

Measuring Happiness and Suffering

SIMON KNUTSSON

Foundational Research Institute

simon.knutsson@foundational-research.org

May 2016

Abstract

Is the balance of happiness versus suffering in the future net positive or net negative (in expectation)? Is the aggregate happiness and suffering in a group of wild or domesticated non-human animals positive or negative?¹ For such questions to have factual answers that are free from value judgements, happiness and suffering would need to be objectively measurable to a very high degree. That is to say, they would need to be objectively measurable on a specific kind of scale.² In principle, we would need to establish a plausible way to add and subtract magnitudes of happiness and suffering across individuals and get a sum; a way that does not involve a value judgement on the part of the person doing the arithmetic.³ However, such a degree of measurability is widely (although not universally) rejected.

Contents

1	Introduction	2
2	Terminology	2
3	Measurement on one interpersonally additive ratio scale is required	4
4	The required degree of measurability is widely rejected	4
5	Why doubt such a high degree of measurability?	5

¹I assume for the sake of this essay that both happiness and suffering exist, although some would dispute that happiness (or pleasure) exists: “One [view] holds that pleasure is nothing real but only the absence of pain; if we think it has positive qualities, we are only being fooled by the transition from a greater to a lesser pain. This view has been asserted by Plato, Epicurus, and Arthur Schopenhauer.” Hurka (2010, 200).

²As explained below, by ‘objectively measurable to a very high degree’ I do not mean that the measurement needs to be precise; it can be approximate. The degree of measurability needs to be high in the sense that happiness and suffering need to be measurable on the same interpersonally additive ratio scale. If happiness and suffering were measurable on such as scale, one would be able to say that Ann’s happiness is about twice the magnitude of Ben’s suffering, and that the total amount of happiness minus suffering among Ann and Ben is strictly positive.

³Strictly speaking, the aggregation procedure need not be simple addition and subtraction, it could, for example, involve ratios (Brülde 1998, 70-71), but simple addition is probably the most common idea and ratios would require the same degree of measurability.

6	How a solution to the measurement challenge could look	7
7	Even if happiness and suffering <i>are</i> sufficiently measurable in principle, the practical application is still a challenge	9
8	If happiness and suffering are not sufficiently measurable in principle	11
9	A call for clarity	13
	References	13
	Appendix: Objective, non-evaluative measurement of happiness and suffering	15
	Further Reading	17

1 Introduction

If happiness and suffering are not sufficiently measurable, which they do not seem to be, then we are plausibly left with making subjective value judgements about whether, say, a future scenario or a wild animal population with both happy and suffering individuals contains more happiness than suffering. Such claims could then be understood as being colored by the value judgement that the scenario or population is overall good.⁴ Similarly, many disagreements about whether the future, a scenario, a life or a group of individuals contains a net positive or net negative balance of happiness versus suffering are probably not mainly disagreements about facts but rather instances of people making different value judgements, for example about which lives are worth living.⁵

⁴Overall good, at least considering only the happiness and suffering and disregarding how it is distributed.

⁵To clarify, I do not mean that our assessments of amounts or magnitudes of happiness versus suffering are biased. I mean that for such assessments, at least in the kinds of cases that I am focusing on (such as future scenarios and populations of wild animals), there appears to be no plausible factually correct answer (and hence no such answer that our assessments can be biased compared to). I am saying that such assessments appear to be partly value judgements, for example about the goodness and badness of the mental states in question.

⁶This idea about the quantification of pleasure dates back at least to Edgeworth (1881).

Given how controversial it is to claim that happiness and suffering can objectively be added together across individuals, and how unclear it is how one should understand talk of ‘amounts,’ ‘units’ and ‘magnitudes’ of happiness versus suffering, I suggest that one should clarify what one means when talking about such things. For example, does one mean that the magnitude of, say, happiness should be conceived of as the number of smallest perceivable increments of pleasure?⁶ Or is one making up magnitudes and corresponding numbers as an expression of one’s values?

2 Terminology

By ‘happiness,’ I mean roughly to feel good, and by ‘suffering,’ I mean roughly to feel bad. So when I say ‘happiness’ and ‘suffering’ one could roughly replace those words with ‘plea-

sure' and 'displeasure.'⁷ Happiness and suffering are some of the things that have been conceived of as **finally** good and bad for an individual (i.e. good and bad for their own sake, as ends rather than as means). Other things that have been said to be finally good or bad for an individual include preference satisfaction and preference frustration, achievement, friendship, and so on. In other words, these things, and many more, have been claimed to be the facts that constitute an individual's 'well-being.' How high or low someone's level of well-being is refers to how good or bad life is for the individual who leads it. If one says that someone's well-being is high, one says that the individual is well-off, and to say that someone's well-being has increased is to say that life has become better or less bad for the individual who leads it. That makes 'well-being' an evaluative term. In contrast, terms such as 'happiness,' 'suffering,' 'pleasure,' 'preference satisfaction,' and so on are commonly assumed to be factual terms, and I will take that assumption as my starting point.⁸ The idea is

that well-being depends on, is grounded on, or consists in, various facts, such as happiness and suffering. Although I focus on happiness and suffering, what I say in this essay generally applies to whatever is taken to constitute well-being, whether it is happiness and suffering, or other facts such as preference satisfaction, achievement, or friendship.

I will also speak of measuring 'utility.' There are at least two different conceptions of utility, but I will just speak of 'utility' for simplicity, and will treat the measurement of either of the two senses of utility as being part of the same broad topic as measuring happiness and suffering.⁹ The first sense of utility is roughly the sense of the early utilitarian Jeremy Bentham. Utility in this sense is understood in terms of pleasure, displeasure, pain, happiness, unhappiness, suffering, or the like. This is roughly how a hedonist, for example a hedonistic utilitarian, thinks about utility.¹⁰ The second sense is von Neumann-Morgenstern utility, which is understood in terms of preferences and choice. An individual is said to get higher utility from

⁷However, an affective happiness theory of well-being could say that happiness and to feel good should not be understood in terms of pleasure, but rather some other affective state, such as mood, but we can disregard that complication.

⁸Deciding on a definition or conception of 'happiness,' 'suffering,' 'pleasure,' and so on plausibly involves evaluative or moral judgements, if one wants a conception that is as evaluatively or morally relevant as possible, or that makes a theory of well-being or morality as plausible as possible. The terms are still factual in the sense that, given a definition or conception, they refer to various facts, such as various mental states. See Brülde (1998, 13–16, 95–96).

