

University of Kent
 Canterbury CT2 7NF
 l.goldstein@kent.ac.uk

References

- Goldstein, L. 1986. Epimenides and Curry. *Analysis* 46: 117–21.
 Goldstein, L. 1994. A Yabloesque paradox in set theory. *Analysis* 54: 223–27.
 Kripke, S. 1975. Outline of a theory of truth. *The Journal of Philosophy* 72: 690–716.

Intuitions about large number cases

THERON PUMMER

1.

An intuition that proposition P is true is a mental state of its seeming to be the case that P (which is distinct from the belief that P), and a *reliable* intuition that P is an intuition which gives those who have it a reason to believe P.

Many of us have intuitions about certain large number cases in ethics. For example, consider:

Hangnails for Torture. For any excruciatingly painful torture session lasting for at least two years to be experienced by one person, there is some large number of minute-long very mildly annoying hangnail pains, each to be experienced by a separate person, that is, other things equal, worse.¹

In this case, many of us have the strong intuition that Hangnails for Torture is false. That is, we have the intuition that there is *no* number of such mild hangnail pains that would be worse than 2 years of excruciating torture. While it is true that not everyone shares the intuition that Hangnails for Torture is false, my aim here is not to argue that those who do not are broken normative thermometers. Rather, my concern is whether this intuition is reliable, that is, whether those who have it thereby have a reason to believe that Hangnails for Torture is false.

Though I will focus on Hangnails for Torture, most of what I will argue applies, *mutatis mutandis*, to similar large number cases, particularly the

1 This sort of case is discussed by Larry Temkin (1996), Alastair Norcross (1997) and Stuart Rachels (1998).

Repugnant Conclusion.² Because such large number cases are at the centre of several deep puzzles in contemporary ethical theory, the epistemic status of our intuitions about these cases is of great importance.

2.

Several excellent philosophers have claimed that since we cannot relevantly imagine very large numbers, we should discount the intuition, which many of us have, that Hangnails for Torture is false.³ And some of these philosophers sometimes suggest that if we *could* relevantly imagine very large numbers, we might have a different intuition; namely, we might have the intuition that Hangnails for Torture is true. I call such philosophers *Large Number Sceptics*.

To help explain the Large Number Sceptic's position, I will first highlight what it is not. Large Number Sceptics do not maintain that all of our claims about cases involving large numbers should be discounted. For instance, consider the claim that 7 billion mild hangnail pains would be worse than 100 thousand mild hangnail pains, other things equal. The Large Number Sceptic would have no complaint here, since we can plausibly justify this claim without having to imagine large numbers. We can simply appeal to the plausible principle that 'more pains are worse, other things equal'.

On the other hand, claims about the truth of Hangnails for Torture seem different, in that we cannot plausibly justify them by appealing to the principle that 'more pains are worse, other things equal'. Why not? Because in Hangnails for Torture, it is not the case that other things are equal. In comparing the 2 years of torture with the very many hangnail pains, we have to determine whether the difference in intensity between torture and mild hangnail pain is compensated by the difference in number, or duration. There is 'only' one 2-year-long torture, but arbitrarily many hangnail pains. For the many minute-long hangnail pains to be worse than the 2 years of torture, it seems there would need to be the equivalent of billions or even trillions of *years* of them (it would take about 525,000 such pains to make up just 1 year). To determine whether the many hangnail pains *are* worse, it

2 The *Repugnant Conclusion* states that:

For any possible population of at least ten billion people, all with a very high quality of life, there must be some much larger imaginable population whose existence, if other things are equal, would be better, even though its members have lives that are barely worth living (Parfit 1984: 388).

3 They include: J.M.E. McTaggart (1927: 453), Alastair Norcross (1997: 146–52), Joshua Greene (2001), Torbjörn Tännsjö (2002: 344), John Broome (2004: 55–59), Michael Huemer (2008: 907–11) and Adam Cureton (2009). While not all of these philosophers have addressed Hangnails for Torture in particular, their remarks imply a general scepticism about the reliability of intuitions about large number cases which would clearly apply to this particular large number case.

seems we would need to relevantly imagine not just the difference in intensity between torture pain and hangnail pain, but also the difference between 2 years of torture and the equivalent of at least billions of years of hangnail pains.

