

Aggregation and Problems about Trade-offs Within Lives

In chapter four, I argued that we should reject Sidgwick’s conception of individual self-interest, because it assumes a simple additive aggregationist approach for assessing the overall value of a life. Specifically, I argued that in some cases anti-additive aggregationist principles are applicable within lives, as well as between lives, including analogues of the Minimize Great Additional Burdens View, the Consolidate Great Additional Benefits View, and the Second Standard View. However, as should be evident, like chapter two and three’s anti-additive aggregationist principles, chapter four’s analogues of those principles are incomplete. They are only applicable to, and generate rankings for, certain cases. In other cases, they are silent, and we must rely on other principles to rank alternative lives. As we will see, often the principles that are appropriate for ranking alternative lives require precisely the sort of trade-offs that are prohibited in those cases where the anti-additive principles apply. In this chapter, we will see that, as with rankings involving different lives, the relevance of different principles for ranking individual lives raises deep problems about aggregation and trade-offs within lives.

5.1 Two More Standard Views

In chapter two, I argued that most people accept:

The First Standard View—Trade-offs between Quality and Number Are Sometimes Desirable: in general, it is preferable if a larger number of people are benefited to a lesser degree, than if a smaller number of people are benefited to a greater degree, if the number of people who are thereby benefited is “sufficiently” greater, and if the differences in the initial situations of the people benefited and the degrees to which they are benefited are not “too” great.

Most people accept a similar view within lives. However, within lives the trade-offs can be put, more aptly, in terms of quality versus duration, rather than quality versus number. Moreover, for reasons of exposition, it is worth breaking the kind of view I have in mind into two separate principles, one where the quality involves benefits, and the other where it involves burdens. Specifically, I claim that most people accept both of the following:

The Fourth Standard View--Trade-offs between Benefits and Duration Are Sometimes Desirable Within Lives: in general, it is better receive a lesser benefit for a longer duration, than a greater benefit for a shorter duration, if the longer duration is “sufficiently” longer, and if the difference in the quality of the benefits is not “too” great.

The Fifth Standard View-- Trade-offs between Burdens and Duration Are Sometimes Desirable Within Lives: in general, it is better receive a greater burden for a shorter duration, than a lesser burden for a longer duration, if the longer duration is “sufficiently” longer, and if the difference in the quality of the burdens is not “too” great.

One could give countless examples of each of these. Here are three examples of the first. First, I must choose between two pleasurable experiences, a 10 minute back and neck massage followed by 10 minutes of “nothing,” or a 20 minute back massage. Assuming that neither has any long-term effects, that a back and neck massage feels better than a back massage alone, but that the latter is almost as good as the former, the second choice seems clearly better than the first. Second, in exchange for a week’s work, I can be paid \$100 for one week, or \$90 dollars for two weeks, but, let us suppose, I can’t save my pay from week to week. Assuming that the difference in the quality of the benefits that I can buy with \$90 dollars is not too much less than what I could buy with \$100 dollars, the second option seems clearly better than the first. Surely it is better to have slightly lower quality benefit for twice as long, than a slightly higher quality benefit for half as long. Third, suppose I could live a rich and rewarding life for 25 years,

followed by 75 years of suspended animation, or a life that was almost as rich and rewarding for fifty years, followed by 50 years of suspended animation. Assuming the difference in the overall quality of the life was not “too” great, the second seems better than the first. Of course, people will have different views about when the differences in lives would be “too” great. Some would happily forgo a larger house, or fancier cars, for twice as many years of life. Others would make trade-offs in their job satisfaction, health, or relationships. But virtually all would agree that some loss in the quality of one’s life, even in terms of things that matter, would be worth sufficient gain in the duration of one’s life, as long as the difference in quality between the two lives was not too great.

Consider the last example closer. I suspect that even those who most love their jobs, would readily accept a less satisfying one for an extra twenty-five years of pretty good life! And as much as people rightly value robust health, would it not be worth ongoing allergies, some loss of hearing, or perhaps a limp, for an extra twenty-five years of pretty good life? Likewise, as important as it is to have wonderful relationships with one’s friends and family, mightn’t it be worth a few less satisfying relationships—or even a few less friends or family members!—for an extra twenty five years of pretty good life. Not sure? Make it an extra fifty years of life, or seventy-five! And suppose it isn’t your best friend, or a member of your immediate family, but just a pretty good friend, and maybe a cousin, aunt, or uncle that you like, but don’t see all that much.

Does this sound crass, or cold-blooded? Perhaps it is, but it reflects an important truth. There are many sources of value in life, and as long as the quality of one’s life

remains positive, some losses in quality that are not “too” big, can be more than made up for, overall, by sufficient gains in the duration of one’s life

The preceding examples range from the trivial—trade-offs involving different massages for less than an hour—to the substantial—trade-offs involving important components of well being and significant differences in the length of one’s life.

Together, they help indicate the robustness of the Fourth Standard View. In a wide range of cases, involving alternatives for individuals, it would be better to receive a lesser benefit for a longer duration, than a greater benefit for a shorter duration, if the longer duration is “sufficiently” longer, and if the difference in the quality of the benefits is not “too” great.

Next, let me note three examples of the Fifth Standard View. Borrowing from alternatives presented in chapter two, I can be brief. Let us suppose a given case of warts is worse than a given case of acne, but not “too” much so. Then, other things equal, most people would prefer n weeks of acne to n weeks of warts, but they’d rather suffer a given duration of warts than a “sufficiently” longer duration of acne. How much longer the duration of acne would have to be would, of course, depend on the details of the alternatives. This need not concern us here. The important point is simply that if the negative impact of the acne was close “enough” to that of the warts, the warts would be better than the acne, if the acne lasted two, or three, or five times as long. Moreover, this would be true whether the bouts of skin ailments were measured in months, years, or decades. Thus, in general, it would be better to suffer warts for a month, year, or decade, than acne for two (or three, or five) months, years, or decades, respectively.

Similar claims would hold for trade-offs between broken arms and broken legs, or bouts of psychosis versus bouts of manic-depression. Even if, other things equal, a broken arm is worse than a broken leg, or psychosis is worse than manic-depression, duration involving trade-offs between such illnesses makes a difference. Thus, while final judgments would, of course, depend on the details of the cases, in general, it would be better to suffer a broken arm for a month, a year, or a decade, than a broken leg for two (three, or five) months, two (three, or five) years, or two (three, or five) decades, respectively. Similarly, for trade-offs between psychosis and manic-depression.

As with the examples illustrating the Fourth Standard View, the preceding examples range from the trivial—trade-offs involving warts versus acne—to the substantial—trade-offs involving psychosis versus manic-depression. They also range over different lengths of time, from months, to years, to decades. Correspondingly, they suggest that the Fifth Standard View is very robust. I conclude, therefore, that for a wide range of cases, involving alternatives for individuals, it would be better to receive a greater burden for a shorter duration, than a lesser burden for a longer duration, if the longer duration is “sufficiently” longer, and if the difference in the quality of the burdens is not “too” great.

5.2 Another Worry about Consistency

In section 2.3, we illustrated an inconsistency between the First Standard View, the Second Standard View, a plausible assumption, and the transitivity of the “all-things-considered better than” relation. As should be evident, similar worries arise about consistency for chapter four’s Third Standard View, section 5.1’s Fourth and Fifth Standard Views, a very plausible assumption, and the transitivity of the “all-things-

considered better than” relation. This should not be surprising. Problems of consistency can arise for comparing alternatives involving a single life for the very same reason that they can arise for comparing alternatives involving different lives; namely, that for some such comparisons trade-offs between quality and duration seem appropriate, but for others they do not. Thus, in some cases one employs an additive aggregationist approach when comparing alternatives, in others one does not.

Let us spell out the purported inconsistency further for two possible cases; one involving trade-offs between benefits and duration, the other involving trade-offs between burdens and duration. We’ll begin with the case involving benefits.

To start, let me explicate the plausible assumption that underlies the purported inconsistency. It should look familiar, as it is analogous to the assumption underlying section 2.3’s purported inconsistency. The assumption involves two elements. The first is just that there is, or a least could be, a spectrum of alternatives involving benefits that might accrue to an individual, such that at one end of the spectrum the alternative would involve the individual receiving a great benefit for a “sufficiently” long period of time, while at the other end the alternative would involve the individual receiving a minor benefit for a much longer period of time. The second is just that one could move from one end of the spectrum to the other in a finite series of small steps, such that the Fourth Standard View would apply to all pairwise comparisons involving alternatives that were “near” each other on the spectrum, while the Third Standard View would apply to all pairwise comparisons involving alternatives that were at opposite ends of the spectrum. Naturally, alternatives that were “near” each other on the spectrum might only be one small step, or perhaps a few small steps, apart, while alternatives at opposite ends of the

spectrum would be more than a few small steps, and perhaps, in fact, a great many small steps, apart.

So, for example, one might set a scale ranging from 1, representing the quality of an oyster-like life that is barely worth living, to 1000 representing the quality of life of the highest form of human existence. Then, one might believe that for differences in quality of life of, say, 3 or less, the Fourth Standard View would apply, while for differences in quality of life of, say, 900 or more the Third Standard View would apply. Intuitively, then, the idea is that the difference in the quality of lives between, say, 779 and 776, or 238 and 235, or 27 and 24, is “sufficiently” small that, in accordance with the Fourth Standard View, it would be better to live the lower quality life for n years, rather than the higher quality of life for m years followed by $n-m$ years of suspended animation, as long as n is “sufficiently” greater than m , so that the quality of life that is experienced lasts for a “sufficiently” longer duration in the former alternative than in the latter alternative. But, intuitively, the difference in the quality of lives between, say, 950 and 50 is so significant that, in accordance with the Third Standard View, no matter how long one lived, say n years, it would be better to live a life of level 950 for a significant period of time, say 100 years followed by $n-100$ years of suspended animation, than to live all n years of one’s life at level 50. But, of course, for familiar reasons, these two sets of views are incompatible with the transitivity of the “all-things-considered better than” relation, as the first view entails the rejection of the second if the “all-things-considered better than” relation is transitive.

It follows that we must reject the Fourth Standard View, the Third Standard View, the transitivity of the “all-things-considered better than” relation, or the plausible

assumption that there could be a spectrum of lives ranging from one involving an oyster-like quality that persisted for many years to one involving the highest form of human existence that persisted for a significant number of years, such that the Fourth Standard View would apply to lives near each other on the spectrum, while the Third Standard View would apply to lives at opposite ends of the spectrum. Giving up any of these views will not be easy.