⁹My distinction between the two conceptions of utility roughly corresponds to the distinction Kahneman et al (1997, 375) makes between experienced utility and decision utility: "The concept of utility has carried two quite different meanings in its long history. As Bentham [1789] used it, utility refers to pleasure and pain, the 'sovereign masters' that 'point out what we ought to do, as well as determine what we shall do.' This usage was retained in the economic writings of the nineteenth century, but it was gradually replaced by a different interpretation [Stigler 1950]. In current economics and in decision theory, the utility of outcomes and attributes refers to their weight in decisions: utility is inferred from observed choices and is in turn used to explain these choices. To distinguish the two notions, we shall refer to Bentham's concept as *experienced utility* and to the modern usage as *decision utility*. With few exceptions, experienced utility is essentially ignored in modern economic discourse."

¹⁰More specifically, a hedonist thinks about utility in terms of pleasure and displeasure (or pleasant and unpleasant experiences).

¹¹See Binmore (2009). I think that for our purposes we can treat measuring utility in either of the two senses as belonging to the same broad topic in part because I refer to measuring utility mainly in the section "The required degree of measurability is widely rejected." In that section, I refer to economists' (and others') skepticism about measuring utility interpersonally on one ratio scale, and to my understanding, this is the view among these skeptics whether one speaks of utility in the sense of von Neumann-Morgenstern utility or in the sense of

something if she prefers or chooses it. This is how economists typically think of utility.¹¹ A preferentialist, such as a preference utilitarian, would have something similar to this second sense of utility in mind when saying that an individual's well-being consists in satisfied and frustrated preferences; that is, when saying roughly that it is good for an individual to get what she wants and bad for her to get what she has an aversion to.

3 Measurement on one interpersonally additive ratio scale is required

What degree of measurement is needed to make claims about whether the balance of happiness and suffering among several individuals is net positive or net negative?¹² Such claims require one of the highest degrees of measurability: that happiness and suffering can be measured interpersonally (across different individuals) on the same additive ratio scale. This means roughly that one would be able to say that Ann's suffering is twice the magnitude of the happiness of each of Ben, Chloe, and Diana, and that we can add all these together and get a strictly positive sum, for example $-2 + 1 + 1 + 1 = 1$.¹³ This would be a high degree of measurability in the sense that happiness and suffering would be highly quantifiable. I do not mean that the measurement needs to be precise; it can be approxi-

mate of the kind that 'Ann is suffering roughly ten times as much as Ben is happy.'

In other words, the kinds of claims that we are concerned with about balances of happiness versus suffering require that happiness and suffering can be quantified similarly to concepts like length or weight.¹⁴ Length is a typical example of an attribute that is generally considered to be measurable on an additive ratio scale: two meters are twice as long as one meter, two plus one meters equals three meters, and so on. Mass is another common example: one can meaningfully speak of the total weight of several things or persons, such as Ann and Ben.¹⁵

4 The required degree of measurability is widely rejected

That happiness and suffering are measurable, in principle, to the extent that is required to talk about the net balance among several individuals is highly controversial and widely rejected. That is, it is controversial that they are (in principle) measurable to the required degree in an objective, non-arbitrary, scientific way that does not involve value judgements on the part of the person doing the measurement. One could say "I assign number -10 to Ann's suffering and $+5$ to Ben's happiness, and then I add them together. These numbers are intertwined with my values, and others might assign different numbers depending on their val-

happiness, suffering, pleasure, displeasure, and so on. Assuming that this is correct, we can simply just speak of 'utility' as an unspecific term in that section.

¹² That is, assuming that some of the individuals are happy and that some suffer. If, for instance, all individuals under consideration would suffer, it would suffice to make weaker measurability claims. But for the cases that we are concerned with (far-future scenarios and groups of non-human animals), we can assume that some of the individuals are happy and that some suffer.

¹³ Representing these individuals' suffering and happiness with e.g. the following numbers would work just as well: $-4 + 2 + 2 + 2 = 2$. The requirement is that the ratios are preserved (hence the term 'ratio scale'), and with these new numbers, Ann's suffering is still twice the magnitude of the happiness of each of Ben, Chloe, and Diana.

¹⁴ However, there are complicating dissimilarities: for instance, length and weight are normally only zero or positive but for happiness and suffering, we are talking about a scale with a zero, a positive and a negative side.

¹⁵ For more on measurement and different scales, see Roberts (1985), Krantz et al. (1971), and Carlson (forthcoming).

ues." Although that might be the best we can do, it does not count as measurement in the objective sense that we are concerned with in this section (see the appendix for more on objective measurement versus measurement colored by value judgements).

Economist Yew-Kwang Ng says that the following statement is representative of the typical economics textbook view: "Today, no one really believes that we can actually measure utils."¹⁶ He continues that the probably most widely used textbook says that "economists today generally reject the notion of a cardinal, measurable, utility."¹⁷ For the kinds of claims discussed above, we would need utilities to be measurable on a strong kind of interpersonal cardinal scale (an interpersonally additive ratio scale), which is widely rejected by economists.¹⁸ Ng adds that this skepticism about measurability is also "very common" among "sociologists and psychologists who study happiness."¹⁹

Some history is interesting here. The early utilitarians of the 19th century, such as John Stuart Mill and Henry Sidgwick, did not seem worried by interpersonal comparisons of utility.²⁰ But as Bergström (1982a) points out, "in the 20th century, things have changed a great deal. Now the dominant view — at least among economists — seems to be that interpersonal comparisons of utility are impossible or necessarily subjective and unscientific."²¹

Of course, there is not complete consensus

on the matter; some believe that happiness and suffering (or utility) can be measured to the extent required to talk about the net balance or amounts among several individuals. Philosophers generally seem to be somewhat more optimistic than economists.

5 Why doubt such a high degree of measurability?

Why would one doubt that happiness and suffering can be measured, in principle, on the same additive ratio scale? That is, why doubt that Ann's happiness can be said to be five times Ben's suffering and that we can add these magnitudes together into a sum?²² One can believe that because happiness and suffering are so different, it does not make sense to compare magnitudes of happiness and suffering on the same scale. To illustrate this point with a more clear case, we can consider other facts that might be finally good and bad for individuals, besides happiness and suffering. For example, one of the things that might be finally good for individuals is to have close personal relationships, and one that might be finally bad is to be deceived. These are plausibly not measurable on the same ratio scale. It hardly makes sense to say that the extent to which I have close personal relationships is about five times the extent to which I am being deceived. Similarly, one can say that even if an instance of suffering can be twice the mag-

¹⁶Ng (2015, 10), quoting Miller (2011, 436-37).

¹⁷Ng (2015, 10), quoting Samuelson and Nordhaus (2010, 89).

¹⁸Another such statement about the views of economists can be found in Binmore (2009): "What do modern economists mean when they talk about units of utility? How can such utils be compared? It is widely thought that the answer to the second question is that utils assigned to different individuals cannot sensibly be compared at all." Note, as I mentioned above, that economists usually think of utility in terms of preferences and choice and not happiness, pleasure, suffering, and the like.

