

How Infinitely Valuable Could a Person Be?

Levi Durham and Alexander Pruss

Abstract

Many have the intuition that human persons are both extremely and equally valuable. This seeming extremity and equality of value is puzzling: if overall value is the sum of one's final value and instrumental value, how could it be that persons share the same extreme value? One way that we can solve the Value Puzzle is by following Andrew Bailey and Josh Rasmussen (2020) and accepting that persons have infinite final value. But there are some significant downsides to their way of thinking about values, which relies on the extended real numbers. We offer a different approach: if we model values using the hyperreal numbers, we can capture many of our intuitions about the extremity and relative equality of human value without incurring the substantial theoretical costs of using the extended real numbers. We also examine other ways of modeling the infinite value of persons and evaluate the strengths and weaknesses of these other accounts compared to ours.

Introduction

Many have the intuition that human persons are both extremely and equally valuable. This seeming extremity and equality of value is quite strange: if overall value is the sum of one's final value and instrumental value, how could it be that persons share the same extreme value? Call this the Value Puzzle.

One way that we can solve the Value Puzzle is by following Andrew Bailey and Josh Rasmussen (2020) and accepting that persons have infinite final value—infinite construed as $+\infty$. But we think that there are some very significant downsides to their approach. We offer a different suggestion: if we model values using the hyperreal numbers, we can capture enough of our intuitions about the extremity and relative equality of human value without incurring the theoretical costs of using the extended real numbers. Additionally, we compare a hyperreal model to other models and evaluate whether they satisfy the desiderata we would want from a model of infinite value. Our approach does not allow for fully equal total value, but it does at least some justice to the intuition by allowing for approximately (in a sense to be defined) equal final value. That is, the hyperreal model can satisfy the practical desiderata we would want out of a model of infinite value but not all of the theoretical desiderata.

In section one of this paper, we introduce the Value Puzzle. In section two, we consider Bailey and Rasmussen’s proposed solution to the Value Puzzle. In section three, we argue that modeling values using hyperreal numbers is superior to modeling values using the extended reals. In section four, we offer a solution to a modified version of the Value Puzzle where persons have relatively equal final value. In section five, we argue that even if Bailey and Rasmussen model their arguments using infinite cardinal numbers, it falls prey to the same problems as using the extended reals. And in section six, we compile a list of desiderata that any ideal system for the value of persons should be able to satisfy and evaluate how well different systems on offer perform.

1 The Value Puzzle

Suppose that human persons are extremely valuable. This seems quite plausible. We certainly speak and, when we are at our best, *act* as if the loss of any

human life is a loss of great value. When governments assign dollar amounts to the value of a human life for regulatory purposes, the values are usually quite high, regardless of one's instrumental value.¹ Among both philosophers and non-philosophers, we find many instances where people claim that human life is immeasurably valuable. Thus, *prima facie* we have at least some reason to think that persons are extremely valuable.

Many also have the intuition that persons are all equally valuable. Many of us find ourselves hesitant to claim that one person is worth more than another overall, where overall value is the sum of instrumental value and final value. By instrumental value, we mean the value a thing possesses but not for its own sake. Something possesses instrumental value not because of what it *is* but because of what it can *do*. (For instance, a pulmonologist's instrumental value greatly outweighs a typical layperson's during a respiratory illness pandemic, and Herbie Hancock's instrumental value greatly outweighs ours in regard to creating music.) By final value, we mean the value a thing possesses for its own sake. Something possesses final value not because of what it can *do* but because of what it *is*. (For instance, Monet's *Woman with a Parasol* would be valuable even if it served no instrumental value, and a toddler has extreme final value though she can't even drop a cookie without bursting into tears.)²

The equality and extremity of human persons' value seems quite strange, however. Even if we all share the same final value, our instrumental values differ wildly; some persons are much more capable of adding value to the world than others. If overall value is the sum of instrumental and final value, then, so long as our instrumental values differ, should not our overall values differ? Why should all persons share the same, extreme value? And is it even plausible that

¹For the U.S, the value depends on the agency, but the Federal Aviation Agency values a human life at \$9.6 million (Frakt 2020).

²For more on the distinction between instrumental value and final value, see Korsgaard (1983), Kagan (1998), and Rabinowicz and Rønnow-Rasmussen (1999).

Albert Schweitzer and Al Capone share the exact same final value? We can call the conglomerate of these questions the Value Puzzle.

2 Extended Real Values

Bailey and Rasmussen (2020) think that it would be bizarre if all persons shared the same *finite* value. Whether our shared value were 10 or 10,000,000 value points, it would seem horribly arbitrary if a person's final value were some finite number. They think that this would be far less puzzling if a persons' final value were *infinite*.³ This is their proposed solution to the Value Puzzle. They offer two arguments for the infinite value of a human person. The first is an argument from confirmation. The second is a deductive argument that draws on differences in persons' instrumental values.

Their Confirmation argument goes like this: it seems that all persons have both extreme and equal value; if a person's value were finite, this seeming equality and extremity would be highly unexpected, but if a person's value were infinite, this would be expected. These two factors add up to a strong evidential argument for the infinite value of persons:

$$\frac{P(\text{Equal Human Value} \mid \text{Infinite Value Hypothesis})}{P(\text{Equal Human Value} \mid \text{Finite Value Hypothesis})} = \frac{1}{\text{low}}.$$

Therefore, they conclude, the fact that human persons have both extreme and equal value seems to be very good evidence for humans persons' having infinite value.