Consider next the case of burdens. The Third Standard View was put in terms of benefits, but its anti-additive aggregative reasoning is also plausible regarding burdens. Indeed, it may be more plausible regarding burdens. Thus, the Third Standard View may be spelled out further as follows.

The Third Standard View—Even within Lives, Trade-offs between Quality and Duration Are Sometimes Undesirable Even When Vast Differences in Duration Are at Stake: if the differences in the quality of benefits is “sufficiently” large, it would be preferable for someone to receive the larger benefit, as long as it persisted for a “sufficiently” long duration, than for that person to receive the smaller benefit for any period of finite duration. Likewise, if the difference in the quality of burdens is “sufficiently” large, it would be preferable for someone to receive the smaller burden for any period of finite duration, than the larger burden for a “substantial” period of duration.

As should be apparent, the Third Standard View, so understood, is inconsistent with the Fifth Standard View, the transitivity of the “all-things-considered better than” relation, and a plausible assumption. Not surprisingly, the plausible assumption is analogous to the one concerning benefits, and it, too, has two components. The first is just that there is, or a least could be, a spectrum of alternatives involving burdens that might accrue to an individual, such that at one end of the spectrum the alternative would involve the individual receiving a great burden for a “sufficiently” long period of time, while at the other end the alternative would involve the individual receiving a minor

burden for a much longer period of time. The second is just that one could move from one end of the spectrum to the other in a finite series of small steps, such that the Fifth Standard View would apply to all pairwise comparisons involving alternatives that were “near” each other on the spectrum, while the Third Standard View would apply to all pairwise comparisons involving alternatives that were at opposite ends of the spectrum. Naturally, alternatives that were “near” each other on the spectrum might only be one small step, or perhaps a few small steps, apart, while alternatives at opposite ends of the spectrum would be more than a few small steps, and perhaps, in fact, a great many small steps, apart.

As in the case of benefits, at least one of the inconsistent positions must go. But here, as elsewhere, it is by no means evident which position(s) should be given up.

5.3 A Powerful Example—From Torture to Mosquito Bites

On reflection, I believe both sets of inconsistencies noted above are troubling. But they are not equally so. In truth, I believe that there are asymmetries in our thinking about benefits and burdens that may make the situation regarding burdens more problematic than the situation regarding benefits. Correspondingly, let me put the problems concerning benefits to the side, for a while, and focus on the more troubling problems concerning burdens.

The discussion in section 5.2 is fairly abstract. This is particularly true of the discussion regarding burdens, where I have basically presented a schema for recognizing the inconsistency of our views. To aid in the exploration of this topic, let us consider in detail a powerful, more concrete, example illustrating the kind of inconsistency identified in section 5.2.

The example I shall present is a variation of an example sent to me by Stuart Rachels. I have altered the details of Rachels's example in ways he objects to, and the discussion of it is my own. So, I shall refer to the following example as my own, and Rachels deserves no blame for it. Still, if the reader finds the following compelling, Rachels deserves much credit. Rachels offered his example as a counterexample to the transitivity of the betterness relation, and in an earlier paper of mine I followed him in this regard.¹ Indeed, at the time, I thought Rachels's argument was the simplest, most powerful, counterexample to the transitivity of the betterness relation that had yet been offered. But my aim here is not to argue for the intransitivity of the betterness relation, but rather to illustrate the incompatibility of four deeply plausible views.

I begin with a rough statement of the four views.

View One: for any unpleasant or "negative" experience, no matter what the intensity and duration of that experience, it would be better to have that experience than one that was only a "little" less intense but twice (or three or five times) as long.

View Two: there is, or could be, a spectrum of unpleasant or "negative" experiences ranging in intensity, for example, from extreme forms of torture to the mild discomfort of a mosquito bite, such that one could move from the harsh end of the spectrum to the mild end in a finite series of steps, where each step would involve the transformation from one negative experience to another that was only a "little" less intense than the previous one.

View Three: the mild discomfort of a mosquito bite would be better than two years of excruciating torture, no matter how long one lived and no matter how long the discomfort of a mosquito bite persisted.

View Four: “all-things-considered better than” is a transitive relation.

That is, if it is better to experience A than to experience B, and better to experience B than to experience C, then it must be better to experience A than to experience C.

To avoid unnecessary complications, let me emphasize that the argument I am presenting does not turn on differences in intensity that are imperceptible, or on “least barely noticeable” differences. Specifically, I take it that one negative experience could count as only a “little” less intense than another, in the sense required by Views One and Two, even if, at least when attended to, it were clearly and noticeably less intense. So, while the logic of my argument does not require it, throughout this discussion when I speak of one negative experience as only being a “little” less intense than another, I shall assume that it is nevertheless true that the one is clearly and noticeably less intense than the other.

As should be apparent, View One reflects the Fifth Standard View. View Three reflects the Third Standard View. Most people find both views intuitively compelling. Indeed, I have presented variations of the example to many hundreds of people in countless audiences over the years, and the overwhelming majority have accepted both Views One and Three. In fact, while I suspect that some people might have doubts about the general formulations of the Sixth and Fifth Standard Views, and hence might wonder about the full scope of those views, almost everyone regards Views One and Three as

“obvious,” “uncontroversial,” or “undeniably true.” View Two reflects the plausible assumption regarding a spectrum of burdens presented in section 5.2. It seems to assert a straightforward (modal) fact. View Four is a position widely assumed. Unfortunately, however, Views One, Two, and Three are incompatible with View Four.

The argument for this is straightforward. To see this, we will consider a large series of alternative lives, A1 through An, each of which is lengthy—perhaps, indeed, very lengthy—and each of which includes, as a kind of steady background condition of low level annoyance, 15 mosquito bites per month. First, compare two lives, A1 and A2, and suppose that A1 and A2 are similar, except that A1 contains two years of excruciating torture, A2 four years of torture whose intensity is almost, but not quite, as bad as A1's. Think of the intensity of A2's pain as clearly and noticeably less bad than A1's, yet of roughly the same order of magnitude, so that it would be accurate to characterize A2's pain as only a “little” less intense than A1's. In accordance with View One, most would regard A1 as clearly better than A2, meaning that it would be better to live life A1 than to live life A2. Next, compare A2 with A3, where A3 stands to A2 as A2 stands to A1. Given the choice between two long lives which are otherwise similar, except that one contains four years of very intense pain, and one contains eight years of pain whose intensity is almost, but not quite, as bad, most would judge the former as clearly better than the latter. That is, most would judge A2 better than A3, in virtue of the same considerations that led them to judge A1 better than A2. Iterations of this reasoning imply that A3 would be better than A4, A4 better than A5, A5 better than A6, and so on, with the intensity of the unpleasant experiences slowly, but steadily, decreasing in each successive life. Eventually, in accordance with View Two, one would

be comparing two alternatives, say, A_n and A_{n+1} , such that in addition to the 15 mosquito bites per month present in each outcome, A_n involved a very minor discomfort for a very long time, while A_{n+1} involved a mild discomfort that was almost, though not quite, as unpleasant as A_n 's, but that lasted twice (or three or five times) as long. (Recall that each life being compared is, perhaps, very lengthy.) Finally, we may assume that the extra mild discomfort that obtains in A_{n+1} is just one extra mosquito bite per month, so that for the duration of that discomfort A_{n+1} is bearing 16 mosquito bites per month rather than the “usual” 15.

Summing up, View Two implies that there could be a gradual and finite spectrum of equally lengthy lives, from A_1 to A_{n+1} , such that A_1 contained two years of excruciating torture, together with a lifetime of 15 mosquito bites per month, A_{n+1} contained many years of the mild discomfort of one extra mosquito per month, in addition to the “standard” 15 mosquito bites per month, and each spectrum member after A_1 contained an unpleasant or “negative” experience that was a “little” less intense than its predecessor's but lasted twice (or three or five times) as long. View One implies that for each pair of adjacent members of the spectrum the first is better than the second. Thus, A_1 is better than A_2 , A_2 is better than A_3 , ... A_{n-1} is better than A_n , and A_n is better than A_{n+1} .

Given the foregoing, View Four implies that A_1 must be better than A_{n+1} . But if View Three is right, A_1 is not better than A_{n+1} . To the contrary, in accordance with View Three, the real but mild discomfort of one extra mosquito bite per month would be better than two years of excruciating torture, no matter how long one lived, and no matter how long the discomfort of one extra mosquito bite per month persisted.

This completes my example. I believe it raises serious questions about the consistency of our views about aggregation even within a life. More importantly, I believe that the questions it raises are deeply problematic as I think it is deeply implausible to deny any of the four views underlying the example. More accurately, even if there are reasons to question one or more of the Views as stated, and as we shall see below there may be, I think some version of this example will be compelling. And I am confident that I am not alone in thinking this. After discussing versions of this example with literally hundreds of people over the years, it is apparent that almost everyone, at least initially, accepts some versions of Views One through Four that are incompatible.

5.4 Refining Views One Through Four and Clarifying the Example

Before going on, it should be emphasized that, as indicated, I was only offering a rough characterization of Views One through Four. Likewise, as presented, the example is somewhat under-described, leaving it open to various uncharitable interpretations and objections. Correspondingly, I recognize that both the Views and the example may need to be tightened up and revised in various ways to avoid objections. In most cases, I think it is evident how the Views could be revised and the example spelled out to avoid possible objections, and rather than present the reader with all the necessary nuanced qualifications I shall simply trust the reader to recognize the refinements that might be in order. Still, it might be useful to give a few samples of what I have in mind, to prevent unnecessary misunderstandings or premature rejection of my argument, as well as to indicate lines of refinement that are available.

Someone might assume that a life containing extreme torture is not worth living, while a life that is merely very painful is. Similarly, someone might claim that the most

extreme forms of torture will result in a disintegration of the self—perhaps in the form of a psychotic breakdown—and that a life that involves a tortured life with a disintegrated self would be worse than a merely very painful life with an integrated self. On either of these views one might claim that long before one reaches A_{n+1} one will reach a point, say between A_{27} and A_{28} , where the slight decrease in intensity of pain transforms the case from one involving extreme torture, where one’s life is not worth living, to one "merely" involving great pain, where one’s life is worth living, or from one involving a tortured disintegrated self to one involving a painful integrated self. Correspondingly, one might reject the judgment that A_{27} is better than A_{28} . This, in turn, might appear to undermine View One, and allow one to argue that the transitivity of the “all-things-considered better than” relation would only entail that A_1 is better than A_{27} , not that A_1 is better than A_{28} . Correspondingly, on either of the views in question one might be able to avoid the deeply implausible claim that given the transitivity of the “all-things-considered better than” relation, A_1 , the life involving two years of torture, is better than A_{n+1} , the life involving one extra mosquito bite per month, for many months.