Of course, Large Number Sceptics do not deny that we can mathematically calculate the number of minutes in a billion years. But they'd insist that, to intuitively weigh the badness of the many minute-long hangnails against the badness of the torture, we must also have some sort of *intuitive grasp* of the extremely large number of hangnails. Most of us have a fairly decent intuitive grasp of numbers under 1,000; we have a *feel* for the differences between such familiar, 'medium-sized' numbers. Do we have a comparably decent intuitive grasp of *billions* of years of minute-long hangnail pains? It is not clear that we do. And it is perhaps not unreasonable to claim that we *cannot* imagine such large numbers, no matter how hard we try. John Broome writes:

[W]e have no reason to trust anyone's intuitions about very large numbers, however excellent their philosophy. Even the best philosophers cannot get an intuitive grasp of, say, tens of billions of people. That is no criticism; these numbers are beyond intuition. But these philosophers ought not to think their intuition can tell them the truth about such large numbers of people (2004: 56–7).

In a nutshell, the Large Number Sceptic worries that we will fail to have an accurate sense of the differences between the torture session and the hangnail pains. While we will appreciate the difference in intensity between the torture and the hangnail pain, we might imagine the equivalent of only *hundreds* of years of hangnail pains, instead of billions. Our limited grasp of large numbers, it seems, inappropriately biases our intuitions in favor of the claim that the 2 years of torture is worse, rendering these intuitions unreliable. Here, then, is the Large Number Sceptic's argument:

- (1) We cannot relevantly imagine very large numbers.
- (2) In order to have reliable intuitions about Hangnails for Torture, we must relevantly imagine very large numbers.

So, (3) We cannot have reliable intuitions about Hangnails for Torture.

3.

Recall that to deny Hangnails for Torture is to claim that there is *no* number of mild hangnail pains that would be worse than 2 years of excruciating torture. I will argue that we can support this claim, on the basis of our intuitions, even if we cannot relevantly imagine very large numbers of hangnail pains (e.g. the equivalent of billions of years of such pains). Though I am not sure there isn't a relevant sense in which we can grasp very large

numbers, I will simply grant (1) to the Large Number Sceptic. But, crucially, I think we should ask: what *would* our intuitions about Hangnails for Torture likely be, *if we could* relevantly imagine very large numbers of hangnail pains?⁴ And I will argue as follows:

- (4) We are justified in believing that, if we could relevantly imagine any number of mild hangnail pains, we would have the intuition that Hangnails for Torture is false.
- (5) If (4) is true, then we have a reason to believe that Hangnails for Torture is false.

So, (6) We have a reason to believe that Hangnails for Torture is false.

First consider (5). The Large Number Sceptic claims that our intuitions about Hangnails for Torture are unreliable because they are formed under bad conditions, where we fail to relevantly imagine very large numbers. Thus, if our intuitions about Hangnails for Torture were formed under suitably idealized conditions, where we did not fail to relevantly imagine very large numbers, then these intuitions would not be unreliable, and so they would give us a reason to deny Hangnails for Torture. (I am here assuming that these intuitions are reliable apart from the worry raised by Large Number Sceptics). Finally, it seems that if we are now justified in believing that we would have the intuition that Hangnails for Torture is false under these suitably idealized conditions, then we now have a reason to believe that Hangnails for Torture is false.

More generally, we can have reasons to believe P, not on the basis of whether P seems true in our actual, non-idealized condition, but on the basis of our justified belief that P would seem true to us under suitably idealized conditions. Still more generally, if we are now justified in believing that reliable detectors of the truth of P (which include suitably idealized versions of ourselves) would say ‘P’, then we’ve now got a reason to believe P. It thus seems hard to deny (5).

I will now argue for (4), the claim that Hangnails for Torture would seem false to us under suitably idealized conditions. (I’ll be dropping the ‘we are justified in believing that’ clause at the beginning of (4) merely for convenience).

4 In granting (1) to the Large Number Sceptics, I am not granting that it is logically impossible for us to relevantly imagine large numbers, but only that it is impossible in some weaker sense, e.g. psychologically impossible. Thus, in asking what our intuitions would be if we could relevantly imagine large numbers, I am inquiring about what would be true in logically possible worlds in which our capacities for imagining large numbers were suitably improved.

4.

Here are two competing hypotheses:

- (4) If we could relevantly imagine any number of mild hangnail pains, we would have the intuition that there is *no* number of such pains such that it is worse than 2 years of excruciating torture.

And (predictably),

- (7) If we could relevantly imagine any number of mild hangnail pains, we would have the intuition that there is *some* number of such pains such that it is worse than 2 years of excruciating torture.