Bailey and Rasmussen's Instrumental Difference argument goes like this:

³They are not alone in entertaining the possibility that persons have infinite value. Kant suggests as much when he says that human dignity is such that it is above all objects of price (Kant, 1958). Some personalists, such as Karol Wojtyła (2013), suggest the same, as do Rabbi Immanuel Jakobovits (1979, p.379) and Christian theologian Matthew Anderson (forthcoming).

since human persons clearly differ in instrumental value but share the same overall value, it must be the case that human persons share the same infinite final value. They present their case as follows:

- (I1) We are all equally overall valuable.
- (I2) We differ with respect to instrumental value.
- (I3) If so, then either we differ with respect to final value, or we are all overall infinitely valuable.
- (I4) Therefore, either we differ with respect to final value, or we are overall infinitely valuable. (from I1-I3)
- (I5) We do not differ with respect to final value.
- (I6) Therefore, we are all overall infinitely valuable. (from I4-I5)

Bailey and Rasmussen think that they have sufficiently motivated (I1); however much human persons may differ, they do not differ in overall value. They take (I2) to be uncontroversially true. They say that (I3) is a consequence of being equally overall valuable; if persons are equally overall valuable (with overall value being the sum of final value and instrumental value) but differ in instrumental value, then persons must differ either in final value or have infinite final value. We do not, they assert in (I5), differ in final value, thus establishing their conclusion in (I6).

Both of Bailey and Rasmussen's arguments work only if there is only one candidate positive, infinite value in the domain in which we are modeling values. This is the case in the extended real numbers, represented by the symbol $\bar{\mathbb{R}}$. Now, Bailey and Rasmussen do not explicitly invoke the extended real numbers. However, the Instrumental Difference argument only works as stated if there is only one positive infinite value. For the purposes of this paper negative infinite

values are unimportant. Given that Bailey and Rasmussen offer no reason to deviate from standard decision theory except in respect of the positive infinite value, it is reasonable to model their account with the extended reals—i.e., finite real values together with a unique value $+\infty$ infinitely bigger than all the finite values and, in case it is needed for something beyond the scope of our paper and theirs, a unique value $-\infty$ infinitely smaller than all the finite values. It does not matter whether the positive infinite value is “really” the infinity of mathematicians’ extended reals, but for simplicity we will assume it is.

There are only two infinities in $\bar{\mathbb{R}}$: $-\infty$ and $+\infty$. This works well for modeling values in some areas, but it delivers bizarre results in others. We need to be able to compare infinite values in a way that the model does not allow. Some infinite values are greater than others.⁴

Let’s illustrate with an example. Consider the following case:

Kantian General

A general is deciding whether to bomb military complex A that contains 10 enemy soldiers or complex B that contains 10,000 enemy soldiers. She knows that bombing either one will win the war. The general thinks that persons have infinite final value. She runs her calculations: if she bombs A , there will be an infinite loss of value; if she bombs B , there will be an infinite loss of value.⁵

If the general were taking the extended reals as her domain, the options would be equivalent.⁶ Whether the general were to choose to bomb A or B , the model says that value loss would be equal.⁷ But this is deeply problematic:

⁴For a good introduction into difficulties that arise when dealing with infinities in decision theory, see Hájek (2012) and Bartha, Barker, and Hájek (2014).

⁵We are assuming proportionality and that all the conditions for a just war are met.

⁶This is part of why standard Expected Value Theory (EVT) is restricted to finite values.

For a discussion of infinite values and EVT, see Beckstead and Thomas (2020).

⁷This is assuming that when one dies their value is annihilated and not merely relocated,

the options are not equivalent. Given that B contains 9,990 more persons than A , it seems that there would be a greater value loss if the general bombed B . However, modeling infinite values using the extended reals does not allow us to demonstrate this. Adding $+\infty$ to itself still equals $+\infty$. And taking the lives of n people—regardless of how large n is—results in no net loss of value so long as there is still one person remaining with value $+\infty$.

The upshot of this is that, if the final value of one person does not always equal the final value of n persons, modeling values using the extended reals cannot capture this. The problem here is not in our intuitions about infinities, even though our intuitions about infinities are notoriously bad.⁸ The problem is the model. This model commits us to saying that if there were a global thermonuclear war that wiped out the entire human race except just one person, there would be no loss in final value. This is a significant problem for the view.

Additionally, there are problems with dealing with probabilities and infinite values lurking here: $+\infty$ swamps all non-zero probabilities.⁹ Consider this case:

Kantian Lifeguard

A lifeguard is deciding whether to save (a) a group of 10 beachgoers who are being dragged out to sea very slowly or (b) a single beachgoer who is being dragged out to sea very quickly. Having an accurate sense of her own abilities, the lifeguard knows that there is a 99% chance of saving every member in (a) and a 1% chance of saving the sole member of (b), and she knows that there is no chance at all of saving both (a) and (b). She runs her calculations: if she attempts to save (a), the expected utility is infinite; if she attempts to save (b), the expected value is infinite.

as may be the case if there's an afterlife. Thanks to [] for pointing this out.

⁸For more on this, see Arntzenius, Elga, and Hawthorne (2004).

⁹See Hájek (2003) and Colyvan et. al (2006).

Something has gone badly wrong here. The lifeguard should clearly choose (a), but on Bailey and Rasmussen’s model, (a) and (b) have the same expected utility: $+\infty$. Even if (a) involved saving 1,000,000,000 lives and (b) only one life, the expected utility would be the same. Making rational decisions involving infinite values as Bailey and Rasmussen understand them is simply unworkable.

3 Non-Archimedean Values

We need a way of comparing infinite values. A non-Archimedean ordered field allows us to do just that: this is a field V that satisfies the ordinary algebraic and order axioms of the real numbers, but contains “infinite” values α such that $|\alpha| > n$ for every (standard) natural number n .

Among philosophers, the most well-known non-Archimedean ordered fields are the fields of hyperreals, so we will focus on those.¹⁰ The hyperreals, represented by the symbol ${}^*\mathbb{R}$, are an extension of the real numbers \mathbb{R} developed by Abraham Robinson (1966) to include infinitesimals and infinities.

Mathematically, non-Archimedean ordered fields have the standard arithmetical properties that we are familiar with from high school, but allow for infinite and non-zero infinitesimal values. A positive infinite is a value M such that $N < M$ for every (ordinary) integer N and an infinitesimal is a value α such that $-1/N < \alpha < 1/N$ for every (ordinary) positive integer N . Every real number is finite and every infinitesimal is finite. A positive infinite M is the reciprocal of a positive infinitesimal $\alpha = 1/M$.