¹⁹Ng (2015, 11).

²⁰According to Bergström (1982, 284).

²¹Page 284.

²²There are formal requirements that are said to be necessary for relations such as 'happier than' to allow additive measurement (on a ratio scale). The requirements include transitivity, monotonicity, and so on. For more depth on such requirements, see Roberts (1985), Krantz et al. (1971), and Carlson (forthcoming). In this section, I just bring up some informal reasons to doubt such a high degree of measurability.

nitude of another instance of suffering, and an instance of happiness can be twice the magnitude of another instance of happiness, an instance of happiness still cannot be twice the magnitude of an instance of suffering because happiness and suffering are not measurable on the same scale.

Another case is if we assume that what is good for an individual is preference satisfaction and what is bad is aversion fulfillment (aversion fulfillment is when an event or situation occurs that an individual did not want to happen).²³ One understanding of an individual's preferences and aversions is in terms of the choices that the individual makes or would make. A way to say how much stronger a preference is than another is to consider what the individual chooses or would choose. An argument against comparing and adding strengths of preferences and aversions across individuals (or within a life over time) is that the strengths of an individual's preferences and aversions are to be understood in terms of that individual's choices at a point in time. But across individuals (and within a life over time), the argument

goes, there is no choosing super subject that spans across individuals (or across time) that can be used to compare preferences and their strengths.²⁴

The point of this section is to exemplify, for presentation purposes, why one would doubt that happiness and suffering (and utility) are measurable to the required degree (on one interpersonally additive ratio scale); I do not mean that this is the end of the discussion for and against measurability. My take on the topic is that for our kind of cases (far-future scenarios, large populations, etc.) factual judgements about amounts of happiness and suffering typically cannot be distinguished from value judgments.²⁵ I agree with Griffin (1979) that there are situations in which the judgement about amounts of happiness and suffering is "not independent of, not distinguishable from" one's preferences.²⁶ In other words, I essentially sympathize with the "objector" in the following passage from Hurka (2010):

Moore and Mayerfeld assume, as Bentham

²³Brülde (1998, 36).

²⁴My formulation of this argument is roughly taken from Brülde (2003, 156). See, for example, Binmore (2009) for a nice overview of interpersonal utility comparisons. Binmore believes that they can be made.

²⁵Parfit (2011) also speaks of comparisons of amounts of happiness and suffering as normative, although I am not sure what exactly he means: "According to Hedonistic Utilitarians, the past has been in itself good if there has been, in the lives of all conscious beings, a positive total sum of happiness minus suffering. To explain such claims, we must explain the sense in which some amount of happiness might be *greater* than some amount of suffering. This sense of 'greater' is normative, in a way that is often overlooked. We can first consider brief pleasures and pains. Some pleasure would be in itself greater than some pain if the nature of these two experiences would on balance give us reasons to choose to have both rather than neither. It might, for example, be worth enduring intense cold on some mountain's summit for the sake of seeing a sublime view. When some pleasure is in this sense *greater* than some pain, these experiences would together give us a positive sum of pleasure minus pain, or as we can say more briefly a *net* sum of pleasure. Such claims need more explanation, and should be qualified in various ways. The relative value of such experiences would be very imprecise. Despite these facts, we can often truly believe that some pleasure is greater than some pain. In some longer part of our life, or our life as a whole, we might similarly have a positive sum of happiness minus suffering. That would be true if it would be worth enduring this suffering for the sake of this happiness." (Page 610).

²⁶Griffin (1979) says, "Perhaps in many situations the judgement whether amounts of happiness and unhappiness are *equal* is not independent of, not distinguishable from, one's being indifferent to a situation in which they are combined. Similarly, perhaps in these cases the judgement whether a certain amount of happiness is just greater than a certain amount of unhappiness is not independent of the judgement that the point has just been reached where one would prefer a situation in which they are combined" (page 53). Further down on that page, he omits 'perhaps' and implies that he is suggesting that there *are* such situations where judgements about relative happiness and suffering are not independent from one's preferences.

and Sidgwick also do, that we can compare the intensities of pleasures and pains independently of assessing their values. But an objector may challenge this assumption, saying the claim that a pain is more intense than some pleasure merely says the pain is more evil than the pleasure is good, without pointing to some independent psychological fact that makes it so. There is no such fact, she may argue, and thus no possibility of nonevaluative comparisons of pleasures and pains.²⁷

I try to avoid saying that the magnitude of an instance of happiness is equal to, greater than, or smaller than an instance of suffering, but if I would say that, I would think of my statement as being partly a value judgement. It would be colored by how bad I believe that the suffering is compared to how good I believe that the happiness is.

6 How a solution to the measurement challenge could look

Assume that one disagrees with me and believes that happiness and suffering are, in principle, objectively measurable on one interpersonally additive ratio scale. One would then need to provide an account of what a unit and a magnitude of happiness and suffering is.²⁸ What is this unit of happiness and suffering we are to add and subtract? When

is a degree of happiness of the same magnitude as a degree of suffering? One such account is provided by classical hedonistic utilitarian [Torbjörn Tännsjö](#); he argues that pleasure and displeasure can, in principle, be measured to the degree that is required for our purposes. He roughly endorses the idea from Edgeworth (1881) that the magnitude or intensity of pleasure and displeasure could be measured in principle by the number of smallest perceivable increments. One could then, in principle, count such increments of pleasure and displeasure, add and subtract them across individuals, and end up with a net sum that is positive, zero, or negative. Actually, Tännsjö develops Edgeworth's idea and says the unit is not exactly the smallest perceivable increment, but rather that "sub-noticeable," or more accurately, "only indirectly perceivable" increments should count.²⁹ However, for simplicity, I will just speak of perceivable increments, and I will call this kind of conception of units of pleasure and displeasure the 'Edgeworth-Tännsjö' view, although that doesn't do justice to Tännsjö's view.

If this Edgeworth-Tännsjö conception of the magnitude of pleasure and displeasure were plausible, it would, in principle, allow objective non-evaluative measurement of these magnitudes to the degree required for our purposes³⁰ (assuming increments of pleasure and displeasure are units on the *same* scale, which

²⁷Hurka (2010, 201–202).

²⁸And one would need an account of why there is a zero-level and where it is located.

²⁹It is not crucial for our purposes, but the following passage from Tännsjö (2000) explains what he means by 'sub-noticeable,' i.e. 'only indirectly perceivable' increments: "How are we then to measure sub-noticeable differences? If we stick to the interpretation where the relation of having the same value is transitive, we may use the method of introducing 'extra' states for comparison as a method of arriving (in principle) at a measure of differences in well-being that are, in the situation, not noticeable (which are, that is, *indirectly* noticeable). Consider again the situation with three possible states for a person, A, B, and C. When comparing A and B the person notices no difference, when comparing B and C the person notices no difference, but when comparing A and C the person notices a difference: C is in a noticeable way worse (directly worse) than A. We conclude, then, that B is indirectly worse than A but indirectly better than C" (page 87). That is, according to Tännsjö, the difference (increment) between A and B would not be noticeable, but still real and morally relevant despite being "only indirectly noticeable" or "only indirectly perceivable." The same goes for the difference between B and C.