I think one might challenge either of the positions in question. But I’m inclined to grant them. Even so, however, I don’t think they undermine the example. More specifically, such positions suggest that View One may be limited in scope, and need to be revised accordingly, but the example can be spelled out in such a way that the problems still arise with the properly revised version of View One. Let me explain.

Both positions imply that there can be some abrupt transition points, between one pain and another that is only a “little” less intense, where the slight difference in pain levels could, as it were, be like the proverbial straw that broke the camel’s back, making

the difference, literally, between whether the life is worth living or not, or whether the life was integrated or not. If we grant this, we should also grant that View One fails for such transition points. Clearly, it would not be better to have a more intense pain for a given duration, than a “little” less intense pain for two (or three, or five) times as long, if the latter life was worth living but the former life was not. And it is certainly arguable that an integrated life has value that is lacking in a disintegrated life, so that a very painful integrated life might be better than a tortured disintegrated life, even if the pain in the former lasted much longer than the torture in the latter. Still, this limitation of View One’s scope is not especially restrictive. Even if we agree that View One does not apply to the transition points in question, we may continue to find View One compelling for comparing relevant alternatives on each side of the dividing lines between the different kinds of lives. Correspondingly, my example might be spelled out in such a way that all of the lives being compared, from A_1 to A_{n+1} , are understood to be on the same side of each dividing line. Thus, all of the lives might be assumed to be integrated or not (presumably, lives can be disintegrated for reasons not involving torture), and likewise all of the lives might be assumed to be worth living or not.

In my example, I was assuming that all of the lives were, on the whole, worth living. This can be a plausible assumption. For example, I believe that most survivors of the Nazi concentration camps probably had lives that were, on the whole, worth living, even if the portion of their lives when they were in the concentration camps may not have been. Similarly, lives that are sufficiently good for a sufficiently long period of time will, on the whole, be worth living, even if there are significant periods of those lives which are not themselves worth living, and I was assuming that such would be the case for each

of the alternative lives in my example. Likewise, in my example I was assuming that A1 involved two years of the most excruciating torture humanly endurable, compatible with an integrated self. So, as the example is to be understood, throughout the period in which someone would be enduring torture or significant pain, he would retain his sense of self, his awareness that the great pain he was experiencing was happening to him, his desperate hope that the torture would stop and he would be returned to a life of normalcy, and so on. I conclude that while the scope of View One may need to be restricted in light of the positions noted above, those positions aren't relevant to my example as I intend it to be interpreted. That is, I submit that my example can be plausibly fleshed out such that View One is deeply plausible for comparing alternatives A1 through An+1. Our problems remain.

An opponent might persist. Even if one grants that each of the lives of my spectrum may be worth living on the whole, the crucial question is whether there might be some step from an intense pain to one which is only a little less intense, which, though small, nevertheless involves the move from a pain which is unbearable to one which is bearable. If there is, one might argue, then View One should be rejected for that transition point as well. So, for example, if the pain involved in life A27 was literally unbearable, while the pain involved in life A28 was, though only a little less intense, bearable, even if only barely so, it might seem that, contrary to View One, A28 would be better than A27, as it would be better to have to suffer the bearable pain for two (or three or five) times as long as one would otherwise have to suffer the unbearable pain. In this case, the transitivity of the "all-things-considered better than" relation might yield the claim that A1 was better than A27, and that A28 was better than An+1, but since A27

would not be better than A28, it would not imply the unpalatable conclusion that A1, a life containing two years of intense torture, would be better than A_{n+1}, a life containing one extra mosquito bite per month for many years.

The preceding is a particular instance of a general line of argument that is powerfully seductive. Though I've introduced it here, it is sufficiently important to warrant a detailed response in a separate section. Suffice it to say, I believe the argument is mistaken and should be rejected. I will argue for this in chapter six.

Here are some other ways one might attempt to resist the inconsistency I have argued for. One might assume that in A1 the two years of torture takes place at the end of a long life, while in A2 the four years takes place at the beginning. One might then reject View One, arguing that A2 would be better than A1, since it would enable one to avoid a lifetime of dreaded anticipation and since it is more tragic to end one's life in misery than to begin one's life that way. Alternatively, one might assume that year after year after year of an extra mosquito bite might eventually have the effect of the infamous “water torture”—where a steady drip of water, over time, ultimately drives one crazy, even though each drop, by itself, would be innocuous. With this assumption, one might deny the Third View, arguing that A1 would, in fact, be better than A_{n+1}. Or, echoing the first lines broached above, one might assume that even though there might only be a small decrease in the intensity of a pain, say, between A12 and A13, these differences might be correlated with neurological differences that might have tremendously significant consequences. Perhaps, on the “straw that broke the camel's back” analogy suggested above, although A13's pain might only be a little less intense than A12's, the neurological consequences of the latter might involve permanent psychological

scarring—say painful memories and terrifying nightmares for the duration of one’s life—while those of the former might not. On such an assumption, of course, one could reject View One, and its implication that A12 would be better than A13.

The foregoing illustrates that my example could be construed in ways that would render it untroubling. But this does not seriously threaten my argument. Rather, it illustrates that my example, and the Views underlying it, need to be refined and interpreted carefully, and that, among other things, they depend on appropriate "other things equal" clauses. In my example, I explicitly state that A1 and A2 are “similar” except for the small difference in the intensity of their unpleasant experiences, and the length of those experiences. And I claim that A3 stands to A2 as A2 stands to A1, and imply, more generally, that each alternative A_m is to be regarded as “similar” to A_{m+1} , except insofar as their pains differ in intensity and duration. Correspondingly, I was explicitly intending to rule out scenarios such as those suggested above. Similarly, I take it for granted that Views One through Three assume appropriate “other things equal” clauses. Nobody would deny that if A12 has horrendous consequences that A13 lacks, then A13 might be better than A12 even if it involved a pain that was almost as bad as A12’s, but that lasted two (or three, or five) times as long. View One is offered as plausible on the assumption that the alternatives don’t differ in significant ways besides the intensity and duration of the pains in question. Alternatively, one might insist that the example is to be understood so that whatever negative experiences obtain in A_m —including whatever negative consequences might ensue from any other negative experiences— A_{m+1} will have comparable negative experiences that are a little less intense but last two (or three, or five) times as long.

So, my example assumes that the timing of the unpleasant experiences will be comparable in each life. It also assumes that there are no unstated side-effects, or "extrinsic" bads associated with the alternatives. A slight increase in the intensity of torture might lead to a lifetime of sleepless nights and painful memories. But suppose it didn't. Lots of tiny annoyances could drive one to permanent distraction, but embedded in a rich and fulfilling life, they needn't. And so on. I suggest, then, that while my example, and Views One through Three can be interpreted, uncharitably, in all sorts of ways that would render it innocuous, other interpretations are available. Properly spelled out, Views One through Three retain their plausibility, and my example retains its force. The inconsistency of our views is not easily avoided.

Some people acknowledge the general appeal of View One, but suggest that it no longer holds once an unpleasant experience becomes mild enough. Correspondingly, they deny that the transitivity of the "all-things-considered better than" relation implies that A must be better than A_{n+1} , because they deny that A_n is better than A_{n+1} . Frankly, I don't really understand this position. To be sure, a mosquito bite is pretty minor as far as unpleasant experiences go, but it is still annoying to have one. Surely, there is no doubt that, *ceteris paribus*, it would be better to have fewer mosquito bites than more mosquito bites—better to have one extra mosquito bite per month for one year, a decade, or a hundred years, than to have one extra mosquito bite per month for two (or three, or five) times as long. But then, by the same token, it seems clear that it would be better to have an experience that was noticeably more bothersome than a mosquito bite, as long as it was only a little more bothersome (say, for the sake of argument, 10% worse,

though this suggests a level of phenomenological precision that probably doesn't exist), than to have a mosquito bite that lasted two (or three, or five) times as long.

The assumption underlying the above position seems to be that once an unpleasant experience becomes minor enough, its duration no longer matters. As indicated, I see no reason to accept this assumption. Thus, I think View One remains plausible even at the level of mosquito bites, and that A_n is, indeed, better than A_{n+1} , as View One implies. But even if one rejects this, a variation of my counterexample is easily constructed. Take the intensity of a negative experience that is just above the level where duration no longer matters. Suppose this is the level obtaining in A_p . Then we can rerun my argument from A_1 to A_p . View One will presumably apply to the lives from A_1 to A_p , so A_1 will be better than A_2 , A_2 better than A_3 , ..., and A_{p-1} better than A_p . The transitivity of the "all-things-considered better than" relation would then imply that A_1 must be better than A_p . But I now deny this. If A_p 's negative experience is so minor that if it were any less bad it wouldn't matter how long it lasted, then it is virtually insignificant. Given this, I believe View Three would apply. A_p would be better than A_1 . Two years of excruciating torture would be worse than any amount of such virtually insignificant pain.

The preceding suggests that I could limit the scope of View One. I could grant that at the most extreme end of the spectrum of negative experiences, View One would no longer hold. Nevertheless, it seems clear that it would hold for a large enough portion of the spectrum to generate the problem I have highlighted. There will be a range of cases for which our views remain inconsistent.

A similar point might be made to those who have questioned whether View One is plausible for extremely short durations. Thus, for example, it might be claimed that a negative experience that only lasted a nanosecond wouldn't be better than a negative experience that was a little less intense, but lasted two (or three, or five) nanoseconds. Here, it might be claimed that it wouldn't matter which negative experience one had. Or it might even be claimed that for extremely negative experiences, like excruciating stabs of intense pain, it would always be better to experience a less intense pain than a more intense pain, as long as both events were extremely short.

I am not persuaded that we should accept the preceding suggestion, but I also don't think it matters. First, the logic of the argument I presented would allow View One to be revised so that the less intense negative experience might last much longer than the more intense one, even more than five times as long if necessary. And surely an extremely intense stab of pain that only lasted a nanosecond would be better than one that was only a little less intense, but that lasted for one second, or five seconds, or a minute! So View One might be revised to allow for longer lengths of the less intense negative experience for instances of short duration, and the problems of inconsistency would remain. But one might also simply limit View One's scope, and grant that it doesn't apply for durations of nanoseconds, or perhaps not even for durations of less than a second, or even a minute. As above, though I think one needn't grant such a claim, it is enough if there is a spectrum of cases, and durations, for which it does apply. And as my example illustrates, I think there clearly is.