For instance, the familiar arithmetic fact that if $0 < y$ then $x < x + y$ is

¹⁰There are other non-Archimedean ordered fields that would also be good candidates. The surreal numbers (Knuth 1974) and formal power series (Pedersen 2014) would all work just fine. We also do not need all of the field structure: vector valued utilities would do the job as well (Russell 2020).

true for any y , even an infinite one, so adding a positive increment to an infinity gives a *different* infinity. Hence if we have one infinity, we have infinitely many of them. This is unlike an extended real model on which if $x = +\infty$ and y is anything other than $-\infty$, then $x + y = +\infty = x$. A non-Archimedean ordered field thus allows one to capture the intuition that even if you have something of infinite value (say, your own life), adding something of positive value to it (say, reading a good novel) is worthwhile, while adding something of negative value is worth avoiding. In particular ten thousand lives of equal (or approximately equal) infinite value have more value than ten, and hence the Kantian General problem dissolves.

Similarly, we have the familiar algebra fact that if $0 < p < 1$ and $0 < x$, then $px < x$, so a probability p of getting an outcome of value x is less valuable than the certainty of an outcome of value x . That familiar algebra fact carries over to the non-Archimedean setting unchanged, whereas for extended reals we have $p \cdot \infty = \infty$ for any positive p . Similarly, the Lifeguard problem disappears: $0.99 \cdot 10 \cdot M > 0.01 \cdot 1 \cdot M$ for an infinite value M (and the inequality remains if we replace the M on one side with something approximately equal to it).

If utilities range over a non-Archimedean field, and we only have a finite number of sure outcomes to deal with, we can make use of familiar expected utility decision theory. Thus, if the sure outcomes are A_1, \dots, A_n and have non-Archimedean utilities $u(A_1), \dots, u(A_n)$, and the probability of A_i is p_i , then the expected utility is $p_1u(A_1) + \dots + p_nu(A_n)$. Expanding this to infinitely many outcomes carries significant technical challenges, but for simple cases like Kantian General and Lifeguard we do not need this.

At the same time, we need to admit some conceptual issues here. Following Bailey and Rasmussen, we are primarily investigating the *objective value* of persons. Decision theory typically concerns agent *preferences* between “lotteries”, i.e., probabilistic mixtures of sure outcomes. Classically, one starts with such

preferences and proves that if they satisfy certain axioms of rationality, then the sure outcomes can be assigned utilities and the preferences recovered from comparisons of expected utility (von Neumann and Morgenstern, 1947). One can do something very similar in the hyperreal context, as long as one drops the continuity axiom.¹¹ One could attempt to do the same thing in the objective value sphere, by starting with value comparisons rather than numerical or hypernumerical values. But while one can talk about value comparisons between sure outcomes, it is more difficult to engage in objective value comparisons between lotteries. For one, lotteries are mixtures of sure outcomes with respect to agential uncertainty, and hence comparisons between lotteries appear innately agent-centric. We might move to lotteries that are mixtures of sure outcomes with respect to objective chances. Or one might try start with aggregative value comparisons between multiple copies of a valuable thing: three duplicates of this goat are aggregatively (i.e., without taking into account diversity and other arrangement goods) less valuable than two duplicates of this elephant, and a human has infinite value compared to that goat provided that the human is aggregatively more valuable than any finite number of duplicates of the goat.¹² How well some such approach will work is not clear, but the problems here have to be faced by anyone whose theorizing starts with objective value comparisons and then moves to agential decisions.

¹¹Fishburn (1971) gives a simplified exposition of Hausner’s (1954) vector-valued representation theorem in the case of a finite-dimensional space, i.e., one with finitely many sure outcomes. A finite-dimensional vector-valued utility is a vector (a_1, \dots, a_n) of real numbers, with lexicographic comparisons and with expected utilities being computed component-by-component. One can then, up to order-equivalence, express such a vector-valued utility as the hyperreal number $a_1 + a_2\epsilon + \dots + a_n\epsilon^{n-1}$ where ϵ is any fixed positive infinitesimal, and thereby we can obtain a hyperreal-valued representation theorem.

¹²Given some plausible axioms, aggregative value comparisons can indeed be shown to be representable via hyperreal or vector-based values. Unfortunately, the representation is not unique, even up to probabilistic decision equivalence. [reference deleted for anonymity]

That said, it is doubtful that simple more-valuable-than comparisons should be thought to capture all of the realm of value. It is plausible that there are primitive facts not just about one thing being more valuable than another, but about how much more valuable one thing is than another, whether quantitatively (three times, infinitely many times, etc.) or qualitatively so (slightly more, much more, etc.), without this reducing to comparisons between chances or between duplicates. Consider, for instance, a comparison between a deterministic world consisting of a single mindless particle staying still forever and a deterministic world consisting of a single brain thinking about mathematics forever. The latter seems *much* more valuable. But it does not seem that we can account for the “much” here by simple better-than comparisons either in terms of chances or in terms of duplicates. For it does not make sense to say, for instance, that any non-zero chance of the brain world is preferable to certainty of the particle world, because one cannot talk of a *chance* in the context of an expressly deterministic world. And aggregative value comparisons between duplicates destroy an important aspect of value: a world with exactly one particle exhibits great simplicity, an aesthetic value that will be diminished if the world is duplicated into a reality with two or more particles, even if these are in separate universes.

Before moving on, it is worth stopping to address a couple of objections that one might have to our claim that values behave more like hyperreal numbers than extended real numbers. First, one might doubt that agents have preferences that are fine-grained down to the level of infinitesimals. Do people really ever value an object an infinitesimal amount more than another object? Second, one might think that hyperreal values are much more likely to be arbitrary than real values. For instance, why should a person assign some infinite value α to an object rather than $\alpha + \epsilon$? If we are modeling values using $\bar{\mathbb{R}}$ as our domain, this isn't a problem; there are only two infinities. But it is hard to see why an

agent would assign an infinite value α to an object rather than a nearby $\alpha + \epsilon$. This seems to give us some reason to prefer $\overline{\mathbb{R}}$ over $^*\mathbb{R}$ for modeling values.

