp.2**	**Pp.3**		**Pp.98**	**Pp.99**	**Pp.100**
The Number of People	10^{11}	10^{12}	10^{13}	…	10^{108}	10^{109}	10^{110}
How Good Their Lives Are	100	99	98	…	3	2	1

The paradox has to do with which of these populations are *better than* which others. When two populations differ only with respect to the number of people they contain, and the well being of those people, what is a generally applicable procedure for determining which is better? That's a very hard question. Much of population ethics is about that question. I will not give a precise answer here. But I will conjecture this: Some vague one-place terms will figure in our best description of our best procedure for determining which population is better. Our best description will be something like this (this one is crude, but it will do as an example):

Instructions for Comparing Populations

First and foremost what matters is that there be *enough people* with *good enough lives*. Secondarily what matters is total well-being – the product of population size and average well-being. So if one population has enough people with good enough lives then pick it. If both or neither do, then pick the one with the greatest total well-being.

Notice that 'enough people' and 'good enough lives' are one-place vague terms. So the words in the instructions may underdetermine how best to satisfy them. To have a useful heuristic we need to add something about the context in which the instructions are issued. And if we are interested in *ranking* populations the natural one is:

The Global Heuristic

Pick the material that best satisfies the *Instructions for Comparing Populations*, issued in a context in which all possible materials are mutually salient.

So, in the Paradox of Repugnance, which proposition is false? *Ascent* is false.

Why is *Ascent* false? It is false because there is something in the series that is no better than its predecessor. As we move along the series, the populations get better for a while. We are in the realm of populations in which enough people have good enough lives. Each thing satisfies the *Instructions for Comparing Populations* (issued in a context in which all possible populations are mutually salient) better than its predecessor. Then they no longer get better. We are in the realm of populations of people whose lives are on the edge of being good enough. Population (e.g.) 53 satisfies the *Instructions for Comparing Populations* (issued in a context in which all possible populations are mutually salient) no better than population 52. When issued in a context in which all possible populations are mutually salient, the *Instructions for Comparing Populations* are unhelpful. You may as well pick 53. You may as well pick 52. Then the populations get better again. We are in the realm of populations in which people do not lead good enough lives. Population, e.g., 100 satisfies the *Instructions for Comparing Populations* (issued in a context in which all possible populations are mutually salient) better than population 99.

Why is *Ascent* truthy, though false? Because when we compare (eg.) populations 51 and 52, we mistakenly apply the local heuristic. We confuse the populations presently salient to us (just these two) with the populations we should imagine to be mutually salient in the context in which the instructions we are trying to follow were issued (all possible populations.) This is the familiar, very easily made mistake.

Finally, let's turn back to the Paradox of Sacrifice and the spectrum of situations that it involves:

	St.1	**St.2**	**St.3**	**…**	**St.998**	**St.999**	**St.1000**
Marcia's Sacrifice (%)	0.1	0.2	0.3	…	99.8	99.9	100
The number she saves	1	2	3	…	998	999	1000

The paradox has to do with which situations *morally dominate* others (I remind you that one situation morally dominates another when any decent person should prefer that the one come about). What is a general procedure for determining whether one situation morally dominates another? Again that's a difficult question. Again, without fully answering the question, I conjecture that that some vague one place terms will figure in our best description of our best procedure for determining whether one situation morally dominates another. Our best description will be something like this (again this description is crude, but it will do as an example):

Dominance Instructions

First and foremost what matters is that Marcia is entitled to prefer that she retain some autonomy, some semblance of a fulfilled and self-directed life. Secondarily what matters is general human well-being. So pick the situation with higher general human well-being, unless Marcia retains some autonomy only in the situation with lower general human well being.

'in which Marcia retains some autonomy' is a vague one-place term. To have a useful heuristic we need to add something about context. The natural way to go is:

The Global Heuristic
Pick the situation that best satisfies the *Dominance Instructions*, issued in a context in which all possible situations (involving Marcia making sacrifices) are mutually salient.

So which proposition is false? *Ascent* is false.

Why is *Ascent* false? It is false because there is a situation in the series that does not dominate its predecessor. For a while, as we move along the series, each situation dominates its predecessor. We are in the realm of situations in which Marcia retains some autonomy. Then we come to situations that don't dominate their predecessors. We are in the realm of situations in which Marcia is on the edge of retaining some autonomy. Then we come to situations that, once again, dominate their predecessors. Marcia has given up her autonomy, she may as well sacrifice yet more.

Why is *Ascent* truthy, though false? Because when we compare (e.g.) situations 500 and 501, in which Marcia is on the edge of losing her autonomy, we mistakenly apply the local heuristic.