The preceding is related to a point that can be made about both View Three and the Third Standard View. The views in question don't simply maintain that if the

difference in quality of burdens is sufficiently high, then any duration of the smaller burden will be better than any duration of the greater burden; rather, they suggest that any duration of the smaller burden will be better than a “substantial” period of duration of the larger burden. This limitation in the scope of the two views is not ad hoc, rather it is both necessary and appropriate. The truth is that if the duration of a significant burden becomes short enough, it ceases to be a significant burden. Indeed, it may even cease to be a burden. So, while it may well be true that two years of intense torture is less preferable than any duration of a mosquito bite, it needn’t be true that two minutes, or two seconds, or two nanoseconds of intense torture would be less preferable than any duration of a mosquito bite. In fact, it may not be meaningful to speak of only two minutes, or two seconds, or two nanoseconds of pain as intense torture. Stabs of pain, even if willfully and maliciously inflicted, may have to persist for a “significant” duration before they might count as torture.

I have no particular view about how long the willful infliction of pain might have to persist before it should “properly” be counted as torture. Nor do I have a developed view about where the cut off point, or range, might be, such that once intense torture persisted beyond that point, or range, it would be less preferable than any amount of a mosquito bite. But it is perfectly plausible to believe that there is such a point, or range, and that two years of intense torture is well beyond it. This, of course, is all View Three contends, and all I need for my argument.

In sum, View Three and the Third Standard View are already formulated so as to recognize that their scope is limited. And perhaps they need to be further refined to recognize other limitations in their scope. Likewise, as we have seen, the other views I

have presented may also need refinements to reflect ways in which they, too, are limited in scope. However, this does not affect my argument, as long as suitably revised versions of the views are compelling, as I believe they are, for some spectrum of cases like the one I presented.

I conclude that while my example is liable to interpretations that would render it innocuous, properly spelled out it remains deeply problematic. The example highlights significant inconsistency in our thinking; inconsistency that can only be avoided by giving up a deeply plausible position.

5.5 An Objection to View Three: Invoking Principles of Decomposition and Recombination

In the preceding section, I suggested that a negative sensation might have to persist for a “significant” period of time, before it could count as torture. Still, any duration of torture will be composed of a large number of nanoseconds of something, which we may, for lack of a better word, refer to as prototorture. Now I’ve granted that one nanosecond of prototorture might be better than n years of a mosquito bite, for some n . Suppose this is so. Some people might use that as a key premise in a “proof” against View Three. Their argument might run as follows.

There are approximately 63 quadrillion nanoseconds in two years. For simplicity, let us just say that there are q nanoseconds in two years. Now imagine a life with a duration of $q*n$ years. Such a life is composed of q segments of n years each. Likewise, two years of intense torture is composed of q nanoseconds of prototorture. Now, by hypothesis, one nanosecond of prototorture is better than n years of a mosquito bite. So, q nanoseconds of prototorture must be better than $q*n$ years of a mosquito bite.

Therefore, View Three should be rejected. It is not true that any number of years of a mosquito bite would be better than two years of intense torture.

This argument has two related components. It invokes a principle of decomposition, according to which a given whole can be meaningfully decomposed into a set of disjoint parts. In this case, a life of duration $q \cdot n$ years, is regarded as decomposable into q segments of n years each, and two years of torture is regarded as decomposable into q segments of one nanosecond each. Additionally, the argument invokes a principle of additive aggregation, according to which the value of the whole is equal to sum of the values of each part. Assume that the value of A , $V(A)$, can be represented by a real number, x , that the value of B , $V(B)$, can be represented by a real number y , and that if A is better than B , then $V(A)$ is greater than $V(B)$, that is, $x > y$. Given the preceding components and assumptions, it is an easy matter to “prove” the conclusion in question, since it is a simple algebraic fact that for all x , y , and p , if $x > y$ and p is a positive number then $p \cdot x > p \cdot y$.

In other words, the argument supposes that two years of torture and a lengthy lifetime of mosquito bites can each be divided into q parts, and that the value of each whole is just equal to the sum of the values of its parts. So, the value of two years of torture will just be equal to q times the value of a nanosecond of prototorture, and the value of a life of $q \cdot n$ years of a mosquito bite will just be equal to q times the value of n years of a mosquito bite. Since, by assumption, a nanosecond of prototorture is better than n years of a mosquito bite, the value of the former must be greater than the value of the latter, and hence q times the value of the former must be greater than q times the value of the latter, given that q is a positive number. It follows, on this reasoning, that the

value of two years of intense torture must be greater than $q \cdot n$ years of mosquito bites, and hence that the former must be better than the latter. Therefore View Three should be rejected.

Many people are attracted to an argument like the preceding one. But I think it should be rejected. The first point to note is that it assumes, as a starting point, the legitimacy of an additive aggregationist approach in assessing the value of outcomes. View Three obviously rejects such an approach. Hence, the objection in question does not, so much, argue against View Three, as beg the question against it. Of course, advocates of additive aggregation might retort that View Three is in the same boat. That View Three does not, in fact, provide an argument against additive aggregation, but simply begs the question against it. I accept this rejoinder. But it should be noted that View Three was never offered as an argument against additive aggregation. It was offered as a position that virtually everyone finds deeply intuitive. This is important. If it comes down to choosing between two positions, each of which can be construed as begging the question against the other, typically the advantage will lie on the side with the greatest intuitive plausibility, and the burden of proof will lie on the other side.

But advocates of View Three can do much more than charge the objection with question begging. They can directly challenge the plausibility of the assumptions underlying the objection. I have, of course, already offered a series of examples illustrating the intuitive implausibility of applying principles of decomposition and additive aggregation in assessing the value of individual lives, in section 4.3. But other considerations, some of which are more general and theoretical, can also be offered

against the assumptions in question. To do this, it will be useful to first take note of a domain where the assumptions hold, the domain of arithmetic.

To facilitate discussion, let us say that in arithmetic, the value of a number is just equal to the number itself. For example, the value of the number eight is just eight, which we will express as $V(8) = 8$. Now in arithmetic, the principle of decomposition holds. Any numerical “whole” can be decomposed into a set of disjoint “parts.” In fact, any whole can be decomposed into many different sets of equal or unequal size. For example, 8 can be decomposed into two parts of equal size, four parts of equal size, two parts of unequal size, three parts of unequal size, etc. That is, $8 = 4+4$, $8 = 2+2+2+2$, $8 = 5+3$, and $8 = 5+2+1$. Moreover, in arithmetic, additive aggregation holds. Specifically, the value of a whole is equal to the sum of the values of its parts. So, $V(8) = V(5+3+1) = V(5)+V(3)+V(1)$. We can further note that in virtue of the commutative property of arithmetic, according to which $x+y = y+x$, a view that incorporates both the principle of decomposition and the principle of additive aggregation will also entail a principle of recombination, according to which the parts of a whole can be “rearranged” in any order without affecting the value of the whole. So, for example, $V(8) = V(5+3+1) = V(3+5+1) = V(1+3+5)$ and so on, since $V(5)+V(3)+V(1) = V(3)+V(5)+V(1) = V(1)+V(3)+V(5) = 8$. All this is straightforward. The axioms of arithmetic entail principles of decomposition, additive aggregation, and recombination of the sort invoked by the objection to View Three.

Because of the familiarity of arithmetic, and the ubiquitous role that it plays in our understanding of both natural and social sciences, it may be natural to suppose that arithmetic, and its underlying assumptions, hold for all domains. But while no domain

can be, strictly speaking, incompatible with arithmetic, there may be domains to which it doesn't apply, or, more importantly, cases within various domains for which it would be mistaken to directly apply the underlying assumptions of arithmetic. This is a point that has been recognized by philosophers and others throughout the ages. For example, it has been recognized wherever a version of holism has been defended, with its fundamental insight that in some cases the value of a whole, considered as a whole, is greater or less than the sum of the values of its parts, considered separately as parts.² More generally, it has long been recognized that the normative realm is rife with examples where the assumptions underlying the objection to View Three fail.

Holism is often regarded as a mysterious doctrine. But as I am employing the term here, it needn't entail anything more than ordinary facts of the sort that economists describe as interaction effects. Here is a standard example. Consider two meals, each composed of two parts. One meal combines a nice fish, F, with a fine white wine, WW, the other beef, B, with a fine red wine, RW. It is widely accepted that the value of these two meals cannot be assessed along the lines assumed by the objection to View Three. That is, one cannot assess the value of the two meals by decomposing them into their separate parts, determining the value of each part separately, and then summing the resulting individual values. The idea is that the white wine and fish may complement or interact with each other in such a way that each enhances the taste of the other. As a result, the value of a meal which combines the white wine and the fish may be greater than the value of each partaken separately. That is, $V(WW+F) > V(WW)+V(F)$. Similar remarks might apply to the red wine and the beef, so that $V(RW+B) > V(RW)+V(B)$. By the same token, there may be negative interaction effects, such that, when combined, the

tastes of white wine and beef might detract from each other, and similarly for the tastes of red wine and fish. In such cases, $V(WW+B) < V(WW)+V(B)$, and $V(RW+F) < V(RW)+V(F)$. All this is commonplace. Similarly, you don't have to be a gourmand to know that the order of a seven course meal is crucially important to the success, or value, of that meal. So, as with the case just discussed, to determine the value that someone would receive from a meal, it is not enough to consider the value that they would receive from ingesting each item separately, and then add those values together. One needs to know the order in which the items would be consumed, how, if at all, they would combine, and the interaction effects, if any, between them. So, for the domain of gustatory value, at least, one must reject the principles of decomposition, additive aggregation, and recombination underlying the objection to View Three.

Another classic example of holism, or organic unity, for which the principles of decomposition, additive aggregation, and recombination evidently fail concerns beauty. Consider, for example, the beauty of a face.³ It is widely accepted that the most beautiful face, overall, may not be composed of the most beautiful features as judged independently. Specifically, there may be two sets of features—say, eyes, nose, mouth, chin and ears—such that each member of the first set might be judged more beautiful than the corresponding member of the second set, when considered alone, and yet the combination of more beautiful features, might be considerably less beautiful than the combination of less beautiful features. That is, letting $V(\text{Eyes1}) > V(\text{Eyes2})$ express the judgment that, considered alone, the eyes in the first set of features are more beautiful than the eyes in the second set of features, and using a similar notation for the other comparisons, it might well be that $V(\text{Eyes1}) > V(\text{Eyes2})$, $V(\text{Nose1}) > V(\text{Nose2})$,

$V(\text{Mouth1}) > V(\text{Mouth2})$, $V(\text{Chin1}) > V(\text{Chin2})$, and $V(\text{Ears1}) > V(\text{Ears2})$, and yet it may still be that $V(\text{Eyes1}+\text{Nose1}+\text{Mouth1}+\text{Chin1}+\text{Ears1}) < V(\text{Eyes2}+\text{Nose2}+\text{Mouth2}+\text{Chin2}+\text{Ears2})$. Perhaps the most beautiful eyes would be too wide for the most beautiful nose, which might be petite, and neither might fit especially well with the most beautiful chin, which might be pronounced. Moreover, it is evident that nothing like a principle of recombination applies in the case of beauty. Order, in this case in the form of spatial arrangement, makes all the difference, as moving a pair of eyes two inches up, or to the left, or interchanging an eye with an ear might well result in a grotesque appearance that only Picasso could love! Again, all this is commonplace, but it serves to remind one that while the principles of decomposition, additive aggregation, and recombination apply in arithmetic, they often fail in other domains.