How could we determine which of these two hypotheses is more plausible? Many of us who have the intuition that Hangnails for Torture is false will say that (4) simply seems more plausible, and that the Large Number Sceptic has offered no positive support whatsoever for (7). We can, I think, say more than this. We can argue that (4) is more plausible than (7), based on an *extrapolation* from certain reliable intuitions that we have. Consider the following:

Variable Claim. Two years of excruciating torture is worse than X years of very mildly annoying hangnail pains (each a minute long), other things being equal.

Since we can, I believe, relevantly imagine what it would be like to have such very mildly annoying hangnail pains for a year, we can relevantly imagine approximately 525,000 such pains. We can also relevantly imagine a range of familiar, ‘medium-sized’ numbers of people each having very mildly annoying hangnail pains for a year, and, in this sense, imagine the equivalent of many years of such pains. (For instance, we can relevantly imagine something at least as bad as the equivalent of 16 years of hangnail pains, each pain realized in a separate person, by imagining 16 people each having 1 year of such pains.)

Consider the Variable Claim, and imagine substituting in progressively larger numbers for X. For example, six intuitions we might have are as follows:

- 2 years of excruciating torture is worse than 4 years of very mildly annoying hangnail pains (each a minute long).
- ‘...’ is worse than 8 years of ‘...’
- ‘...’ is worse than 16 years of ‘...’
- ‘...’ is worse than 32 years of ‘...’
- ‘...’ is worse than 64 years of ‘...’
- ‘...’ is worse than 128 years of ‘...’

These intuitions cannot plausibly be discounted by the Large Number Sceptic. We can specify that the 2 years of torture is the worst such 2-year

period that we can relevantly imagine, filled with agonizing physical pain, terror, depression, etc. (So, by definition, it is relevantly imaginable). Moreover, we can imagine at least hundreds of people who would each suffer approximately 525,000 hangnail pains. Arguably, we can imagine tens of thousands of people who would each have such pains for a year (some football fans arguably have a decent intuitive grasp of the number of people in a crowded stadium).

One might object: ‘you are here imagining people having many hangnail pains, whereas in Hangnails for Torture each such pain is had by a separate person. So these intuitions you have cited are irrelevant.’ But this objection fails. For if we find the Variable Claim plausible when the X years of minute-long hangnail pains are packed into lives such that each life receives a year’s worth, then we would surely find this claim at least as plausible if the X years of minute-long hangnail pains were spread very thinly across the lives of separate persons such that each life receives just one such pain.

I can now offer the *Extrapolation Argument* for (4):

- (8) If (7) were true, and thus if (4) were false, then we would become less confident in the Variable Claim, the larger we imagine X to be.
- (9) We do *not* become less confident in the Variable Claim, the larger we imagine X to be.

So, (4) is true and (7) is false.

(8) seems plausible because, if (7) were true and thus there were some value for X, call it *n*, such that if we imagined *n* years of mild hangnail pains we would have the intuition that together they are worse than 2 years of excruciating torture, it seems we would gradually lose confidence in the Variable Claim as our imagined value for X gets closer to *n*. This seems true even if *n* were very large; we would presumably lose at least *some* confidence as X gets larger, if (7) were true.⁵

(9) seems plausible because the Variable Claim seems *just* as plausible, whether X is 8, 64 or 128. At least, this is how things seem to many of us who find Hangnails for Torture counterintuitive. Even as the imagined number of very mild hangnail pains gets much larger, we remain *just* as confident that 2 years of excruciating torture would be worse.

5 I should note two points here. First, even if we cannot relevantly imagine very many values for X, the Extrapolation Argument would support (4) *over* (7), if (8) and (9) were true. However, the Extrapolation Argument for (4) seems *stronger*, the more values for X we test out. Second, it might be that the Variable Claim would only seem false when X is infinity – then the explanation of the fact that we do not lose any confidence in the Variable Claim as X gets larger might be that X, though getting larger, would not be getting any *closer* to infinity. Though this is a possibility, I see no reason to believe that it is true. And indeed it seems dubious that those of us who would find the Variable Claim true for every finite value for X would find it false if X were infinity.

Some Large Number Sceptics might raise the following objection: ‘we do not remain just as confident that the Variable Claim is true, as X gets larger. Rather, we become *slightly* less confident that the Variable Claim is true, as X gets larger. So the Extrapolation Argument fails.’ I have two replies to this objection.

