

*Asymmetries In Value*

THOMAS HURKA  
University of Toronto

Values typically come in pairs. Most obviously, there are the pairs of an intrinsic good and its contrasting intrinsic evil, such as pleasure and pain, virtue and vice, and desert and undesert, or getting what one deserves and getting its opposite. But in more complex cases there can be contrasting pairs with the same value. Thus, virtue has the positive form of benevolent pleasure in another's pleasure and the negative form of compassionate pain for his pain, while desert has the positive form of happiness for the virtuous and the negative form of pain for the vicious.

Of each pair we can ask how its elements relate to each other, and the simplest answer is that they do so symmetrically, so that, for example, a pleasure of a given intensity is exactly as good as a pain of that intensity is evil, or benevolence exactly as great a virtue as compassion. But there is no necessity for this. Values can equally well be asymmetrically related, and in several ways. In this paper I ask, of a series of pairs of value, when asymmetries between their elements are plausible, what bases these asymmetries have, and whether there are any patterns among them. I start with the simplest case, that of pleasure and pain.

Utilitarians typically treat these values as symmetrical, so a given quantity of pleasure exactly cancels the disvalue of an equal quantity of pain; this is Jeremy Bentham's view and is also suggested by Henry Sidgwick.<sup>1</sup> But G.E. Moore disagrees. In *Principia Ethica* he says that while pleasure has "at most some slight intrinsic value," pain is "a great evil," adding that "[t]he study of Ethics would, no doubt, be far more simple, . . . if . . . pain were an evil of exactly the same magnitude as pleasure is a good; but we have no reason whatever to believe that the Universe is such that ethical truths must display this kind of symmetry."<sup>2</sup> A similar view has recently been defended by Jamie Mayerfeld. One of his two main claims about happiness and suffering is that


it is more important to prevent suffering than to promote happiness, because “suffering is more bad than happiness is good.”<sup>3</sup> And the view he shares with Moore is intuitively appealing. If one can either relieve one person’s intense pain or give a slightly more intense pleasure to another person, many will say it is right to relieve the first person’s pain. Or imagine first a world containing only intense mindless pleasure like that of the deltas and epsilons in *Brave New World*. This may be a good world, and better than if there were nothing, but it is surely not very good. Now imagine a world containing only intense physical pain. This is a very bad world, and vastly worse than nothing.<sup>4</sup>

The Moore-Mayerfeld view needs to be distinguished from two others. One holds that pleasure is nothing real but only the absence of pain; if we think it has positive qualities, we are only being fooled by the transition from a greater to a lesser pain. This view has been asserted by Plato, Epicurus, and Arthur Schopenhauer, but I will assume that it is false and that the pleasure of eating an oyster or of passionate love involves more than just feeling no pain. The second view holds that pleasure, though real, has no value at all. This is the view of negative utilitarianism, on which our only moral duty is to prevent pain. But negative utilitarianism implies that if we could either bring about a world in which billions of people are ecstatically happy but one person suffers a brief toothache or bring about nothing, we should bring about nothing. The Moore-Mayerfeld view avoids this absurd implication by giving happiness some positive value, so enough of it will outweigh minor pain. But exactly what does the view say?

The symmetry view says that a pleasure of a given intensity is always exactly as good as a pain of that intensity is evil. Its simplest version is represented in Fig. 1, whose vertical axis measures the value of a pleasure or pain just as a pleasure or pain and whose horizontal axis measures its intensity, and which displays the relation between the two as a single straight line. (A pleasure or pain can also have non-hedonic value, for example if it is virtuous and on that basis good or undeserved and on that basis evil. Both Fig. 1 and our current discussion abstract from these possibilities and consider the value of hedonic states only as hedonic states.) Now, the symmetry view will be false so long as there is one case where a pain is more evil than its corresponding pleasure is good, but I take Moore and Mayerfeld to assert a stronger *pairwise asymmetry* thesis:

For any intensity  $n$ , a pain of intensity  $n$  is more evil than a pleasure of intensity  $n$  is good.