Here is a familiar example from the normative realm. Suppose we know four facts about a given outcome. We know that one person has been good, one person has been bad, one of the two people has fared well, and the other has fared poorly. These factors are insufficient to judge the overall goodness of the outcome, because the goodness of the outcome is not just an additive function of the value of those factors considered independently. It matters whether it is the good person who fares well, and the bad person who fares poorly, or vice versa. In the former case the outcome is just, in the latter it is unjust, and this makes a significant difference to the overall goodness of the outcome. Here, too, principles of decomposition, additive aggregation, and recombination fail to apply.

Finally, consider two cases more directly analogous to the objection we are considering. In the first case, we know that there are many pieces of straw that have been

placed on a camel's back, and subsequently removed. In the second, we know that there are many snowflakes that have landed on the ground, and then melted. What can we say about the effects in those two scenarios? Clearly, not much. Timing, here, can make a crucial difference to the interaction effects, if any, that obtain in these scenarios. For example, if each piece of straw that was placed on the camel's back was removed prior to another piece of straw's being similarly placed, it would be utterly inconsequential that the camel had so much straw placed on its back over the course of its life. Things would be markedly different, however, if all the straw were placed on its back at once, or each piece was added sequentially, but none were removed until after the camel's back was broken! Likewise, many snowflakes spread out over many winters might amount to nothing more than a series of inconsequential "dustings" of light snow. While the same total amount of snow occurring in a 3 hour span might produce a devastatingly dangerous blizzard. This is, of course, perfectly obvious. But it is important. The value, or disvalue, of a given total amount of straw on a camel's back, or a given total number of snowflakes that fall, cannot be determined by simply decomposing the total amounts into individual components of single pieces of straw, or single snowflakes, assessing the value or impact of those components considered separately, and then simply summing the resulting values or impacts. In some cases there may be a dispersal, or dissipation of components, so that there are no interaction effects between them, while in other cases there may be an accumulation of factors resulting in significant interaction effects. Given this, the principles of decomposition, additive aggregation, and recombination do not apply for such cases.

Return now to the argument against View Three. It fails, because the relation between a nanosecond of prototorture and two years of severe torture is like the relation between a snowflake and a blizzard. Even if we agree that, considered by themselves, n years of a mosquito bite are worse than a nanosecond of prototorture, and that two years of torture consists of q nanoseconds of prototorture, we needn't agree that $q*n$ years of mosquito bites are worse than two years of torture. One nanosecond of prototorture every n years for $q*n$ years may be utterly inconsequential, just as one light dusting of snowflakes every winter might be. This is because each nanosecond of torture, or light dusting, and its effects would have completely dissipated before the next one occurred. But q consecutive nanoseconds of prototorture would be a very different matter, as would the accumulation of snowfall in a blizzard. In these cases there would be no dispersal or dissipation of individual effects, but instead an interaction, combination, and accumulation of factors and effects that would radically alter their normative significance. In sum, even though there might be the same total number of nanoseconds of prototorture in two years of unremitting torture, as in 1 nanosecond of prototorture every n years for $q*n$ years, the former would be a human tragedy, while the latter would be inconsequential.

The preceding discussion reveals that just because A would be better than B, and C better than D, it does not follow that A+C would be better than B+D. Similarly, just because A would be better than B, it doesn't follow that n instances of A would be better than n instances of B. The key question is how the different factors are related to each other and, in particular, the nature, if any, of the interaction effects between them. On the assumption that one nanosecond of prototorture would be better than n years of a

mosquito bite, which I've granted, View Three expresses the view that there is a powerful accumulated effect of consecutive moments of prototorture, but not between consecutive periods of mosquito bites lasting n years each. This seems right. The effect of a mosquito bite in one year will be completely dissipated long before the effect of a mosquito bite in subsequent years arrives. Correspondingly, we don't believe that while a few years of a mosquito bite would merely be a nuisance, lots of years of a mosquito bite would rise to the level of a human tragedy. They would remain merely a nuisance. A persisting nuisance, to be sure. But still only a nuisance. This is why View Three seems so powerful. While one nanosecond of prototorture would be insignificant, and better than n years of a mosquito bite, q consecutive nanoseconds of prototorture would rise to the level of excruciating torture, while q periods of a mosquito bite, each lasting n years, would remain a mere nuisance. And as View Three rightly reflects, a nuisance is better than a tragedy.

I conclude that we should reject the argument against View Three. It relies on principles of decomposition, additive aggregation, and recombination that apply in the domain of arithmetic, but not to the normative realm generally, or to the sphere of pain in particular.

5.6 A Proportionality Argument Against View Three

Some people might offer the following argument against View Three. I have granted that a nanosecond of prototorture would be insignificant. This, it might be claimed, is because a nanosecond is a teeny tiny insignificant "blip" in a lifespan lasting 70 years. For beings such as us, who live as long as we do, anything that endures for such a tiny fraction of our lives must rightly be regarded by us as completely trivial.

After all, in the absence of long-term side affects, surely, the overall assessment of our lives will be utterly unaffected by what goes on during a single nanosecond of our life, a period which is a mere 2205 quadrillionth of a normal 70 year lifespan.

Similarly, it might be argued, if we lived long enough, specifically, if we lived 4410 quadrillion years, two years would be a teeny tiny insignificant “blip” in our lifespans. Correspondingly, in the absence of long-term side affects, surely, the overall assessment of such lives would be utterly unaffected by what goes on during a single two year period, a duration which would be a mere 2205 quadrillionth of the “normal” lifespan. Thus, anything that endured for such a tiny fraction of such lives would rightly be regarded as completely trivial. Therefore, View Three should be rejected, as an extra mosquito bite per month that lasted for quadrillions of years could easily outweigh something that was completely trivial.

We can call the preceding argument the proportionality argument against View Three. It asserts that since two years of torture within a sufficiently long life stands in the same proportion as one nanosecond of prototorture within a normal 70 year lifespan, the former, within such a life, should have the same normative significance as the latter, within a normal life. At first blush, the proportionality argument has some intuitive appeal. However, on reflection, it is clear that it should be rejected. Let me offer three related arguments against it.

First, a life spanning quadrillions of years has the potential to be vastly more valuable or disvaluable, in absolute terms, than a life spanning a mere 70 years, just as a life spanning 70 years has the potential to be vastly more valuable or disvaluable in absolute terms than a life spanning 7 nanoseconds. Correspondingly, the fact that two

years stands in the same proportion to a life of quadrillions of years as one nanosecond does to a normal life, just means that the former has the potential to be vastly more valuable or disvaluable, in absolute terms, than the latter. Thus, it is perfectly possible that while one nanosecond of prototorture is, in absolute terms, completely trivial, two years of torture is, in absolute terms, a human tragedy. These absolute facts wouldn't be altered by the further fact that the latter might be embedded in a life of vastly greater value, or disvalue, than the former. This explains why one might accept View Three, even if one grants that one nanosecond of prototorture would be better than a normal lifetime of one extra mosquito bite per month. If, as seems possible, one extra mosquito bite a month, say from 15 to 16, never rises to more than a mere nuisance, no matter how long one lives, one should accept View Three, since a life containing a human tragedy won't be better than a life containing an ongoing nuisance. But this is completely compatible with the judgment that a normal life containing the mere nuisance of an extra mosquito bite a month would be better than a life containing the completely trivial occurrence of one nanosecond of prototorture.

Second, the proportionality argument claims that the reason a mere nanosecond of prototorture is completely trivial, within a normal lifespan, is that a nanosecond is but a tiny fraction of such a life. If this were true, that would explain how two years of intense torture would also be completely trivial, if they were embedded within a life that was sufficiently long. But it isn't true. The reason that a nanosecond of prototorture is completely trivial is simply that it wouldn't even be noticed! A nanosecond of prototorture would almost certainly be imperceptible at the conscious level, and probably even at the subconscious level. Hence, a nanosecond of prototorture would remain

completely trivial even if it occurred within a life that persisted only 7 years, or 7 months, or 7 days, or 7 seconds. So, the proportion of the life that is subject to a nanosecond of prototorture is simply not what accounts for its being trivial, facts about perceptibility and human psychology are.

Consider a normal human who is subjected to a nanosecond of prototorture. Such a person would presumably not even notice the event. Moreover, even if there were, at some level, some subconscious awareness of the instantaneous event, there would be no particular reason to even form the desire to be rid of the experience, since it takes time to formulate a desire, and the event would almost certainly have already passed before the desire was completely formulated. Thus, if the occurrence of a nanosecond of prototorture somehow succeeded in prompting the desire to be rid of the prototorture, there would be no moment at which that desire was frustrated, as the desire would have almost certainly already been fulfilled by the time it was fully formed. So there is, indeed, good reason to regard a nanosecond of prototorture as completely trivial, as it would have (virtually?) no negative impact on the quality of a person's conscious mental states, and would involve (virtually?) no frustration of a person's desires.

The situation is rather different in the case of two years of intense torture. Even within a life of extraordinary length, two years of intense torture would involve incredibly negative conscious mental states and give rise to the deep, fervent, and overriding desire to be rid of such states. That desire would, of course, persist and be horribly frustrated throughout the two years. The combination of agony suffered and the frustration of what would almost certainly be, at the time, one's deepest desire, accounts

for why two years of intense torture would be a human tragedy, even within a life of extraordinary length.

The proportionality argument requires that the facts of human psychology would have to change, so that two years of intense torture within a sufficiently long life would have the same psychological impact as one nanosecond of prototorture would within a normal life. But there is no reason to grant this. Even if, as a matter of fact, human psychology would undoubtedly change in countless, unpredictable, ways if we lived for quadrillions of years, one cannot assume that it would have to, as a matter of logical necessity, change in the way the proportionality argument requires. For the sake of my argument, it is enough if there is some possible world in which beings with psychologies like ours persisted for quadrillions of years. Since such a world is possible, the proportionality argument should be rejected.