To respond to the first worry, note that following Bailey and Rasmussen our *primary* concern is with modeling objective value, not with agent preferences, though we use agent preferences to argue for the inadequacy of the Bailey and Rasmussen model. On many plausible stories about objective values, it is unlikely that our mental states can match the full precision of the objective values. Consider, after all, that if time is correctly modeled by the real numbers, the state of some finitely valuable entity, say an elephant, existing for x seconds will increase monotonically with x . But there are uncountably many positive real numbers x , and so there are uncountably many different values, while it is unlikely that a human being can represent uncountably infinitely many values, or even countably infinitely many. Additionally it need not be the case that every quantity in $^*\mathbb{R}$ represents an actually instantiable value. All we need is that the values under consideration can be represented as members of $^*\mathbb{R}$.

To respond to the second worry, we are only attempting to provide a better model for talking about values. Even if deciding between different infinite values is bound to result in an arbitrary value assignment, this isn't a problem if our only concern is modeling values. Moreover, while assigning any particular infinite value to a person may be arbitrary, if we were to assign *some* infinite value, the fact that we chose *this* value might be arbitrary, but it is not arbitrary that we chose *an* infinite value. The arbitrariness concern still remains, but it's not much of a problem for our modeling purposes.¹³ It's also worth noting that in a classical real-valued value setting, it seems reasonable to assuage much of the arbitrariness worry by saying that only ratios of values matter.¹⁴ We can

¹³See Easwaran (2014) and Pruss (2018) for analogous arbitrariness concerns about probability measures when using the hyperreals.

¹⁴In classical decision theory, utility assignments are taken to be equivalent if they are positive affine transforms of each other, i.e., translations and/or positive rescalings. However,

say the same thing in our case. Of course, we still have the question of why the ratio of values of Socrates to that of Alexander’s horse Bucephalus is that particular infinite value that it is. But even if both values are finite, there is the question of why the ratio is what it is.

Again, we are not saying that the hyperreals or any other non-Archimedean fields give us the perfect way to model values.¹⁵ But we are saying that, compared to $\bar{\mathbb{R}}$, ${}^*\mathbb{R}$ is far superior. The hyperreals allow us to do everything that the extended reals allow and more—minus many problems. There is, therefore, very good reason to use ${}^*\mathbb{R}$ for modeling values rather than $\bar{\mathbb{R}}$.¹⁶

4 Relatively Equal Final Value

Now let us consider how Bailey and Rasmussen’s Confirmation and Instrumental Differences arguments fare as arguments against the hypothesis that human value can be modeled as a positive infinite hyperreal.

On the hyperreal infinity hypothesis, we would not expect all humans to have equal value. There are too many infinities in ${}^*\mathbb{R}$. If α is infinite, so is $\alpha + .001$, as are $\alpha + .0001$, $\alpha - 10000$, α^2 , 3^α and so on, but none of these infinities equal one another. Once we have realized there are just as many positive infinite quantities as positive finite quantities,¹⁷ it would be at least as surprising that when we are looking at values of substances—such as persons—translations of assignments are not equivalent. It matters ethically that the value of a human being or a cat is positive rather than negative. Thus we do not expect translation invariance.

¹⁵As Bostrom (2011) points out, not all value theories can be modeled so easily by the hyperreals.

¹⁶We are by no means the first to suggest this approach either. Vallentyne and Kagan (1997) take the hyperreals as their domain when aggregating value across infinitely many worlds.

¹⁷If a is positive finite, then $\alpha + a$ is positive infinite for any fixed positive infinite α . So there are at least as many infinite quantities as finite ones. Conversely, if a is positive infinite, then $1/a$ is positive finite, so there are at least as many positive finite quantities as positive infinite ones.

humans have equal value given the infinite value hypothesis as given the finite value hypothesis. All of the probabilistic force in the Confirmation Argument comes from its expecting all humans to share *equal* and extreme value given the infinite value hypothesis, but the hyperreal infinite value hypothesis doesn't expect equal value in ${}^*\mathbb{R}$. Thus, Bailey and Rasmussen's Confirmation argument works against the hyperreal infinite value hypothesis just as much as against the finite value hypothesis, assuming we accept its premise that all humans have equal value.

Furthermore, the Instrumental Differences argument works against our hypothesis as well. For if overall value is the sum of instrumental value and final value, and overall value of different people is equal and final value of different people is equal, then instrumental value is equal, since hyperreal arithmetic works just like real arithmetic, and unlike extended real arithmetic where we have $x + y = x + z$ despite y and z being different, as long as $x = +\infty$ and neither y nor z is $-\infty$.

This puts us in a dilemma: either we take $\overline{\mathbb{R}}$ as our domain for modeling values and end up with unsettling problems like Kantian General and Kantian Lifeguard; or we take something like ${}^*\mathbb{R}$ as our domain for modeling values, in which case we have to claim that both of Bailey and Rasmussen's arguments for the infinite value of persons are unsuccessful. The latter option is, we think, far better.

If we take the hyperreals as our domain, we can still address the Value Puzzle. We can capture the idea that all persons possess dignity: a value that lets them stand on relatively equal footing with any other person, as in Kant (1958). But we need to modify one of our two assumptions: we need to substitute *relatively* equal final value for *strictly* equal final value: any two persons' final values differ only infinitesimally *relative* to their own final values. That is, suppose a person a 's final value is some infinite amount α and person

b 's final value is β . We propose that $\alpha - \beta$ will be an infinitesimal fraction of α , i.e., $(\alpha - \beta)/\alpha$ will be infinitesimal (which is equivalent to saying that $(\alpha - \beta)/\beta$ is infinitesimal¹⁸).

The final value of person b may be far greater than or far less than person a 's final value, but for both persons, whatever the difference may be, that difference is only infinitesimal compared to their *own* final value. Even if one person is a head of state and the other a hospice patient, both can know that their own final value, what makes them both valuable *qua* person, is far more significant than their difference in final value, thus preserving the intuition that there is a unique dignity of persons. So, qualifying the equality of persons' final value allows us to capture at least some of our intuitions about the value of persons.