³⁰That is, measurement on the same interpersonally additive ratio scale.

requires another argument). Although it may be difficult or impossible in practice to *know* how many increments of displeasure a particular, say, fish can perceive, it would allow measurement *in principle* as long as she can notice increments. It might be impossible for us to know how many, but that is an epistemological or practical problem; not a problem for measurability *in principle*. Furthermore, the measurement would be *objective* (not involving value judgements). Although deciding whether one finds this conception of pleasure and displeasure plausible likely involves some kind of value judgement,³¹ magnitudes of pleasure and displeasure would be objectively measurable to the degree that we need for our purposes.

One problem is that most people aren't likely to accept this conception of magnitudes of pleasure and displeasure as the number of smallest perceivable increments. Myself, I find it implausible. That Adam can perceive twice as many increments of displeasure as Bob does not mean, in my view, that Adam suffers twice as severely, in the evaluatively relevant sense. Rather, Adam is perhaps more discriminate and can distinguish among different experi-

ences in a more fine-grained way. Whether he suffers more, and especially twice as much, would need to be based on some other consideration.³² This problem becomes especially important when comparing across species and even with possible future minds that are very different from ours. In particular, if a mind were designed to be able to notice tiny increments,³³ I would be reluctant to say that such a mind would feel vastly greater magnitudes of pleasure and displeasure (in the evaluatively relevant senses of pleasure and displeasure).

I bring up this Edgeworth-Tännsjö conception of the magnitude of pleasure and displeasure as perceivable increments to illustrate the form that a solution could have. I mention an objection to it to show what a challenge it is to measure happiness and suffering (even in principle) to the degree that one can speak of net positive or negative balances among several individuals (and even in a single life over time). Even an alternative (the Edgeworth-Tännsjö idea about perceivable increments) that has been considered to be the most promising faces a strong objection.³⁴

Other ideas that aim to enable interpersonal comparisons of utility or suffering and hap-

³¹As Tännsjö (2000, 82) says, "Why ought we to conceive of well-being in terms of hedons? Why ought we to say that, when there exist more than one way of changing the situation from something similar to one state to something similar to another state, the maximum number of such changes is what matters? There exists a discussion about these questions but, in the final analysis, it seems to be the case that the hedonistic utilitarian option for hedons, rather than some other possible idea, such as the stipulation that the worst possible state, or best possible state, for each sentient being, is of equal worth, rests on a normative judgement."

³²Others have made similar objections before. For example, Bergström (1982, 308) discusses it and refers to Arrow (1963, 117-8) for a cogent expression of it. Tännsjö's reply to the objection is the following: "it could be objected, of course, that it is unfair that those who have fine discriminatory capacities should count for more, in our moral calculations, than those with less fine discriminatory capacities. This has been considered a *reductio* of hedonistic utilitarianism. However, this seems to be a consequence that classical hedonistic utilitarians are prepared willingly to acknowledge. An objection to this stipulation seems to be an objection to hedonistic utilitarianism as such. Furthermore, it should be noticed that, if sub-noticeable differences [see note X for an explanation of sub-noticeable differences] of well-being are taken into full account, then it is not quite true that classical hedonistic utilitarianism is biased against people with poor discriminatory capacities. It is rather biased against people with a poor *sensibility*. This is different and, indeed, not at all objectionable if, in the final analysis, pleasure and displeasure are what really matter." Tännsjö (2000, 89-90).

³³Or "only indirectly" notice them.

³⁴It has been considered the most promising, at least by Bergström (1982a) and Tännsjö (2000). Bergström (1982, 310) says that "from a strictly utilitarian point of view, the method of minimal units seems to be *the* most acceptable method."

piness have been proposed. That literature is vast, and I don't mean to give a comprehensive overview here nor to establish that happiness and suffering cannot be measured to the required degree.³⁵ I will, however, discuss briefly a particular family of methods that Bergström (1982a) calls 'science fiction methods.'

'Science fiction methods' would say that happiness and suffering can in principle be measured with the help of, say, neurophysiology and biochemistry, or a better understanding of the physical processes that happiness and suffering are identical with, reducible to.³⁶ One can imagine a 'hedonometer' that measures the relevant neurological or physical events.³⁷ However, these methods would not provide the answer; they presuppose it.³⁸ That is, we first need to know which neurophysiological or physical processes we should think of as constituting twice as much suffering as the happiness that another process (in another individual) constitutes.

In other words, the hedonometers would need to be calibrated using as input a conception of which processes correspond to or make up relative magnitudes of happiness and suffering in different individuals.³⁹ 'Science fiction methods' and hedonometers are perhaps best seen as complements or tools given a conception of happiness and suffering (or e.g. pleasure and displeasure), such as the Edgeworth-Tännsjö conception. As such, science fiction methods could perhaps help us extrapolate or improve the internal coherence of our views

on relative magnitudes of happiness and suffering, but they do take such views as input. In conclusion, these methods do not solve the challenge of measuring and comparing happiness and suffering in a way that avoids subjective value judgments or arbitrariness, which is the fundamental challenge.

7 Even if happiness and suffering *are* sufficiently measurable in principle, the practical application is still a challenge

Say that the Edgeworth-Tännsjö view is right; that it is plausible to conceive of and measure, in principle, pleasure and displeasure in terms of smallest perceivable (or only indirectly perceivable) increments. Alternately, consider that some other plausible method exists for measuring, in principle, happiness, suffering, preference satisfaction, aversion fulfillment, or the like to the degree required to add such amounts across individuals. What would the practical implications be? Could we apply those methods in real life when we contemplate the balance of, say, happiness and suffering in a future scenario or a group of wild non-human animals?

One issue is how one should understand the idea of smallest 'noticeable' or 'perceivable' changes in pleasure and displeasure. What first comes to mind may be that one can ask, in an experiment, humans about the increments that they notice. Nevertheless, the practical problems seem huge when it comes to, for ex-

³⁵For example, Ng (2015) argues, based on considerations of evolutionary fitness, that happiness can be measured on one additive interpersonally comparable ratio scale, where (net) happiness is "the excess of positive affective feelings (including pleasure) over negative ones" (page 8). See Binmore (2009) and Bergström (1982a) for good overviews on the topic of measurement and interpersonal comparisons. Two related dissertations are Kloksiem (2009), which focuses on the measurability of pleasure and pain, and Rossi (2009), which deals with interpersonal comparisons of utility understood in terms of preferences. The two dissertations provide plenty of further references. Oosterheld (2015) formalizes the notions of preference satisfaction and preference frustration.