First, some might lose confidence in the Variable Claim, as X gets larger, because they are attending to the plausible thought that the sum of hangnails gets *worse*, as X gets larger. But when we are careful to recognize that the Variable Claim is about how the badness of hangnails *compares* with the badness of 2 years of torture, and that it might be that, while a larger sum of hangnails is always worse, it could never be *worse than* 2 years of torture, our confidence in the Variable Claim might remain constant as X gets larger. Many of us, at least when we are reflecting carefully in this way, do not lose *any* confidence in the Variable Claim as X gets larger (even when we move from a very small X to the largest X we can relevantly imagine). And I doubt many of us lose *any* confidence in the claim that it is better to prevent the torture murder of a child than it is to give L people one lollipop lick each, as our imagined L gets larger.

Second, recall that I am not here arguing that those who do not have the intuition that Hangnails for Torture is false are broken normative thermometers. Similarly, I am not here arguing that those whose confidence in the Variable Claim diminishes as X gets larger are broken normative thermometers. Rather, my concern is whether those of us who *do* have the relevant ‘anti-Hangnails for Torture’ intuitions (including the relevant levels of confidence in the Variable Claim as X increases) thereby have a reason to believe that Hangnails for Torture is false. I only claim that if (9) describes *our* intuitions then *we* have a reason to believe that Hangnails for Torture is false – since (9) and (8) imply (4), which together with (5) imply (6). I here make no claim about the reasons had by those who lack the relevant intuitions.

5.

One might suggest that it is possible to offer a *counter* extrapolation argument, which supports (7) and thus opposes (4). One might have us consider the following:

Second Variable Claim. Y very mildly annoying hangnail pains (each a minute long) is worse than Z seconds of excruciating torture, other things equal.

There are imaginable quantities for Y and Z such that the above claim might seem plausible. For example, it might seem that 10,000 such hangnail pains (about a week’s worth) is worse than 1 second of excruciating torture.

We could then substitute in progressively larger numbers for Y and Z, while preserving the ratio between them (10,000 hangnail pains for each second of torture). For example, here are some intuitions we might have:

- 10,000 hangnail pains (about a week's worth) is worse than 1 second of excruciating torture.
- 20,000 '...' is worse than 2 seconds of '...'
- 30,000 '...' is worse than 3 seconds of '...'
- 40,000 '...' is worse than 4 seconds of '...'

One might then offer the following *Counter Extrapolation Argument* for (7):

- (10) If (4) were true, and thus if (7) were false, then we would become less confident in the Second Variable Claim, the larger we imagine Y and Z to be (holding the ratio between Y and Z constant).
- (11) We do *not* become less confident in the Second Variable Claim, the larger we imagine Y and Z to be (holding the ratio between Y and Z constant).

So, (7) is true and (4) is false.

This argument is structurally similar to the Extrapolation Argument for (4). However, this Counter Extrapolation Argument for (7) is unsound. In particular, (11) is false. Many of us *will* lose confidence in the Second Variable Claim, once Z becomes 'sufficiently' large (and we'll lose confidence in this claim well before Y and Z become so large that we cannot relevantly imagine them). Our intuitions are sensitive to more than just the ratio between Y and Z. Once we start considering days, hours or even minutes of excruciating torture, many of us will find it unclear or implausible whether any number of hangnails would be worse.

At the very least, this Counter Extrapolation Argument will not convince those of us who have the relevant 'anti-Hangnails for Torture' intuitions to believe that we would, under suitably idealized conditions, have the intuition that Hangnails for Torture is true.

6.

Many of us not only find Hangnails for Torture to be implausible, but we also do not lose any confidence in the Variable Claim as X gets larger. We thus have no reason whatsoever to believe (7), and what intuitive evidence we do have to go on supports (4). Moreover, (4) and (5) together imply (6), the conclusion that we have a reason to believe that Hangnails for Torture is false. Since this conclusion is based on our intuitions, it seems our intuitions do provide some positive support for the claim that Hangnails for Torture is false, and that this support has not been defeated by the Large Number

Sceptic's argument. Recall that the Sceptic's argument goes:

- (1) We cannot relevantly imagine very large numbers.
- (2) In order to have reliable intuitions about Hangnails for Torture, we must relevantly imagine very large numbers.

So, (3) We cannot have reliable intuitions about Hangnails for Torture.