Moreover, the relatively equal final value view still allows us to capture our intuitions regarding sums of multiple persons' final value. We can explain why the death of more people likely results in the loss of greater final value, why there is something valuable in having greater numbers of people voting in elections, why it is better for 10 people to exist instead of only 5, and so on. More people means more final value; fewer people likely means less final value.¹⁹

We can give a new two step argument for the view that the final value of

¹⁸Suppose $\epsilon = (\alpha - \beta)/\alpha$ is infinitesimal. Then $\beta/\alpha = 1 - \epsilon$ cannot be infinitesimal. But $(\alpha - \beta)/\beta = \epsilon\alpha/\beta$. This will be infinitesimal unless perhaps α/β is infinite, which can only happen if β/α is infinitesimal. So, if $(\alpha - \beta)/\alpha$ is infinitesimal, so is $(\alpha - \beta)/\beta$. And by the same argument, if the latter is infinitesimal, so is the former.

¹⁹Specifically, if $\alpha_1, \dots, \alpha_n$ and β_1, \dots, β_m are respectively the final values of the members of a group of n persons and of a larger group of m persons (i.e., $m > n$), then the relative equality thesis plus the infinity thesis implies there is a positive infinite value v such that $(\alpha_i - v)/v = \gamma_i$ and $(\beta_i - v)/v = \delta_i$ are all infinitesimal (just let $v = \alpha_1$, and note that the relative difference between the first member of the first group and anybody else is infinitesimal). Thus, $\alpha_1 + \dots + \alpha_n = (\gamma_1 + \dots + \gamma_n)v + nv$ and $\beta_1 + \dots + \beta_m = (\delta_1 + \dots + \delta_n)v + mv$. Since the sum of a finite number of infinitesimals is infinitesimal, the differences due to the γ_i and δ_i are swamped by the at least unit difference between m and n , and hence the value of the larger group is greater.

persons is infinite. First, the claim that the differences in final value between persons are relatively infinitesimal nicely captures our intuitions about human equality after allowing for the obvious fact that Mother Teresa had more final value than Stalin. But now observe that the difference in final value between Mother Teresa and Stalin was large—and in particular not absolutely infinitesimal. Now, however, it is an easy mathematical fact that if the difference between two quantities is relatively but not absolutely infinitesimal, the two quantities must be infinite. For if α is finite and $\epsilon = (\alpha - \beta)/\alpha$ is infinitesimal, then $\alpha\epsilon$ will be infinitesimal (the product of a finite and an infinitesimal value is infinitesimal), and since $\alpha - \beta = \alpha\epsilon$, the difference between α and β will be absolutely infinitesimal. Now we earlier saw that $(\alpha - \beta)/\alpha$ is infinitesimal if and only if $(\alpha - \beta)/\beta$ is infinitesimal. So the above argument with α and β swapped shows that β is finite and ϵ is infinitesimal, the difference between α and β is infinitesimal. Hence, the only way there can be a non-infinitesimal absolute difference, given a infinitesimal relatively difference, is if both values are infinite.

We end this section by discussing a question that Bailey and Rasmussen do not appear troubled by. Apart from a specifying context, to say that a chair has value ten is as nonsensical as saying that Usain Bolt's height is 6.4. Heights and values need to have units. This remains true even when a value or height is infinite. If we measure Usain Bolt's height in infifeet, which are some specific infinitesimal fraction of a foot, then Bolt is infinitely tall. And similarly if we measure values in infidollars—where an infidollar is some specific infinitesimal fraction of a dollar—then a chair has infinite value.

Thus, simply saying that humans have infinite value has to be elliptical for a more precise statement involving some specific kind of unit of value. If the value of human persons is measured in millibidens, where a biden is the value of Joe Biden, then trivially at least one human being has a finite value, namely Joe Biden has value 1000. If the value of human persons is measured in infibidens,

though, then it will be similarly trivial that at least one human being has infinite value.

So what does the claim that human persons have infinite value mean? We suggest the relevant units are ones on which where the ordinary value-infused daily events of human life—having lunch, talking to a friend or spraining an ankle—have non-infinitesimal finite value. The conversion rates between lunches, friendly talks, ankle-sprainings (or, for that matter, dollars) are all finite and non-infinitesimal, and so if a human person has infinite value with respect to one of these units, the person has it with respect to all the others.

We can put the claim of infinite value as follows: A human person has a value greater than any finite number of ordinary events. In a numerical (real or non-Archimedean) setting this basically means that if x is the value of a person and y is the value of some ordinary event such that $y \neq 0$, then the ratio $x/|y|$ is infinite.

5 Cardinal numbers

As we noted, Bailey and Rasmussen do not explicitly endorse an extended-real model. What we take as defintory of the extended-real model is that there is only one positive infinity, and Bailey and Rasmussen's arguments as officially formulated do presuppose that.

However, they also note that one might consider infinite cardinal numbers as a model of the infinite value of persons (pp.8-9). Here is a thought that supports this suggestion. If the value of persons were something finite, then because there are many values similar to any given finite value, it would be very unlikely that all human persons would have *exactly* the same value. For instance, if one person has value 100, why couldn't another have value 101 and another maybe only value 75? If there is only one positive infinity, and all

persons have infinite value, then it is guaranteed that all human persons have the same value.

But now consider the cardinal numbers. Considered as infinities, they are very far apart. If A and B are infinite cardinals with $A < B$, then B is not merely a little bigger than A , but infinitely many times bigger: indeed (assuming the Axiom of Choice) B is bigger than infinitely many copies of A , since B is bigger than $A \times A$, as $A \times A$ has the same infinite cardinality (here is where Choice is used) as A . If we model the values of persons on cardinal infinities, we would expect all human persons to have the same value. For while it may not be crazy to suppose *some* differences in value between human persons, it would be *prima facie* quite surprising if one human person were *infinitely* many times more valuable than another—and yet that would be the case if they had different cardinal infinities as their values.