³⁶Or [supervenient](#) upon.

³⁷Edgeworth (1881, 101) proposes the idea of an "hedonimeter." Colander (2007) provides a nice perspective on Edgeworth and hedonometers (hedonimeters). Bergström (1982a) uses the term 'H-meter.'

³⁸As Bergström (1982, 305) points out.

³⁹Bengt Brülde phrased this point similarly in conversation.

ample, how many increments of displeasure a human perceives if he is covered in petrol and set on fire. Even more challenging are cases like an antelope who is eaten alive, or a fish dying from pressure changes. Not to mention future minds very different from ours that may exist thousands of years from now that may be organic or perhaps digital (or non-organic). How many increments will they perceive?

Measuring preferences, attitudes, or desires does seem substantially more feasible in practice. Feldman (2010) argues that it is finally good for an individual "to be pleased about something," which, loosely speaking, is "to want it and to believe that it is happening."⁴⁰ An individual can be pleased and displeased about several things to different degrees simultaneously.⁴¹ Feldman says that being pleased and displeased can, at least in principle, be measured on the same additive ratio scale.⁴²

‘...if you are happy, precisely how happy are you?’... in my view, there is an answer. The answer is this: consider all the states of affairs that Tom is occurrently [right now, actively] intrinsically attitudinally pleased (or displeased) about, at the moment. For each, consider the extent to

which he is intrinsically pleased (or displeased) about it. Add up these numbers, making sure to make use of negative numbers for all the things Tom is intrinsically displeased about at the time. This sum represents how happy Tom is at the moment.⁴³

Even if we could conceive of a method to, in principle, at a single point in time, measure the strength of Tom’s desires and aversions regarding various states of affairs and somehow come up with how pleased or displeased he is on the whole at that point in time, this method seems completely impracticable for the scenarios and situations at hand.⁴⁴

Although I have only discussed measurability in practice from the perspective of two theories of well-being, Tännsjö’s and Feldman’s, it seems intractable regardless of which of the common theories that one considers.⁴⁵ So what are we to do in real life? (Still granting that happiness and suffering are, in principle, measurable to the required degree.) One could try to do what Neurath (1912) says could be one option, namely, to "try to imagine simultaneously the pleasure of all those concerned, in order to arrive at a decision about the total plea-

⁴⁰Page 116. However, this is not the complete description of what it is to be pleased, according to Feldman. For example, he adds, "It seems to me that the combination of belief and desire is not sufficient for attitudinal pleasure. It is possible to want a thing and to think that you’re getting it without having a favorable emotional outlook on that thing" (page 116).

⁴¹"Surely it is possible for a person to be pleased about several things at once. Tom might be pleased to be living in Massachusetts, pleased to be sitting on his deck, and pleased to be smelling the roses, all at one time. And, equally surely, a person can be pleased about several things at a time, and simultaneously displeased about several other things. Tom is still displeased about his latest speeding ticket." Feldman (2010, 117).

⁴²Although he also says that he "would not want to take the numbers very seriously." Feldman (2010, 112). He believes that the ratio scale is comparable and additive at least over time in a life (pages 120–121).

⁴³Feldman (2010, 118).

⁴⁴For example, Oesterheld’s (2015) formal mechanism of extracting and comparing preferences involves complete knowledge of the physical configuration of an agent, perfect ability to simulate a large set of possible scenarios and is probably not even computable in theory, i.e. with arbitrarily powerful computers.

⁴⁵Brülde (2003, ch. 6) goes through the standard types of theories of well-being (hedonism, preferentialism, objective list, and a happiness theory) and concludes that regardless of theory, well-being is not measurable, not even in principle, to anything close to the degree that would be required for us to sum, for example, happiness and suffering across individuals.

⁴⁶Neurath says that this could be one of the available options for a statesman "faced with the choice of enhancing the happiness of either one group of men or another at the expense of the remaining ones." Page 119.

sure."⁴⁶ It is a good idea to thoroughly contemplate and to try to imagine happiness and suffering in outcomes, but for our purposes, this gets us essentially nowhere. Another option is to make 'total comparisons' of situations or scenarios and decide which of them one prefers, without calculating the balance of happiness versus suffering.⁴⁷

If this is roughly what measurement in practice of balances of happiness and suffering comes down to for our cases of future scenarios and non-human animal populations with large numbers of individuals, that would be a meager conclusion. However, it doesn't look much better as far as I can tell. I guess if one believes that happiness and suffering can be measured to the required degree in principle, one could perhaps at least start from that method and see if one can get anywhere in practice, and the specific practical method would seemingly depend on how happiness and suffering are measurable in principle. This looks intractable for our cases, which are some of the most challenging measurement-wise because they involve different species and kinds of minds, many individuals, and far-future scenarios.

We can distinguish between two questions, both still assuming that happiness and suffering are sufficiently measurable in principle: (1) Can we conceive of, or come up with, a method that in practice would give us the tools and information that would be substantially better than the options discussed above to "imagine simultaneously the pleasure [and displeasure] of all those concerned" or to make 'total comparisons' without calculating balances of hap-

piness versus suffering; a method that would work satisfactorily for our purposes in practice? (2) Whether or not such a method could be found in the future, to what extent can we *right now* in practice measure the relevant amounts of happiness and suffering that would enable us to, in a reasonably objective, value-free, non-arbitrary way, contemplate the balance of happiness and suffering in far-future scenarios and wild and domesticated animal populations? Both questions are interesting, but (2) is most relevant for our current ability to measure, which is very important for our current talk about balances of happiness and suffering. The answer to (2) looks bleak. The answer to (1) may be up for discussion based on how exactly happiness and suffering is supposed to be measurable in principle.

In the next section, I will consider the practical implications of the idea that happiness and suffering are not sufficiently measurable even in principle. That discussion overlaps with, and works as a continuation of, the discussion here because the practical implications are similar whether measurement cannot be made in principle or whether it could be made in principle but is completely impracticable.

8 If happiness and suffering are not sufficiently measurable in principle

What if, as seems to be the case, happiness and suffering are not sufficiently measurable, not in principle and thus also not in practice? That is, not sufficiently measurable for us to, in an objective and non-arbitrary way, add the

⁴⁷I here refer to the other option that Neurath (1912, 119) brings up: to choose the constellation that one "prefers, without being able to base this preference on a calculation of the pleasure sums." Similarly, utilitarian J. J. C. Smart talks about "comparison of total situations" instead of a summation of happiness. Smart (1973, 32) says, "In order to help someone decide whether to do A or to do B we could say to him: 'Envisage the total consequences of A, and think them over carefully and imaginatively. Now envisage the total consequences of B, and think them over carefully. As a benevolent and humane man, would you prefer the consequences of A or those of B? That is, we are asking for a comparison of one (present and future) *total* situation with another (present and future) *total* situation. So far we are not asking for a *summation* or *calculation* of pleasures or happiness. We are asking only for a comparison of total situations.'" Bergström (1982a) uses the term 'total comparisons' and refers to Neurath and Smart.

happiness and suffering among different individuals and get a sum that is either positive, zero, or negative? Some have worked on ordinal versions of utilitarianism that make weaker measurability assumptions. An ordinal scale only ranks the things in question in terms of more, less, higher, lower, and the like. For example, ordinal measurability allows the ranking that an instance of happiness is more happiness than another, but an ordinal scale does not say *how much* more.