Finally, suppose that while hell is a place of intense torture, heaven simply turns out to be an earth-like place, where people can experience normal earth-like lives indefinitely. On the proportionality argument, it would be completely trivial if God chose to send us to hell for billions and billions of years, as long as He sent us to the earthly afterlife for the rest of eternity. Indeed, he needn't even send us to the earthly afterlife for an eternity. He merely needs to send us for a sufficiently long period, that the proportion between the time spent in hell and that spent living the earthly existence would be equal to the proportion between a nanosecond and a normal lifespan. This view is absurd. If God were to send us to hell for billions and billions of years it would be a great tragedy; and He could not avoid the charge of sadism, or inflicting grievous harm

on us, merely by making sure that he also provided us with many years of an earthly existence. The proportionality argument should be rejected.

5.7 Trusting Our Intuitions Regarding Inordinate Lengths of Time

Some people offer a different challenge to View Three. Specifically, while they grant the intuitive plausibility of the claim that no matter how long the discomfort of a mosquito bite persisted, it would be better than two years of excruciating torture, they insist that our intuitions are not to be trusted in such matters. Such people offer numerous explanations for why we might naturally, but nonetheless wrongly, accept View Three; explanations that ultimately rest on the ground that our intuitions were not developed, and are poorly equipped, to adequately respond to the overall significance of small amounts spread out over innumerable years.⁴

The claim here is that to move from A_1 , the state involving two years of excruciating torture, to A_{n+1} , the state involving one extra mosquito bite per month, would require so many intervening steps, that we would have to imagine what it would be like to have an extra mosquito bite for billions or zillions of years. (Recall, that each step involves a slightly less unpleasant experience, but one that lasts twice--or three or five times--as long as the previous one.) But, it is argued, we have no intuitive conception of what it would be like to live such a lengthy life, and in particular no intuitive conception of how bad it would be to experience an extra mosquito bite for such a long time. Correspondingly, it is argued, our intuitions about such wild science-fiction type cases cannot be trusted. Instead, we should trust the judgment yielded by theories like expected utility theory, which seem to work perfectly well in the ordinary cases for which our intuitions seem plausible, and figure it is those judgments—the judgments of our

considered theory—that we should go by in cases involving inordinately large numbers of small amounts. And what expected utility theory tells us, it is contended, is that no matter how slight a mosquito bite’s discomfort may be, if it lasts long enough, eventually the total amount of discomfort one experiences will outweigh, and then be less preferable than, the pain of two years of excruciating torture. On this view, then, our understanding supposedly leads us to recognize a truth that our imagination fails to appreciate, namely that View Three should be rejected.

I have some sympathy for this kind of response. People often appeal to intuitions in cases where our intuitions cannot be trusted. Moreover, our intuitions are notoriously suspect in cases involving small amounts or large numbers.⁵ Nevertheless, I do not find this response compelling. Let me indicate three reasons why.

First, the response turns on a dubious assumption, namely, that moving from A_1 to A_{n+1} would necessarily involve so many steps that the mosquito bite’s duration would be unimaginably long—billions, or perhaps zillions, of years. This may be false. Consider the following example. A man who is 5' 2" is very short. A man who is 6' 4" is very tall. In this range, a difference of 2 inches is noticeable, but relatively small—less than 3.3% for someone 5' 2", less than 2.6% for someone 6' 4". Yet, one can move from the very short to the very tall, in just seven relatively small steps.⁶ If, by analogy, one could move from the pain of torture to the discomfort of a hangnail in seven steps, the relevant trade-off would be between 2 years of torture and 256 years of an extra mosquito bite. 256 years is a long time, but it is hardly unimaginable. And I am confident that were I fortunate enough to live 256 years, I would much prefer an extra mosquito bite throughout my life to two years of torture somewhere within it.

Moreover, if one seriously doubts whether one can vividly imagine 256 years of an extra mosquito bite, one might change the torture's length. If one starts with 1 year of torture instead of two, the mosquito bite might only last 128 years. If one starts with one day of torture, the mosquito bite might only last 128 days! One day of torture is clearly imaginable; so, too, is 128 days of an extra mosquito bite. When I vividly think of these possible experiences, I know which I prefer. The mosquito bite would be vastly better than the torture. These are not cases where my imagination gives out. Nor do I think it fails me here. It rightly guides my judgment. All things considered, the life with the torture would be much less preferable than the life with the extra mosquito bite.⁷

A perceptual analogy might be useful here. The visual spectrum ranges across the seven colors of the rainbow: red, orange, yellow, green, blue, indigo, and violet. The differences between adjacent members of the spectrum are clearly perceptible, yet still relatively small. Phenomenologically, there is not that great a difference between red and orange, or indigo and violet. But, then, one can get from one end of the visual spectrum all the way to the opposite end in six relatively small steps.

Do the steps seem too big? Add intervening ones. Consider the following spectrum: red, reddish orange, orange, orangish yellow, yellow, yellowish green, green, greenish blue, blue, indigoblue, indigo, indigoviolet, violet. It is hard to deny that, phenomenologically, the gaps between adjacent members of this spectrum are pretty small, though still clearly perceptible. Yet it would only take 12 steps to get from one extreme of the spectrum to the other.

The preceding suggests that although, phenomenologically, the pain of torture and the discomfort of a mosquito bite are at opposite ends of the pain spectrum, it may take a

surprisingly small number of short steps to get from one to the other. To be sure, each step, though short, may involve differences that are clearly noticeable and of some significance. But there is no reason to deny this. Given the choice between a greater pain for a certain length of time, and a lesser pain for twice as long, I prefer the former, unless the latter is much less intense. Clearly, my choice requires that the difference between the two pains not be too great, but it certainly does not require that it be either imperceptible or the least perceptible difference.

In sum, there is reason to doubt whether my example depends on having to imagine what it would be like to have an extra mosquito bite for billions, or perhaps even zillions, of years. The relevant tradeoff might be between two years of torture and a few hundred, or maybe thousands of, years of a mosquito bite. Or, if we change the scale, between two months of torture and a few hundred, or maybe thousands of, months of a mosquito bite. Such tradeoffs are not the sort typically made, but they are not unimaginable.⁸

Let me acknowledge that the preceding considerations are purely suggestive. I certainly have not established that the case of pain is analogous to those of height and color. Although a small number of fairly small steps gets one from a very short to a very tall person, and from one end of the visual spectrum to the other, it remains possible that vastly many steps would be necessary to get from the pain of torture to the discomfort of a mosquito bite. Still, given facts about human evolution, and the roles our different sense modalities play in survival, one might not expect our phenomenological powers of discrimination to differ radically across our sense modalities. That is, it isn't clear that we should need vastly greater powers of discrimination among pains, than among colors.

Indeed, from the perspective of armchair psychology, it might seem sufficient if we are able to distinguish between three kinds of pains: the innocuous, which we can safely ignore; the threatening, which we might want to avoid and should keep an eye on; and the dangerous, which we definitely need to avoid if possible, and otherwise remedy. Moreover, from an evolutionary standpoint the phenomenological differences between such signals might be slight, as long as the signals are clear, reliably detected, and accurately interpreted.

In any event, it seems an empirical fact whether our sensory apparatus for pain operates analogously to our sensory apparatus for color, and that may be all I need for my argument. After all, unless it is metaphysically impossible for the pain and color modalities to be analogous in the way suggested, there are some possible creatures for whom my argument would work--even if it does not, in fact, work for humans. This is all one needs to show the inconsistency between Views One, Two, Three, and Four.

But this is absurd, some will insist. Torture really is very bad. And the discomfort of a mosquito bite is very mild. Surely, it is obvious that one can't get from the very bad to the very mild in a small number of steps. But is it obvious? Or is it possible that this is a place where our intuitions lead us astray? After all, as we've seen, a 6' 4" man really is very tall, and a 5' 2" man is very short, but we certainly can get from the one to the other in seven short steps!

So, I am not convinced that in moving from torture to mosquito bites we would necessarily end up comparing alternatives of such a massive humanly incomprehensible scale that our intuitions can't be trusted. And I find the example particularly gripping and compelling. Still, a second response to the objection in question is to change the

example. So, let me present four different cases where similar problems arise regarding the consistency of our judgments. I shall discuss the first case at length, and the three other cases minimally.

Case I. Imagine a scale of uncomfortable experiences, ranging from a level of 1, the discomfort of a mosquito bite, to a level of 100, extreme torture. Suppose a moderately uncomfortable limp is an 11--significantly worse than a mosquito bite, but not nearly as bad as extreme torture. Start with a choice between A, level 100 discomfort for 2 days, or B, level 80 discomfort for 4 days. B's discomfort is 20% less than A's, but lasts twice as long. Many believe is A better than B. Next compare B to C, where C stands to B as B stands to A. C is 20% less intense, level 64, but lasts twice as long, 8 days. Again, many think B better than C. The tenth choice would be between J, discomfort of level 13.4 for 1024 days, and K, discomfort of level 11 for 2048 days. Again, many would think J better than K.

Given these rankings, transitivity entails that A is better than K. Is it? Given our assumptions, A involves extreme torture for 2 days, K involves a moderately uncomfortable limp for 2048 days, or 5.6 years. Which would you choose for your child--that for the next 5.6 years she suffers an uncomfortable limp, or for 2 days she suffers extreme torture? (Here, as always, I assume there are no relevant side-effects. A pill, or hypnosis, removes all memories of the torture, other children do not mercilessly ridicule your child as a "gimp," etc.) I would choose the limp over the torture for someone I loved. And I would do so because I think it the preferable alternative. But I would also think A is better than B, B better than C, ..., and J better than K.

It might be claimed that the preceding necessarily involves some kind of conceptual confusion. Specifically, it might be claimed that if the discomfort of 5.6 years of a moderate limp seems better than a mere two days of intense torture, then the gap between the pain of torture and the discomfort of a moderate limp must be much greater than my example imagines it to be, such that one could never bridge it in only ten steps of the sort described. On this view, the pain of a moderate limp would have to receive a much lower score than eleven on our scale from one to a hundred, and the real trade-off we would have to imagine would be between two days of torture and a much longer period of moderate limp. But I don't see why we should believe this. B's pain is 20% less intense than A's. That is quite a significant difference. And C's is 20% less intense than B's, also quite a significant difference. The combined result of ten such changes, could be a very large gap indeed, as large as the phenomenological difference between torture and a moderate limp.

Think about it carefully. Imagine being in a state of intense torture, and then having the intensity of that pain reduced by 20%. The result would still be a quite painful state; painful enough, I think, that it might be better to have the former state for only two days, than the latter for four days, twice as long. That is, painful enough to warrant the judgment that A would be better than B. Still, there is no denying that the move from A to B would involve significant improvement regarding the intensity of the pain. Next imagine the pain being reduced by another 20%. This, again, would involve significant improvement regarding the intensity of the pain, but not so much that we should prefer the latter alternative to the former if it lasted twice as long. But then imagine the intensity of the pain being reduced by another 20%. And then another. And another.