However, this broad spacing of cardinal numbers makes the use of cardinal infinities be subject to exactly the same objection as the extended real model is. Adding any finite number of copies of a cardinal infinity yields the same cardinal infinity (assuming the Axiom of Choice), so moving to cardinal infinities does not help with Kantian General. Nor does it help with probabilistic reasoning. For presumably the reason we prefer certainty of a life to a chance of $1/2$ of a life, say, is because we think that half of the value of a life is less than the value of a life. But if B is an infinite cardinal and $B = 2 \times A$, then $B = A$ (again using the Axiom of Choice), so half of B is still B .

More generally, the Bailey and Rasmussen argument works only when infinities are guaranteed to be very far apart. The only models we know where this is true are ones where there is only one infinity, or where a bigger infinity is infinitely many times bigger than a smaller one, and neither kind of model fits well with Kantian General and probabilistic reasoning.

6 Alternatives

While the critique of Bailey and Rasmussen’s model of human value appears conclusive, developing an adequate model is no simple task. We suggest a non-Archimedean ordered field, because this is particularly familiar to philosophers. But of course, there are other well-established ways of modeling values on offer. For infinite values, one might use Paul Bartha’s relative expected utility (RUT)²⁰ or some other qualitative utility theory, or a lexicographic ordinal series like that of Russell (2020). And we should also consider the tried-and-true real-valued model of classical expected utility theory.

Regardless of the system, there are a number of desiderata that an ideal modeling system for the value of persons should be able to satisfy:

- (1) Prefer larger probability of obtaining some value v over smaller probability of obtaining v
- (2) Prefer saving larger numbers of equally valued persons over smaller numbers of persons
- (3) Allow for significant variation in instrumental value
- (4) Have either no variation in final value or in some appropriate sense relatively infinitesimal variation in final value
- (5) Have a preferred “privileged” infinite final value.

We do not claim that this list is exhaustive. However, it gives us a good starting point for how we might assess the relative strengths and weaknesses of different models. Desiderata (1) and (2) are very plausible practical desiderata: when choosing between two options of equal value but one of which has a greater chance of obtaining, one should choose the option with the greater chance of

²⁰See Bartha (2007) and Barth and DesRoches (2017) for more on RUT.

obtaining, even if the options have infinite value, and when choosing between a greater number of equally positively valuable options and a smaller number, we should go for the greater.

Desideratum (3) says that one should acknowledge that there is a difference between the amount of instrumental value between some people. Desideratum (4) says that the final value of different people is the same (or nearly the same), which captures deeply seated egalitarian intuitions, at least to a very close approximation. And Desideratum(5) says that there should be a privileged infinite value over other candidates for infinite value. Desideratum (5) also specifies that there is a non-arbitrary choice of what infinite final value human persons have rather than a multiplicity of candidate values.

Notice that (1) and (2) are practical desiderata. They deal with how we ought to make decisions and act in the world. Next, (3), (4), and (5) are theoretical desiderata. They deal with having an elegant, non-arbitrary model of values.

While an ideal model would satisfy each of these desiderata, we think that, since the first two actually affect how we make decisions, a model that can satisfy (1) and (2) would be better than rival models that cannot, even if the rival models perform better on (3), (4), and (5). The rest of this essay compares the how well different models satisfy these desiderata.

The extended real model of Bailey and Rasmussen fails severely in regard to (1) and (2). It does succeed in (3), but perhaps not quite as well as we would like. The $+\infty$ instrumental value of a person who saves one or more lives beats that of a person whose effects on the world are finite, which in turn beats the value $-\infty$ of a person who destroys lives. But there is no variation between the person who saves one and who saves a million, or destroys one and destroys a million. And we have a technical difficulty in dealing with a person who saves one and destroys ten: $\infty - \infty$ is undefined on the model.

On the other hand, the model succeeds perfectly in regard to (4) and (5), including the no-variation version of (4). We thus have a clear failure in two of the practical desiderata and one of the three theoretical desiderata, but success in the other two theoretical ones.

A cardinal number model is even worse than the extended-real model: it fails equally with regard to (1) and (2),²¹ likely does no better than the extended real model with respect to (3), seems to succeed with regard to (4) at least as restricted to human persons (maybe there can be non-human persons that are infinitely more valuable than us), but fails with respect to (5), since there are now infinitely many choices for the infinite values of persons.

The classical real-valued model succeeds perfectly with regard to (1), (2) and (3). But the only way the model can allow for at-most infinitesimal variation in final value between persons is by supposing that there is *no* variation in final value between persons. However the complete lack of variation when the final value is finite and hence seemingly mundane appears implausible, since the mundane differences between people would seem to ground some differences in value, unless the mundane is overshadowed by something supramundane, as on infinity-based models. So while the model can be made to accept (4), that version of the model is less plausible. And, finally, (5) does not appear to be satisfied: there isn't a clearly privileged finite value for persons to have.

The non-Archimedean ordered field approach succeeds perfectly with regard to (1), (2) and (3). It allows the infinitesimal variation version of (4), but fails with regard to (5): there are infinitely many different non-Archimedean ordered fields, and within each one there are infinitely many possible choices for the value of persons. However, (4) and (5) are merely theoretical desiderata.

Russell's ordinal sequence model works as follows. A value v is identified with a sequence of real numbers $\{v_\alpha\}_{\alpha < \beta}$ in $[-1, 1]$, where the index α runs

²¹We already saw this in the case of (1) and (2).

over ordinals, and β is an ordinal that may be different in the case of different values. The values are compared lexicographically, so that differences with respect to lower indices are more important. One can define expected values for finite lotteries component-wise, treating sequences as being all zeroes past the end of their indices if we like.²² For instance, in a simple case—sufficient for modeling the infinite value of people—if the sequences all have length two²³, then we can essentially define $E[(X, Y)] = (E[X], E[Y])$. And after defining expected values, we will have (1) if lives have positive value.

Because the sequences are required to be bounded by -1 and 1 , we cannot just add values arbitrarily, since then we could exceed the bounds, and we cannot say that n lives each of which has value v have value nv where $(nv)_\alpha = n(v_\alpha)$. However, we can still aggregate monotonically, so that the value of n lives is greater than the value of m lives whenever $n > m$, and we have hope of getting (2). We feel that although we get (2), we have here a weakness: the value of saving n lives should intuitively scale linearly in n , but this is not so on the Russell view. Nonetheless, Russell can allow for approximately linear scaling for small n .