If happiness and suffering were only measurable on an ordinal scale, we would not be able to say that an instance of happiness is twice the magnitude of another instance of happiness, and we would not be able to talk of sums of happiness and suffering. One formulation of ordinal utilitarianism results in a kind of maximum rule, according to which "we are to look to the worst possible outcomes in evaluating actions and institutions" and "to ameliorate the condition of the worst-off moment of phenomenal experience in the world."⁴⁸ Another treatment of ordinal utilitarianism seems to imply that the action or policy recommendation in our cases would be 'not applicable.'⁴⁹

Anyhow, say that we still want to say something about the overall goodness or badness of outcomes based on the "amounts" of happiness and suffering in them, despite that, strictly speaking, suffering and happiness (taken to be factual terms) cannot be measured in a way that allows addition and subtraction. Plausibly, such judgements can still be epistemically better or worse. It is an epistemic virtue to gather relevant information, think through the scenarios that one makes judgements about, and to be internally coherent.

Still, such epistemic constraints probably leave room for a wide range of views on "how much" happiness and suffering an outcome contains.

Assuming that one has informed oneself and has thought through the outcomes that one is considering, how should one come to an overall conclusion about the "amounts" or "balances" of happiness and suffering? One could, as mentioned above, make 'total comparisons'; that is, to consider outcomes as wholes and decide which one prefers, without calculating sums of happiness and suffering. As Bergström says, this "method does not seem utilitarian at all. We have absolutely no reason to believe that the situation which an agent *prefers* is identical with the situation which contains a maximum of happiness. The agent's preferences are morally irrelevant according to utilitarianism."⁵⁰ Of course, if happiness and suffering are not measurable in principle in amounts, there is no maximum amount of net happiness for one's value judgement (or preference) to hit or miss. Therefore, this informed overall judgement (total comparison) method seems fine, assuming happiness and suffering are not sufficiently measurable.

An alternative is to think *as if* happiness and suffering were sufficiently measurable, even though one believes that they are not *objectively* measurable to that degree. That is, to inform oneself about the facts and then, in combination with one's value judgements, make up numbers representing "amounts" and "units," and then do the arithmetic. This is what Brian Tomasik does.⁵¹ He *says* "It's my party, and I'll use the numbers I want to." With this method, if an instance of suffering

⁴⁸Mendola (1990, 86).

⁴⁹Hardin (2009). For example, "An inherent implication of an ordinal utilitarianism, as opposed to a Benthamite cardinal utilitarianism, is that we must be more modest in our claims about the good we might do in the world. Often, our theory yields ambiguous implications or, essentially, no implications at all for what policy we should adopt" (p. 47).

⁵⁰Bergström (1982, 288).

⁵¹But see the appendix for more on this approach and in which senses it is objective and not.

is stipulated to have magnitude -100 , it is, at least partly, a value judgement whether a specific instance of happiness is assigned magnitude $+1$, $+10$, or $+100$. Is this method of making up numbers and adding them better than the 'total comparisons' (overall judgement) method? It is unclear what 'better' should mean here, especially if there is no fact of the matter about what the actual amounts are, since neither method is then more true than the other. At least making up numbers is perhaps more transparent and less conservative in that it can lead to conclusions that are surprising to oneself, and that can conflict with one's overall judgment, and then one would need to decide whether to adjust the numbers or one's overall judgement.⁵²

9 A call for clarity

I want to end with a call for more clarity in discussions about how much happiness versus suffering an outcome contains. It is too unclear to just talk of amounts, net balances, sums, etc. of pleasure, happiness, suffering, preference satisfaction, or the like. It would be better if the author states that she, for example, is talking about addition and subtraction of number of smallest perceivable increments or something else. Another way to be clear to a satisfactory degree is to say, "I am essentially making up numbers that express my subjective values." A reason why I earlier brought up the Edgeworth-Tännsjö conception of the magnitude of pleasure and displeasure as perceivable increments is to highlight it as a good example of clarity. Tännsjö admirably clarifies how he thinks we should conceive of units and magnitudes of pleasure and displeasure and hence what we are supposed to add and sub-

tract.⁵³

References

- Arrow, Kenneth J. (1963). *Social Choice and Individual Values*. 2nd ed. New York: Wiley.
- Bergström, Lars. (1982a). "Interpersonal Utility Comparisons." *Grazer Philosophische Studien* 16/17, 282–312.
- . (1982b). "Om Välfärdspolitikens Målsättning." *Filosofisk Tidskrift* 3, 16–34.
- Binmore, Ken. (2009). "Interpersonal Comparison of Utility." In *The Oxford Handbook of Philosophy of Economics*, edited by Don Ross and Harold Kincaid, 540–59. Oxford/New York: Oxford University Press.
- Brülde, Bengt. (1998). "The Human Good." PhD diss., University of Gothenburg.
- . (2003). *Teorier om Livskvalitet*. Lund: Studentlitteratur.
- Carlson, Erik. (forthcoming). "Value Theory (Axiology)." In *Handbook of Formal Philosophy*, edited by Sven Ove Hansson and Vincent Hendricks. Springer. Especially sec. 1 on measuring goodness and badness.
- Colander, David. (2007). "Retrospectives: Edgeworth's Hedonimeter and the Quest to Measure Utility." *The Journal of Economic Perspectives* 21 (2), 215–26.
- Edgeworth, Francis Ysidro. (1881). *Mathematical Psychics*. Kegan Paul, London.

⁵²This method of making up numbers can also have the benefit of leading to interesting discussions, as Bengt Brülde pointed out in conversation.

⁵³I am very grateful to Caspar Oesterheld, and especially to Bengt Brülde and Brian Tomasik for helpful discussions, pointers, and feedback. I am also very grateful to Gordon Hanzmann-Johnson and Robin Raven who improved the language of this text. In writing this essay, I made heavy use of Bergström (1982a), Brülde (1998), Brülde (2003) and works by Erik Carlson on measurement, e.g. Carlson (forthcoming).