And another. And another. And another. And another. And yet another. While each reduction is sufficiently small, that a state involving the former pain might be better than a state involving the latter pain of twice the duration, each reduction would be quite significant, and the combined effect of ten such reductions would be to produce a state, such that the gap between the first state and the last might be so large, that, intuitively at least, we would regard 5.6 years of the latter as better than only two days of the former.

To be sure, our intuitions about such a case might still be mistaken. I have not argued otherwise. But we certainly can imagine what it would be like to suffer extreme torture for two days. And likewise we can certainly imagine what it would be like to suffer a vastly reduced pain, 89% less intense, for 5.6 years. If, as I believe, we intuitively believe the latter would be better than the former, such intuitions cannot simply be dismissed by appeal to the fact that our intuitions cannot be trusted in cases involving inordinate lengths of time.

Case II. A involves random interruption of one's electrical service 20 minutes per time, 9 times a day, for 24 months; B the same kind of interruptions 7 times a day for 42 months; C 5.66 times a day for 63 months; D 4.33 times a day for 94 months; E 3.33 times a day for 141 months; F 2.33 times a day for 210 months; G 1.6 times a day for 315 months; H 1.1 times a day for 473 months; I .76 times a day for 710 months; J .53 times a day for 1065 months.

Case III. A involves driving delays of 4.5 hours per day for 24 months; B delays of 3.5 hours per day for 42 months; C delays of 2.83 hours per day for 63 months; D delays of 2.17 hours per day for 94 months; E delays of 1.67 hours per day for 141 months; F delays of 1.17 hours per day for 210 months; G delays of .8 hours per day for

315 months; H delays of .55 hours per day 473 months; I delays of .33 hours per day for 710 months; J delays of .27 hours per day for 1065 months.

Case IV. A involves 48 garbage pick ups missed per year (against the backdrop where garbage would normally be picked up once a week) for 24 months; B 40 missed pick ups per year for 42 months; C 32 missed pick ups per year for 63 months; D 250 missed pick ups per year for 94 months; E 19 missed pick ups per year for 141 months; F 13.5 missed pick ups per year for 210 months; G 9.23 missed pick ups per year for 315 months; H 6.35 missed pick ups per year 473 months; I 4.42 missed pick ups per year for 710 months; J 3.1 missed pick ups per year for 1065 months.

Cases II, III, and IV are obviously similar. They each present different sets of alternatives involving different kinds and degrees of frustration of different durations. Ranking the various alternatives involves making trade-offs between given levels of frustration for a given length of time, and decreased levels of frustration for longer periods of time. Having asked various people about such cases, I believe that for each of these cases many would judge A better than B, B better than C, C better than D, and so on. But many would also judge, for each of these cases, that J is better than A. Their reasoning appears to be straightforward. When comparing A with B, B with C, C with D, and so on, the degree of frustrations in the different alternatives are sufficiently similar, that it seems reasonable to rank the alternatives merely on the basis of which has the greater sum total of frustrations. But when comparing A with H, the difference in the degree of individual frustrations is so great, that the distribution of the frustrations and its overall effect on the quality of the lives seems more pertinent than the sum total of the frustrations.

Now my aim here is not to defend the intuitive judgments noted above. It is, rather, to illustrate that the same kinds of worries I raise about the spectrum from torture to mosquito bites can arise in other cases that do not involve lives of inordinate lengths. I can certainly imagine what it would be like to have my garbage not picked up 48 weeks out of 52 (Case IV's A). Likewise, I can certainly imagine not having my garbage picked up approximately once every four months for the duration of an 89 year life (Case IV's J). And I can imagine all the intervening alternatives (Case IV's B-H). Such alternatives do not involve unimaginably large periods of billions or zillions of years, for which our intuitions were not developed. Hence, it does not appear that one can readily dismiss our intuitions about such alternatives as obviously suspect. But, for many, their intuitions about Cases I through IV mirror the firm intuitions most have about my original example. I suggest, then, that the conflict in our views that my original example points to is real, and must be squarely faced. Dismissing it as an artifact of wild science-fiction type cases for which our intuitions cannot be trusted is, I think, a mistake.

View Three is an example of a particularly strong anti-additive position. In claiming that the mild discomfort of a mosquito bite would be better than two years of excruciating torture, no matter how long one lived and no matter how long the discomfort of a mosquito bite persisted, it naturally gives rise to the concern about large numbers and lives of inordinate length that we have been discussing. But it is worth emphasizing that one doesn't have to make such a bold claim for the problems I have been raising to arise. It is enough if one believes that there are some cases where alternatives can be rightly judged on the basis of sum totals, while there are other cases where sum totals are not all that matters, as the pattern of distribution is also relevant. The latter claim is a relatively

weak anti-additive position, but it enough to give rise to my concerns. One will either have to reject the first kind of view, that sometimes an additive approach is appropriate for ranking outcomes, the second kind of view, that sometimes distributions are relevant, the transitivity of the “all-things-considered better than” relation, or the empirical assumption that there can be spectrums of cases such that the first criterion is relevant for comparing the alternatives near each other on such spectrums, and the second is relevant for comparing alternatives at the opposite ends of such spectrums. I continue to believe that each of the views is extremely plausible, and that giving any of them up would require major revision in our thinking.

Finally, let me conclude this section with a few words about expected utility theory. I agree that expected utility theory has proved to be an extremely useful and illuminating tool. But it is only that, a useful tool that humans have constructed to help model, explain, and in some cases guide, human behavior. I am not convinced that expected utility theory has to endorse an additive aggregative approach to assessing outcomes, or the value of human lives. If it doesn't, then it obviously can't be invoked to rule out our anti-additive aggregationist intuitions, like those expressed in View Three. But if it does, then it must face all of the objections and worries that have been raised in this chapter and the preceding ones against additive aggregation.

If the Repugnant Conclusion really is repugnant, no appeal to expected utility theory can remove that repugnance. Similarly, if it really wouldn't be a better outcome to give each of innumerable people one lick of a lollipop, if this unavoidably involved one person suffering unbearable agony for countless years, it is no use pretending otherwise because expected utility theory would judge it so. And if expected utility theory cannot

reflect the view that the goodness of a life is determined, in part, not just by how much pain occurs in that life, but by the quality of the pain and the way it is distributed, then so much the worse for expected utility theory.

I am, in fact, a fan of expected utility theory. But the problems I have been raising require a response. They cannot, I think, simply be rejected because they are incompatible with expected utility theory. Decisions need to be made about which of our views should be rejected, and why. And in making those decisions the costs of doing so must be fully realized. When this is done, we may see that expected utility theory needs revision or is limited in scope. Or we may see that it remains intact. But even if it remains intact, we will have a much better understanding of why this is so, including a much deeper understanding of the nature of moral ideals and practical reasoning, if we confront the issues I am raising head on, than if we simply sweep them under the rug as “obviously” misguided, because incompatible with expected utility theory. Ultimately, the issues I am raising challenge some of the very foundations of expected utility theory. Correspondingly, one cannot respond to such issues simply by an appeal to the very theory that they challenge.

5.8 Summary of Chapter Five

In this chapter, I claimed that most people accept both the Fourth Standard View—Trade-offs between Benefits and Duration are Sometimes Desirable Within Lives, and the Fifth Standard View—Trade-offs between Burdens and Duration are Sometimes Desirable Within Lives. On the Fourth Standard View, in general, it is better to receive a lesser benefit for a longer duration, than a greater benefit for a shorter duration, if the longer duration is “sufficiently” longer, and if the difference in the quality

of the benefits is not “too” great. On the Fifth Standard View, in general, it is better to receive a greater burden for a shorter duration, than a lesser burden for a longer duration, if the longer duration is “sufficiently” longer, and if the difference in the quality of the burdens is not “too” great. Considering various examples, I showed that both the “all-things-considered better than” and Fifth Standard Views are quite robust. Both apply to a wide range of cases, involving benefits or burdens of varying significance and durations of varying length.

I next showed that just as there is an inconsistency between the First Standard View, the Second Standard View, a plausible assumption, and the transitivity of the “all-things-considered better than” relation, so there is an inconsistency between the Fourth and Fifth Standard Views, Chapter Four’s Third Standard View, a plausible assumption, and the transitivity of the “all-things-considered better than” relation. This should not be surprising, given that the Fourth and Fifth Standard Views are analogous to the First Standard View, the Third Standard View is analogous to the Second Standard View, and the two plausible assumptions are similarly analogous. More particularly, problems of consistency arise with respect to trade-offs between quality and duration within a life, for the very same reason that they arise between quality and number between lives; namely that for some such comparisons trade-offs between quality and duration seem appropriate, but for others they do not. As a result, in some cases one employs an additive aggregationist approach when comparing alternatives, while in others one does not.

After arguing for the theoretical inconsistency between the positions in question, and noting that it is by no means evident which position(s) should be given up, I next

illustrated the inconsistency using a variation of a powerful example developed by Stuart Rachels. I noted that virtually everyone accepts View One, the view that for any unpleasant experience, of any intensity and duration, it would be better to have that experience than one that was only a “little” less intense but much longer. I also noted that virtually everyone accepts View Three, the view that the mild discomfort of a mosquito bite would be better to two years of excruciating torture, no matter how long one lived, and no matter how long the discomfort of a mosquito bite persisted. But I noted that these views are incompatible with View Four, the view that “all-things-considered better than” is a transitive relation, given an empirical assumption, View Two, which implies that there could be a spectrum of alternative lives, ranging from one involving two years of excruciating torture to one involving many years of an extra mosquito bite, such that View One applies when comparing alternatives near each other on the spectrum, while View Three applies to the alternatives at the opposite ends of the spectrum. Views One through Four are inconsistent, and together they raise serious questions about the consistency of our views about aggregation even within a life. But Views One through Four are each deeply plausible, and giving any of them up will not be easy.

I next considered various ways of challenging my example. I granted that my example could be interpreted in various ways that would render it implausible. I also granted that, as stated, Views One and Three were liable to various objections. But I argued that my example could be formulated or interpreted, and Views One and Three suitably refined and limited in scope, so as to successfully avoid the worries raised.

For instance, in assessing View One's plausibility, I noted that my example assumes that there would not be a significance difference in the timing of when the unpleasant experiences occurred. Similarly, in assessing View Three's plausibility, my example assumes that there would not be any indirect side-effects from year after year of an extra mosquito bite. So, one should not assume that many years of an extra mosquito bite would eventually make one go crazy, or have the same psychological impact as torture.