And recall that I am granting (1) to the Sceptics. Does this mean that, in resisting this argument, I deny (2)? Not necessarily. Depending on which intuitions count as intuitions *about* Hangnails for Torture, (6) might be consistent with (3). We might think that, for intuitions to be about Hangnails for Torture, they must be strictly based on direct consideration or imagination of the relevant features of this case. I am not sure whether this is true. Fortunately, it does not matter here.

Either the intuitions which support (6) are intuitions about Hangnails for Torture, or they are not. If they are, then (2) and (3) are false. If they are not, then (2) and (3) might be true, but (3) and (6) are compatible. It is fairly painless to concede that we cannot have reliable intuitions which are in some sense *about* Hangnails for Torture. What matters is that we can still plausibly claim that we have intuition-based reason to deny Hangnails for Torture.

While I have argued that the Large Number Sceptics have not defeated the intuitive support for the denial of Hangnails for Torture, I have not addressed other possible defeaters. I have also not addressed the important question of *how strong* the intuitive support for the denial of Hangnails for Torture is relative to various reasons for accepting Hangnails for Torture. This is just the sort of question we may have to answer if we hope to solve the deep ethical puzzles surrounding large number cases like Hangnails for Torture and the Repugnant Conclusion.⁶

University of California, San Diego
La Jolla, CA 92093-0119, USA
tpummer@ucsd.edu

References

- Broome, J. 2004. *Weighing Lives*. Oxford: Oxford University Press.
- Cureton, A. 2009. Moral intuitions about large numbers. Presentation, Southern Society for Philosophy and Psychology Conference. Savannah: GA.

6 For their helpful comments on earlier drafts of this article, I would like to thank Craig Agule, Caley Anderson, Dick Arneson, Amy Berg, John Broome, Adam Cureton, Mike Huemer, Charlie Kurth, Per Milam, an anonymous referee at *Analysis*, and audience members at the 2012 UC San Diego Graduate Philosophy Conference. I am also very grateful to Dick Arneson for his Valtz Chair support during spring and summer 2012. Lastly, this article was accepted for the 2012 British Postgraduate Philosophy Association Conference, which unfortunately I was unable to attend.

- Greene, J. 2001. A psychological perspective on Nozick's experience machine and Parfit's repugnant conclusion. Presentation, Society for Philosophy and Psychology, Annual Meeting. Cincinnati: OH.
- Huemer, M. 2008. In defence of repugnance. *Mind* 117: 899–933.
- McTaggart, J.M.E. 1927. *The Nature of Existence*, vol 2. Cambridge: Cambridge University Press.
- Norcross, A. 1997. Comparing harms: headaches and human lives. *Philosophy and Public Affairs* 26: 135–67.
- Parfit, D. 1984. *Reasons and Persons*. Oxford: Oxford University Press.
- Rachels, S. 1998. Counterexamples to the transitivity of better than. *Australasian Journal of Philosophy* 76: 71–83.
- Tännsjö, T. 2002. Why we ought to accept the repugnant conclusion. *Utilitas* 14: 339–59.
- Temkin, L. 1996. A continuum argument for intransitivity. *Philosophy and Public Affairs* 25: 175–210.

The logic of theological incompatibilism: a reply to Westphal

JOHN MARTIN FISCHER AND NEAL A. TOGNAZZINI

In our forthcoming paper, 'Omniscience, freedom and dependence' (Fischer and Tognazzini, Forthcoming), we argued (among other things) that Jonathan Westphal's (2011) critique of the basic argument for the incompatibility of God's foreknowledge and human freedom (in the sense that requires freedom to do otherwise) is question-begging. We also presented a similar critique of arguments offered by other philosophers, including Storrs McCall (2011).

Westphal has recently responded (2012), arguing that we have completely missed the point of his original paper. We had focused on what we contend is a dialectical infelicity in Westphal's argument; more specifically, we accused him of begging the question against the incompatibilist. Westphal contends in the new paper that we ignored his main point, which is a critique of a crucial premiss of the incompatibilist's argument.¹ He claims that his 'solution' to

1 Although Westphal himself claims to be offering a critique of the incompatibilist argument, one might instead interpret him as offering only a competing argument for compatibilism, which starts by presuming that Jones is free and that God knows what Jones will do, and then goes through the claim that God's belief depends on Jones's action to the conclusion that the two presumptions are not clearly inconsistent. (Thanks to an anonymous referee for this suggestion.) But as we point out in our original article (2012, section 5), at best this results in