- Feldman, Fred. (2010). *What is This Thing Called Happiness?* Oxford: Oxford University Press.
- Griffin, James. (1979). "Is Unhappiness Morally More Important Than Happiness?" *The Philosophical Quarterly* 29 (114), 47–55.
- Hardin, Russell. (2009). "Utilitarian Aggregation." *Social Philosophy and Policy* 26 (1), 30–47.
- Hurka, Thomas. (2010). "Asymmetries in Value." *Noûs* 44 (2), 199–223.
- Kahneman, Daniel, Peter P. Wakker, and Rakesh Sarin. (1997). "Back to Bentham? Explorations of Experienced Utility." *The Quarterly Journal of Economics* 112 (2), 375–405.
- Klockslem, Justin Allen. (2009). "On the Measurability of Pleasure and Pain." PhD diss., University of Massachusetts Amherst.
- Krantz, David, Duncan Luce, Patrick Suppes, and Amos Tversky. (1971). *Foundations of Measurement, Vol. I: Additive and Polynomial Representations*. New York : Academic Press, Cop. Especially chapters 1–3. Reprinted 2006.
- Mendola, Joseph. (1990). "An Ordinal Modification of Classical Utilitarianism." *Erkenntnis* 33 (1), 73–88.
- Miller, Roger Leroy. (2011). *Economics Today*. 16th ed. New York: Pearson.
- Ng, Yew-Kwang. (2015) "Some Conceptual and Methodological Issues on Happiness: Lessons from Evolutionary Biology." *The Singapore Economic Review* 60 (4).
- Oesterheld, Caspar. (2015). "Formalizing Preference Utilitarianism in Physical World Models." *Synthese*, 1–13.
- Ofstad, Harald. (1982). *Ansvar och Handling: Inledning till Moralfilosofiska Problem*. Translated to Swedish by Marlene Öhberg. Stockholm: Prisma.
- Parfit, Derek. *On What Matters*, vol. 2. Edited and introduced by Samuel Scheffler. Oxford: Oxford University Press.
- Portmore, Douglas. W. (2005). "Combining Teleological Ethics with Evaluator Relativism: A Promising Result." *Pacific Philosophical Quarterly* 86 (1), 95–113.
- Roberts, Fred S. (1985). *Measurement Theory: With Applications to Decisionmaking, Utility, and the Social Sciences*. Cambridge: Cambridge University Press. Reprinted 2009.
- Rossi, Mauro. (2009). "Interpersonal Comparisons of Utility: The Epistemological Problem." PhD diss., London School of Economics and Political Science.
- Samuelson, Paul A., and William D. Nordhaus. (2010). *Economics*, 19th ed. New York: McGraw-Hill/Irvin.
- Smart, J. J. C. (1973). In J. J. C. Smart and Bernard Williams, *Utilitarianism: For and Against*. Cambridge University Press.
- Tännsjö, Torbjörn. (1998). *Hedonistic Utilitarianism*. Edinburgh: Edinburgh University Press.
- . (2000). "The Least Sub-Noticeable Difference." In *Imperceptible Harms and Benefits*, edited by Michael J. Almeida, 75–93. Springer Netherlands.
- . (2015). "Utilitarianism or Prioritarianism?" *Utilitas* 27 (2), 240–50.

Appendix: Objective, non-evaluative measurement of happiness and suffering

A key distinction running through this essay is that between (a) in principle measuring happiness and suffering in an objective, non-evaluative, non-arbitrary, empirical, factual way,⁵⁴ and (b) that the person making statements about magnitudes of happiness versus suffering is, at least partly, colored by subjective value judgements. Why does this distinction matter?

The distinction seems to matter to utilitarians. Utilitarianism has been more focused than other moral doctrines on providing an *empirical* criterion for the rightness of acts.⁵⁵ A criterion of rightness states necessary and sufficient conditions for an act to be morally right. For example, classical hedonistic utilitarian Tännsjö formulates a rough classical utilitarian criterion of rightness as follows, and in the same passage, he expresses the desire for the criterion to be *empirical*:

According to the [classical utilitarian] criterion of rightness, a particular

(concrete) action is right if, and only if, there is no other way the agent who performed it could have acted, such that, had he or she done so, the world, on the whole, would have been better. If possible, what makes the world better or worse should be spelled out in purely empirical terms.⁵⁶

He also says that “classical hedonistic utilitarians have often prided themselves on the belief that their criterion of rightness is purely empirical.”⁵⁷ A historical criticism of many kinds of utilitarianism is that comparisons of magnitudes (intensities) of happiness and suffering are necessarily subjective, and not an objective weighing of facts.⁵⁸ I take it that it is acknowledged that it would be a problem for typical utilitarian calculations of happiness versus suffering if they were not in principle objectively measurable.⁵⁹

Why would one want, as Tännsjö says, to spell out in purely empirical terms what makes the world better or worse (such as happiness and suffering)? One reason why utilitarians such as Tännsjö want to is perhaps because otherwise utilitarianism

⁵⁴This would not involve value judgements on the part of the person making the statements about magnitudes of happiness versus suffering (once the definition or conception of happiness and suffering is decided upon).

⁵⁵According to Bergström (1982b, 27).

⁵⁶Tännsjö (1998, 65).

⁵⁷Tännsjö (1998, 63–64).

⁵⁸For example, Ofstad (1982) critically discusses fundamental ideas that are common to many forms of utilitarianism. He says, “The question then becomes: Can we compare the intensity of experiences? Before we take a position we have to be clear about that the question concerns whether comparisons can be made in an objective way, as when we with the help of a thermometer compare the degree of two liquids” (page 51, my translation from Swedish). He also says critically that in normal situations the “weighing” of values and their dimensions such as intensity and time of experiences against each other is “not a weighing but a value-determined judgement... Utilitarianism has not discussed this weighing process... If the utilitarian starts to dig into it, his simple” model collapses (pages 69–70, my translation from Swedish).

⁵⁹For example, Bergström (1982b) says, “of course if there is no rule for how different values or dimensions [e.g. intensity and duration of experiences] are to be ‘weighed’ against each other, then, to the corresponding degree, it opens up a space for arbitrariness and subjective variations. But why would it be impossible to provide such a rule?... It is precisely such a utilitarian rule that I have sketched above” (page 27, my translation from Swedish). The rule that he sketches is essentially the Edgeworth view described above that a unit of well-being is to be conceived of as the smallest perceivable difference in well-being.

would not be distinguishable from “rival” theories such as [prioritarianism](#) and [negative](#) utilitarianism in the simple, standard way.⁶⁰ A second related reason, which is probably the primary reason, could be that many prefer ethical theories to be formulated in a general, simple, impersonal, non-relative, or non-subjective way. For example, Bergström (1982b) seems to mean that without a purely factual conception of things such as happiness and suffering, there is room “for arbitrariness and subjective variations.”⁶¹ Similarly, economists who try to formulate purely empirical conceptions of utility that are comparable across individuals presumably do so because they want conceptions of utility to be scientific, elegant, systematically applicable to policymaking, useful in the social sciences, or the like. How does Tomasik’s approach relate to all this? As mentioned above, based on facts,⁶² reasons, and his values, he “makes up” numbers representing relative magnitudes of happiness and suffering depending on how bad he thinks that some specific suffering is compared to some specific happiness, and then he does

arithmetic with those numbers. For example, he might say that the suffering of dying in a specific way is 100 times as much suffering compared to the pleasantness of some specific pleasant experience. He might also say things such as “I think a human who gets burned moderately suffers about X times as much as an ant who is burned moderately, by which I mean that I care (or should care) X times as much about the human burning as the ant burning.”