My example also assumes that small differences in the intensity of pains would not be correlated with sudden changes of great normative significance. So, for example, it assumes that in each of my example's alternatives the lives being compared are worth living, all things considered, and involve an integrated self. Thus, even if one believes that View One is implausible as stated, because in some cases a small change in the intensity of pain might be the difference between a life worth living, and a life not worth living, or between an integrated self, and a disintegrated self, this does not show that it is implausible for the kinds of cases my example envisages. The point is that View One may need to be limited in scope, to accommodate such considerations, but so limited it remains deeply plausible, applies to a version of my example, and remains inconsistent with other deeply plausible positions.

Similarly, I granted, for the sake of argument, that View One might seem less plausible when comparing levels of discomfort that were extremely insignificant, or when comparing pains of duration that were extremely short. But, as above, this merely shows that View One might be limited in scope, and need to be revised so that it doesn't apply to such alternatives. This does not affect my argument, since the revised View One

would be deeply plausible, and would apply to a suitably rendered version of my example.

I next considered another objection to View Three, based on the view that if, as I grant, a nanosecond of prototorture would be better than n years of a mosquito bite, then q nanoseconds of prototorture must be better than $q*n$ years of a mosquito bite, where q consecutive nanoseconds of prototorture just is two years of intense torture. I claimed that such an objection embodies an additive aggregationist position that begs the question against View Three's intuitively plausible anti-additive aggregationist position. More importantly, I argued that such an objection invokes principles of decomposition, additive aggregation, and recombination that hold in the domain of arithmetic, but fail in other domains. For example, they fail wherever holism is plausible; that is, wherever the value of the whole is greater or less than the sum of the values of each part, considered separately. This, I suggested, may occur whenever there are interaction effects between a whole's parts.

The importance of interaction effects has long been recognized in normative domains, and I noted familiar examples of this from both aesthetics, involving taste and beauty, and ethics, involving proportional justice. Then, drawing on an analogy with both straw and snow, where it clearly makes a huge difference how a given total amount of straw or snow is distributed in space and time, I noted that it would make a huge difference, whether q nanoseconds of prototorture were grouped together, consecutively, for two straight years, or spread out over time, one nanosecond per n years, for $q*n$ years. Due to interaction effects, the first would rise to the level of extreme torture, and be of great normative significance, while due to the lack of interaction effects, the second

would be normatively insignificant. In sum, I showed that there is powerful reason to reject the principles of decomposition, additive aggregation, and recombination in the normative domain generally, and the pain sphere particularly. Correspondingly, there is powerful reason to reject the objection to the Third View based on such principles.

I next considered a proportionality argument against View Three. This argument claims that since a nanosecond of prototorture within a normal lifespan is insignificant, then two years of torture would also be insignificant, if it were embedded within a life whose length stood in the same proportion to two years as a normal lifespan stands to a nanosecond. I argued that this argument wrongly assumes that the reason a nanosecond of prototorture is insignificant is that it would constitute but a tiny fraction of a normal lifespan, whereas in fact the reason is that it wouldn't even be noticed! I noted that facts about perceptibility and human psychology explain why a nanosecond of prototorture wouldn't matter even within a life of much shorter duration, and why two years of torture would matter even in a life of extraordinary length. In addition, I argued that the proportionality argument should be rejected because it implausibly entails that it would be completely trivial if God were to subject us to extreme torture for billions and billions of years, as long as he also provided us with an earthly situation for a sufficiently long time that the proportion between the two equaled the proportion between a nanosecond and a normal lifespan.

Next, I considered a different objection to View Three, based on the claim that we can't trust our intuitions to accurately grasp the overall significance of small amounts spread out over innumerable years. Thus, it is claimed that we can't intuitively grasp how bad it would be to experience an extra mosquito bite for billions or zillions of years,

and hence that we can't trust our intuition that such an outcome would be better than one involving two years of intense torture.

I granted that our intuitions are often suspect in cases involving small amounts or large numbers. However, I suggested that just as one can get from a short man to a tall man, or from one end of the visual spectrum to another, in relatively few steps, so it is possible one could get from one end of the pain spectrum to the other in relatively few steps. If this were the case, then my example might only require one to compare a life with two years of torture to one with a few hundred, or maybe thousands of, years of a mosquito bite or, if one changed the scale, between two months of torture and a few hundred, or maybe thousands of, months of a mosquito bite. I suggested that while such comparisons are not typical, they are not unimaginable.

Moreover, even if the example from torture to mosquito bites were suspect, I suggested that other examples could be presented that would range over clearly comprehensible periods of time, and yet raise similar worries as those of my original example. I presented four such examples, the first involving a series of uncomfortable experiences ranging from extreme torture for two days to the moderate discomfort of a limp lasting for 2048 days, and the second, third, and fourth examples involving, respectively, different scenarios of inconvenience ranging from random interruption of electrical service 20 minutes per time 9 times per day for 24 straight months to 20 minutes per time .53 times per day for 1065 months, from driving delays of 4.5 hours per day for 24 months to .27 hours per day for 1065 months, and from 48 garbage pick ups missed per year for 24 months to 3.1 missed pick ups per year for 1065 months. I suggested that in each of these cases, people might intuitively judge the first alternative to

be better than the second, the second to the third, and so on, yet they might not judge the first alternative to be better than the last. If this is right, it suggests that there is, indeed, a conflict in our intuitions and beliefs that my original example was intended to highlight; a conflict that cannot simply be dismissed as an artifact of wild science-fiction type cases involving unimaginably large numbers for which our intuitions were not developed and cannot be trusted.

I next noted that View Three is an example of particularly strong anti-additive position, in contending that no matter how long a mild discomfort lasted, it would be better than two years of intense torture. While I have argued for the plausibility of positions like View Three, I noted that problems of the sort I am raising can arise with much weaker anti-additive positions. Specifically, as long as one believes that certain kinds of alternatives can be rightly judged mainly on the basis of sum totals, while others can only be judged by giving significant weight to other factors besides sum totals, like the pattern of distribution, one will be faced with a dilemma. One will have to either reject the first belief, the second belief, the transitivity of the “all-things-considered better than” relation, or the empirical claim that there can be spectrums of cases such that the first belief is relevant to comparing all the alternatives that are near each other on such spectrums, while the second belief is relevant to comparing alternatives at the opposite ends of such spectrums.

Finally, I discussed the view that we might simply appeal to expected utility theory to resolve the dilemmas I have been raising. I acknowledged that expected utility has proved to be a useful and illuminating tool, but noted that if expected utility theory is compatible with anti-additive aggregation, then it won't be able to help us, and if it is

incompatible with anti-additive aggregation, then it is subject to all the worries that have been raised about additive aggregationist positions. If we truly want to understand the nature of moral ideals and practical reasoning, we must confront the worries I am raising head on, and decide clearly and forthrightly what should be said about them. We cannot simply avoid my worries, by appealing to a theory whose very foundations may be challenged by the worries I am raising.

In sum, problems of aggregation arise within lives, as well as between lives. Specifically, within lives, as well as between lives, there is an inconsistency between two deeply plausible and widely-held views about how to compare alternative outcomes, an apparent empirical fact, and the transitivity of the “all-things-considered better than” relation. Something must be given up, but to this point, at least, it is by no means clear what that should be.

In this chapter, I have considered various possible responses to my spectrum cases. Other important responses still need to be considered. One of the most important of these is the claim that my spectrum argument can be understood, and dismissed, as yet another example of the well-known, and fallacious, Sorites Paradox. Another challenges the transitivity of the “all-things-considered better than” relation. I consider these, and other objections, in later chapters. In doing so, a clearer understanding of the nature of spectrum cases and their implications will emerge.

¹ Cite to “A Continuum Argument for Intransitivity”

² This might be a place for a footnote with all sorts of historical and contemporary references, Plato, Hegel, Brentano, Moore, Kamm, Korsgaard, Kagan, Hurka, etc. Check on.

³ Insert cite to classic example from Plato—The Protagoras?

⁴ Fred Feldman has sent me a long list of such explanations, many of which are quite ingenious and sophisticated. I don't find these explanations ultimately convincing, but they are worthy of careful consideration, and their appeal is certainly understandable.

⁵ See chapter three of Derek Parfit's Reasons and Persons, Oxford University Press, 1984, also the striking results in Kahneman, Slovic, and Tversky's Judgments under Uncertainty: Heuristics and Biases.

⁶ I am grateful to Peter Unger for this example.

⁷ Note, I do not believe that any amount of torture, no matter how short, must be worse than any amount of a hangnail, no matter how long. Nor does my argument commit me to this. It might well be that two seconds worth of torture would be better than many years of a hangnail, and it is almost certainly true that two nanoseconds of torture would be better than many years of a hangnail. Correspondingly, there may be perfectly transitive orderings from outcomes involving very short amounts of torture to one's involving very long amounts of a hangnail. I don't deny this, nor do I need to. All I need to establish is one case of intransitivity, and for this it is enough if torture that lasts long enough, would be worse than a hangnail of vastly longer (and even perhaps any) duration. My model for this view is noted later in the text, and note 16.

⁸ Readers who still think it will take too long to get from extreme torture to a hangnail, may want to consider variations of my example. Here is one some find convincing. Imagine a scale of "uncomfortable" experiences, ranging from a level of 1, a hangnail, to a level of 100, extreme torture. Suppose a moderately uncomfortable limp is an 11--significantly worse than a mere hangnail, but not nearly as bad as extreme torture. Start with a choice between A, level 100 discomfort for 2 days, or B, level 80 discomfort for 4 days. B's discomfort is 20% less than A's, but lasts twice as long. Many think A better than B. Next compare B to C, where C stands to B as B stands to A. C is 20% less intense, level 64, but lasts twice as long, 8 days. Again, many think B better than C. The tenth choice would be between J, discomfort of level 13.4 for 1024 days, and K, discomfort of level 11 for 2048 days. Again, many would think J better than K. Given these rankings, transitivity entails that A is better than K. Is it? Given our assumptions, A involves extreme torture for 2 days, K involves a moderately uncomfortable limp for 2048 days, or 5.6 years. Which would you choose for your child--that for the next 5.6 years they suffer an uncomfortable limp, or for 2 days they suffer extreme torture? (Here, as always, we assume there are no relevant side-effects. Other children to not mercilessly ridicule your child as a "gimp," a pill, or hypnosis, removes all memories of the torture, etc.) I would choose the limp over the torture for someone I loved. And I would do so because I think it the better alternative. But I would also choose A over B, B over C, ..., and J over K.

This example involves an imaginable range of pains, and finite graspable times. Here, as before, I reject transitivity.