There is no special difficulty about (3).

We can get (4) in a variety of ways. Variation in values at larger indices α is lexicographically less important, and consequently when two values agree at smaller indices, and are non-zero for at least one of these indices but disagree at larger indices, e.g., $(0, 0.2, -0.1)$ and $(0, 0.2, -0.2)$, we can say that they differ “relatively infinitesimally.” Thus, to get (4), we can suppose that there is an ordinal sequence $z = (z_i)_{i < \beta_0}$ of numbers in $[-1, 1]$, with the first non-zero number in the sequence being positive, such that every person’s value is

²²More precisely, if we have a lottery with probability p_i of outcome $\{v_{i,\alpha}\}_{\alpha < \beta_i}$ for $i = 1, \dots, n$, we can first stipulate that $v_{i,\alpha} = 0$ for $\alpha \geq \beta_i$, and then say that the expected value of the lottery is $\{\sum_{i=1}^n p_i v_{i,\alpha}\}_{\alpha < \cup_i \beta_i}$.

²³We are grateful to the referee for suggesting considering length-two sequences.

a sequence of length at least β_0 , and matches the sequence z for indices less than β_0 , but then the values may differ at indices β_0 and greater. The values of persons thus differ relatively infinitesimally.

Finally, we do not have (5). There are infinitely many choices of the initial sequence z that different human values agree on. Moreover, the theory has additional degrees of freedom in the choice of aggregation function.

Because of the aggregation issue, we think a more general non-Archimedean approach is preferable to Russell's. But in any case, we can also see Russell's approach as a special case of the non-Archimedean ordered field approach. For one example of a non-Archimedean ordered field is the surreals, and we can embed Russell's values in the surreals. For the ordinals embed in the surreals, and then we can identify an ordinal-indexed sequence of reals $(v_\alpha)_{\alpha < \beta}$ with $\sum_{\alpha < \beta} v_\alpha \varepsilon^\alpha$ for some positive infinitesimal ε .

Bartha's relative utility model also allows one to capture the intuition of infinite final value. This model is qualitative in nature: instead of assigning numerical or quasi-numerical values to outcomes, it has a comparison relation \lesssim on outcomes that satisfies all the axioms of the von Neumann–Morgenstern (1947) representation theorem except for continuity. Given that the comparison of values in the non-Archimedean ordered field approach satisfies the same axioms, it follows that Bartha's model can accommodate anything the non-Archimedean ordered field model can, assuming we replace outcomes with items of value/disvalue. In particular, there is no difficulty about (1)–(3).

The easiest way to get (4) is to suppose equal final value of persons. But we can also define relatively infinitesimal variation between items, at least if we suppose that there is a zero-value item (“nothingness”?) Q . For then we can say that items A and B of positive value, i.e., such that $Q < A$ and $Q < B$, differ relatively infinitesimally just in case:

(a) $[pA, (1 - p)Q] < A$ whenever $p < 1$, and

(a) $[pB, (1 - p)Q] < B$ whenever $p < 1$,

where $[pX, (1 - p)Y]$ is a situation with probability p of X and $1 - p$ of Y . If A and B have infinite value, then this can be done in an even simpler way: we can replace Q with any item of finite value (say, a nice lunch), and don't need to suppose a special "zero value" item.²⁴

Furthermore, we do better with respect to (5) than on the non-Archimedean or ordinal-series models: we do not need to assign specific quasi-numerical values, but simply need comparisons.

However, we have a serious independent problem with the Bartha approach as applied to the value of persons. When we talk of the final value of a person, we are talking of the value of a substance, not of an outcome, and so we need to replace or supplement Bartha's outcomes with substances. However, the Barthan story also requires the formation of items like $[pX, (1 - p)Y]$ for any items X and Y and any real p between 0 and 1. But then we are comparing the values of entries in disparate ontological categories: situations and substances. There is no such substance as a 67% chance of Biden and a 33% chance of Charlesmagne, and yet Bartha's story requires an item like $[0.33 \cdot \text{Charlesmagne}, 0.67 \cdot \text{Biden}]$ if it includes the items Charlesmagne and Biden.

We might try to squish substance value into situation value by talking of the value of the situation where Biden exists and the value of the situation where Charlesmagne exists. However, there is no such thing as *the* situation where Biden exists. There are, instead, infinitely many situations where Biden exists (one where he is a president, one where he is a famous physicist, etc.), and Biden's value is just a part of the value of each of those situations. We might

²⁴We are grateful to a suggestion from the referee that has simplified our definition and allowed for the removal of Q in the special case of A and B with infinite value.

try to identify the value of Biden as the difference in value between otherwise relevantly similar situations in one of which Biden exists and in the other of which Biden does not. However, the Barthan approach does not have a way of representing the differences between values. On the other hand, if we had a classical or non-Archimedean field approach, we would have some hope of starting with outcome value, and then use subtraction to calculate the value of specific parts of the outcome.

Finally a situation of a 67% chance of Biden and a 33% chance of Charlesmagne may well be impossible given essentiality of origins, since Charlesmagne is likely to be an ancestor of Biden, and hence it is impossible to have Biden exist in the absence of Charlesmagne.

Bartha's approach, thus, is specifically tied to the utilities of outcomes, like situations or worlds, in a way that makes it unsuitable for accounting for the value of persons. Our account starts with values of entities, but then has a parallel problem in giving a precise account of utilities of situations. Here there is a major ethical divide: do we take the value of persons and other entities first, or do we evaluate the values of outcomes and lotteries first? I.e., do we proceed patient-first or agent-first? While we acknowledge the difficulty in accounting for the utilities of situations, we think that the idea of starting with human value has significant ethical plausibility: we can think of much of our ethics as a matter of appropriately responding to this value. However, this is an area of big-picture disagreement.