Tomasik is a moral and evaluative antirealist, which means that he doesn’t believe there are moral or evaluative truths. I interpret Tomasik’s endeavour in ethics and value theory as figuring things out for himself and sharing his conclusions and reasons with others. So what does Tomasik prescribe? He does *not* prescribe that “all persons should bring about the best outcomes in terms of the balance of happiness versus suffering as *they themselves* judge such balances.”⁶³

Rather, he prescribes something like the following: “I try, and please join me if you want, to bring about the best outcomes in terms of the balance of happiness versus suffering as *I, Brian Tomasik*, judge such

⁶⁰Tännsjö (2015, 241) says the following when discussing prioritarianism versus utilitarianism: “Now, if we want to keep utilitarianism and prioritarianism apart, regardless of how else we conceive of well-being, we had better see to it that well-being is an empirical (non-normative) notion. Otherwise, it is likely that, if there are prioritarian weights in the first place, then we have already factored them in, when we assess how well off an individual is at a certain moment.” Similarly, Hurka (2010, 201–202) says, “To formulate a pairwise asymmetry about any pair of values we must be able to compare these values non-evaluatively, or identify instances of them as equal in a way that is neutral about their comparative worth. Only then can their relative values be a further issue.” The same is the case if we want to formulate a pairwise *symmetry* about any pair of values, such as pleasure and displeasure. If we cannot compare what has value, e.g. magnitudes of pleasure and displeasure, non-evaluatively, it is not straightforward to distinguish between non-negative and weak negative views.

⁶¹Page 27, my translation from Swedish.

⁶²Facts such as self-reports by those who have suffered severely, revealed preferences, neurophysiological data, personal experience with similar kinds of pain, etc.

⁶³For example, I take it that Tomasik would not endorse Portmore (2005)’s combination of teleological ethical theory with evaluator relativism, which Portmore describes as follows: “We obtain such a theory by applying the principle ‘act always so as to maximize value’ to evaluator relativism, where the value of certain states of affairs varies according to who the evaluator is. The resulting theory is one where agents ought always to bring about what is, from their own individual positions, the best available state of affairs.”(page 96).

balances.” Perhaps such an ethical position could be called ‘Brian utilitarianism.’ Some might dispute that it’s even a form of utilitarianism, but that would just be a question of labels, which is of minor importance.

Does Tomasik, like Tännsjö for instance, have an objective, empirical, factual conception of magnitudes of happiness versus suffering? Yes, seemingly, but in a different way. Tomasik’s conception is factual and objective in the sense that one could seemingly, in principle, specify factually what goes on in his brain when he judges some specific instance of suffering to have a magnitude (or intensity) that is X times that of a specific instance of happiness. Similarly to the discussion of hedonometers above, one could perhaps in principle objectively measure and specify Tomasik’s neural activity to determine which factual neural processes correspond to Tomasik making the judgment that someone else’s suffering is X times the magnitude of some other individual’s happiness (not his own). In contrast to the discussion of hedonometers above, we now have a factual calibration method available: what Tomasik says or thinks. That is, a specific factual neural activity is to count as stating that this suffering that some individual is experiencing is X times the magnitude of some other happiness if it is the neural activity that Tomasik has when he says or thinks the phrase ‘this suffering is X times the magnitude of that happiness.’ Or, more simply, as Tomasik says, “Just ask me. My answer will be objectively measurable.” This approach would probably not satisfy academic philosophers. What they would most likely want is a factual specification of how happy Chloe is and how much she is suffering in terms of *her* mental states,

experiences, choices, or the like, *not* in terms of the evaluator’s (e.g. Tomasik’s) neurological processes or statements. Couldn’t Tomasik factually specify which mental processes *in others* make up various magnitudes of happiness and suffering? For example, through a series of statements such as ‘burning to death is 3 times worse than being eaten alive,’ ‘one human matters 5,000 times as much as one ant,’ etc. Couldn’t he also, at a lower level of description, specify which physical or neural processes are to count as X times the suffering as other physical or neural processes, and so on. He could, in principle, do that, and it might be the best one can do.

Many philosophers would likely object and ask for the *justification* or basis for why, for example, being sick in way X is Y times as much suffering as being pleased in way Z. The philosophers would likely want a general reasonably simple and short rule for determining magnitudes of happiness versus suffering, such as the Edgeworth-Tännsjö method. Tomasik would likely answer, ‘I am not aware of a plausible, simple, general rule. That suffering that happens over there appears to me to be much worse than the goodness of that happiness over there, say 50 times as bad as the happiness is good (i.e. the suffering is 50 times the magnitude of the happiness). I can provide reasons for why I think so about this case and similar cases, and you might or might not find the reasons compelling. Ultimately, it’s up to each of us to decide for ourselves. I don’t have a better answer, and neither do you.’

Further Reading

Bergström, Lars. (1982a). “Interpersonal Utility Comparisons.” *Grazer Philosophische Studien* 16/17, 282–312.

- Binmore, Ken. (2009). "Interpersonal Comparison of Utility." In *The Oxford Handbook of Philosophy of Economics*, edited by Don Ross and Harold Kincaid, 540–59. Oxford/New York: Oxford University Press.
- Brülde, Bengt. (2003). *Teorier om Livskvalitet*. Lund: Studentlitteratur. Especially ch. 6 on measuring well-being.
- Carlson, Erik. (forthcoming). "Value Theory (Axiology)." In *Handbook of Formal Philosophy*, edited by Sven Ove Hansson and Vincent Hendricks. Springer. Especially sec. 1 on measuring goodness and badness.
- Krantz, David, Duncan Luce, Patrick Suppes, and Amos Tversky. (1971). *Foundations of Measurement, Vol. I: Additive and Polynomial Representations*. New York : Academic Press, Cop. Especially chapters 1–3. Reprinted 2006, seemingly without major revisions.
- Roberts, Fred S. (1985). *Measurement Theory: With Applications to Decisionmaking, Utility, and the Social Sciences*. Cambridge: Cambridge University Press. Reprinted 2009, seemingly without major revisions.