7. Wrapping Up

A theme has emerged here. In each paradox the false proposition is *Ascent*. It is truthy, though false, because when we compare things in the relevant way we mistakenly apply a local heuristic, where we should be applying a global heuristic. This mistake is easily made. Applying the proper, global heurist is tricky. It involves carefully distinguishing the things presently salient to you from another set of things – the things that you should take to be part of your context when you apply the heuristic.

I should emphasize that I am not trying to explain why people in general mistakenly believe *Ascent*, or why most people believe *Ascent*. People believe all kinds of things for all kinds of different reasons. I am hoping, rather, to offer you an explanation of why you mistakenly believed *Ascent* before you read this paper – and I am hoping thereby to leave you feeling free to discount that belief.

The explanation works for me. A closing true story: Some time ago I was struck by a very bourgeois desire for a stereo system. I wanted a nice one, though perhaps not so very nice that I would feel like a fool for failing to appreciate its subtleties. So I resolved to 'get something quality, though not crazy-expensive', and went to a store. I found that the owner-manager-audiophile had fourteen different sets of speakers that he would recommend, eight amplifiers that he would recommend, three CD players that he would recommend, and four grades of speaker wire that he would recommend, all at different price points – making for 14x8x3x4 = 1,344 systems for me to buy, all at different price points.

Fortunately he also had patience. So together we plugged in different components, listened for bass, circled through options for quite a while, until we settled on a system that I was happy with. One thing, though, gave me pause. The speaker wire was $110, which was $100 more than I had ever spent on speaker wire. Could it really be worth it? The fellow suggested that maybe it was not worth it. Speaker wire makes a difference, but the difference is *miniscule*. To demonstrate, he played a piece on the system, then swapped out the high grade, $140 wire for medium grade, $50 wire, and played the same piece again. I had to agree. Any differences were *miniscule*. Then, to seal the point, he swapped out the speakers for much lower grade speakers. This, he said, would make a big, obvious difference. And I had to agree again. It sounded much less good.

The conclusion he expected me to draw from this was that, whether I went with the medium-grade or high-grade wire, I was getting a quality system, which is what I wanted, so I might as well spend less money. And that was the conclusion I drew, until thought about it carefully. We were focusing on three systems. There was the high-grade CD player, amp, speakers and speaker wire system. Call it Q+. There was high-grade CD player, amp, speakers, and medium-grade speaker wire system. Call it Q. And there was the low-grade speakers system. Call it L. My intention, walking into the store, was to get a quality system, spending less money rather than more. If I had formed an intention properly expressed in that way in a context in which just Q+, Q and L were salient then indeed Q would have satisfied it best. But, given that I formed my intention in a context in which all possibilities were salient, Q satisfied it no better than Q+.

So I went for Q+, and the guy, of course, thought I was nuts (though nuts in the good, *nuttily committed to spending money at my store* sort of way). But what was interesting was that I too, at some level, thought I was nuts. It hurt to buy that expensive wire. I felt the sort of conflict between cognition and instinct that comes when you are separated from an enormous drop by pane of glass. You believe that you are safe, but at the same time... I believed I was not extravagantly throwing money away, but at the same time... This suggests to me that, at least in my own case, the pull of context-slippage is strong and visceral. If others are like me, then it is no wonder that the various different versions of the Great Spectrum Paradox have for so long seemed so very mysterious.

References

Broome, John. 2004. Weighing Lives, Oxford: Oxford University Press

Hare, Caspar. 2016. *The Limits of Kindness*, Oxford: Oxford University Press

Narveson, Jan. 1967. "Utilitarianism and New Generations", *Mind* 76: 62–72.

Nozick, Robert. 1974. *Anarchy, State and Utopia*, New York: Basic Books

Pummer, Theron. 2013. "Intuitions About Large Number Cases", *Analysis* 73 (1): 37–46

Rayo, Agustín. 2008. "Vague Representation", *Mind* 117: 329-373

Roberts, Melinda. 1998. *Child versus Childmaker: Future Persons and Present Duties in Ethics and the Law*, Lanham, MD: Rowman & Littlefield.

Temkin, Larry. 1996. "A Continuum Argument for Intransitivity", *Philosophy and Public Affairs* 25 (3): 175-210.

Temkin, Larry. 2012. *Rethinking the Good: Moral Ideals and the Nature of Practical Reasoning*, Oxford: Oxford University Press, 2012.

Rachels, Stuart. 1998. "Counterexamples to the Transitivity of Better Than", *Australasian Journal of Philosophy* 76 (1): 71-83

Williamson, Timothy. 1994. *Vagueness*, Abingdon: Routledge Press