Furthermore, it's worth noting that because Bartha's qualitative preferences satisfy the the axioms of the von Neumann–Morgenstern representation theorem minus continuity, they too can be represented by a hyperreal-valued utility function U such that $A \lesssim B$ if and only if $U(A) \leq U(B)$ (Skala 1974). Thus while Bartha's approach does have less arbitrary choice than ours, all the other theoretical benefits of it can be had by us. Furthermore, given that the most

obvious way to extract values of entities from values of outcomes is by subtraction of the value of an outcome without the entity from an outcome with the entity, it seems that some such quasi-numerical representation is likely to be needed to get values of entities out of Bartha's account.

Some of the same points are likely to apply to any other qualitative approach based around a subset of the von Neumann–Morgenstern axioms.

Conclusion

We have very good reason to prefer using a non-Archimedean ordered field like that of the hyperreals rather than the extended reals to model values. If we use the extended reals, we run the risk of falling into some heinous problems. But if we use the hyperreals, Bailey and Rasmussen's arguments that solve the Value Puzzle must be rejected work. We have offered an account of infinite human final value that works using hyperreals or other non-Archimedean ordered fields, and we have given a framework in which we can understand what it means to say that a person has infinite value. The framework is not the only one possible that achieves the same desiderata. Russell's ordinal series approach would, for instance, work just as well. But that approach can be seen as a special case of the one we are considered.²⁵

²⁵We are grateful to [...including anonymous referee...] for [...].

References

- Anderson, M. (2022). “The Imago Dei and the Infinite Value of Human Life.” In (eds.) O’Mathúna, D.P., Just, B.A., Jeremiah, W. and Eppinette, F.M., *Bioethics in Real Life: Lessons We’re Learning from COVID-19*. Deerfield, IL: The Center for Bioethics & Human Dignity: 74-94.
- Arntzenius, F., Elga., A, and Hawthorne, J. (2004). “Bayesianism, Infinite Decisions, and Binding.” *Mind*, 113 (450): 251–83.
- Bailey, A., and Rasmussen, J. (2020). “How Valuable Could a Person Be?” *Philosophy and Phenomenological Research*, (00):1-14.
- Bartha, P. (2007). “Taking Stock of Infinite Value: Pascal’s Wager and Relative Utilities.” *Synthese*, 154(1):5-52.
- Bartha, P., Barker, J., and Hájek, A. (2014). “Satan, St. Peter and St. Petersburg.” *Synthese*, 191 (3):629-660.
- Bartha, P., and DesRoches, C. T. (2017). “The Relatively Infinite Value of the Environment.” *Australasian Journal of Philosophy*, 95(2): 328-53.
- Beckstead, N., and Thomas, T. (2020). “A Paradox for Tiny Probabilities and Enormous Values” Global Priorities Institute, Working Paper, No. 10-2020.
- Bostrom, N. (2011). “Infinite Ethics.” *Analysis and Metaphysics*, 10:9-59.
- Colyvan, M., Justus, J., and Regan, H. (2010). “The Natural Environment is Valuable but not Infinitely Valuable.” *Conservation Letters*, 3:224-228.
- Easwaran, K. (2014). “Regularity and hyperreal credences.” *Philosophical Review*, 123: 1–41.
- Fishburn, P. C. (1971). “A study of lexicographic expected utility.” *Management Science*, 17: 672-678.
- Frakt, A. (2020, May 11). “Putting a Dollar Value on Life? Governments Al-

- ready Do.” *The New York Times*. Accessed May 2022. <https://www.nytimes.com/2020/05/11/upshot/virus-price-human-life.html>
- Hájek, A. (2012). “Blaise and Bayes.” *Probability in the Philosophy of Religion*, (eds. Chandler, Jake and Harrison, Victoria) Oxford University Press.
- (2003). “Waging War on Pascal’s Wager.” *Philosophical Review*, 112: 27–56.
- Hausner, M. (1954). “Multidimensional Utilities.” In (Eds.) Coombs, C.H., and R.L. Davis, *Decision Processes*, New York: Wiley.
- Jakobovits, I. (1979). “Medical Experimentation on Humans in Jewish Law.” In (Eds.) Bleich, J.D., and Fred Rosner *Jewish Bioethics*, Brooklyn, NY: Hebrew Publishing Company: 377-83.
- Kagan, S. (1998). “Rethinking Intrinsic Value.” *The Journal of Ethics*, 2: 277–297
- Kant, Immanuel (1958). *Groundwork of the Metaphysics of Morals*. Trans. H.J Paton. London and New York: Hutchinson’s University Library.
- Keisler, H. J.(2021). *Elementary Calculus: An Infinitesimal Approach*. Dover Publications.
- Knuth, D. (1974). *Surreal Numbers: How Two Ex-students Turned on to Pure Mathematics and Found Total Happiness*. Addison-Wesley.
- Korsgaard, C. M. (1983). “Two Distinctions in Goodness.” *Philosophical Review*, 92: 169–95.
- Pedersen, A. P. (2014). “Comparative Expectations.” *Studia Logica*, 102(4): 811–848.
- Pruss, A. (2018). “Underdetermination of infinitesimal probabilities.” *Synthese*, 198 (1):777–799.

- Rabinowicz, W., & Rønnow-Rasmussen, T. (1999). “A Distinction in Value: Intrinsic and for its Own Sake.” *Proceedings of the Aristotelian Society*, 100, 1: 33–51.
- Robinson, A. (1966). *Non-Standard Analysis*. Princeton University Press.
- Russell, J. (2020). “Non-Archimedean Preferences Over Countable Lotteries.” *Journal of Mathematical Economics*, 88: 180–186.
- Skala, H. J. (1974). “Nonstandard Utilities and the Foundation of Game Theory.” *International Journal of Game Theory*, 3: 67–81.
- Vallentyne, P., and Kagan, S. (1997). “Infinite Value and Finitely Additive Value Theory.” *The Journal of Philosophy*, 94 (1): 5–26
- Von Neumann, J., and Morgenstern, O. (1947). *Theory of Games and Economic Behavior*, 2nd ed. Princeton: Princeton University Press.
- Wojtyła, K. (2013). *Love and Responsibility*. Trans. Ignatik, Gregorz. Pauline Books and Media.