And the simplest view that captures this thesis is represented in Fig. 2, which still has straight lines but with different slopes above and below the horizontal axis. If the slope for pain is twice as steep as for pleasure, a pain of a given intensity is always twice as evil, considered just hedonically, as a pleasure of that intensity is good. And the basis of the pairwise asymmetry



**Figure 1.** Value of pleasure or pain as pleasure or pain.


is a claim about the values of increments of pleasure and pain that I call the *marginal-value* claim:

For any intensity  $n$ , the difference in evil between  $n$  and  $n + 1$  units of pain is greater than the difference in goodness between  $n$  and  $n + 1$  units of pleasure.

Given a fixed initial intensity, the evil of an additional unit of pain is always greater than the goodness of an additional unit of pleasure.

There is a skeptical objection to this view. Moore and Mayerfeld assume, as Bentham and Sidgwick also do, that we can compare the intensities of pleasures and pains independently of assessing their values. But an objector may challenge this assumption, saying the claim that a pain is more intense than some pleasure merely says the pain is more evil than the pleasure is good, without pointing to some independent psychological fact that makes it so. There is no such fact, she may argue, and thus no possibility of non-evaluative comparisons of pleasures and pains. If so, the difference between symmetry and asymmetry views about them disappears.<sup>5</sup>

This objection raises a more general issue. To formulate a pairwise asymmetry about any pair of values we must be able to compare these values



**Figure 2.** Value of pleasure or pain as pleasure or pain.

non-evaluatively, or identify instances of them as equal in a way that is neutral about their comparative worth. Only then can their relative values be a further issue.

For some pairs of values this non-evaluative comparison is unproblematic. To compare virtue and vice, for example, we must be able to equate a compassionate desire of intensity  $n$  to relieve a given pain with a malicious desire of the same intensity to inflict that pain, and we can do so using the familiar economists' measures of intensity of desire. For desert and undesert, we must equate a virtuous person's enjoying pleasure of value  $n$  with his suffering pain of disvalue  $n$ , which we can do using whatever our method is for comparing the values of hedonic states. But in other cases such comparisons do not seem possible. The question whether beauty of degree  $n$  is more good than ugliness of degree  $n$  is evil seems meaningless, since there are no comparisons of beauty and ugliness independent of comparisons of their value. Likewise for equality and inequality, or, since equality does not admit of degrees, for increases in inequality at high and low levels of inequality.<sup>6</sup> But what about pleasure and pain? Do they allow non-evaluative comparisons?


That it is for this pair that symmetry issues have been most often discussed suggests that many philosophers think they do allow such comparisons, and I

share this view.<sup>7</sup> Surely few would extend the skeptical objection to pleasures and pains taken separately. The claim that the pain of being tortured is more intense than that of being lightly pinched does not say only that the former is worse; it points to an independent psychological fact that makes it so. Likewise for the claim that one of two pleasures is more intense. It does not follow that we can make non-evaluative comparisons across the pleasure-pain divide. We can compare temperatures with temperatures and weights with weights, but it is senseless to ask whether a fire is hotter than a stone is heavy. But pleasure and pain are not like temperature and weight. They are not completely unrelated states but states of the same kind, namely hedonic states, and they are such because the painfulness of the one is the contrary of the other's pleasantness. This does not prove that the two can be non-evaluatively compared, but it does suggest that this may be possible and it seems to be something we do.

If asked whether the pain of being tortured is more intense than the pleasure of eating a jelly bean, surely we can say yes and not mean just that the pain is more evil; the same is true if we are asked about the pleasure of orgasm and pain of being pinched. And we can give evidence for these claims. Within the categories of pleasure and pain, mild feelings make only minimal demands on our attention. We can experience the pleasure of eating a jellybean or the pain of being pinched while simultaneously feeling many other sensations, engaging in other activities, and so on. As a hedonic state becomes more intense, however, it becomes more importunate, drawing more attention to itself and starting to disrupt other activities; this can happen with orgasm on the one side and torture on the other. And we can use this fact to make pleasure-pain comparisons. If the pain of being tortured forces itself more on our attention than the pleasure of eating a jellybean, for example, that is evidence that the pain is more intense. I am not suggesting that the intensity of a pleasure or pain just is the demand it makes on attention; the latter is only an expectable effect of the former. But as an effect it can be used to compare the two in a non-evaluative way.<sup>8</sup>


I do not claim that these points decisively answer the skeptical objection, or expect them to persuade every reader. But for the larger purposes of this paper it is not essential that they do so. The pleasure-pain pair is just one of several for which value-asymmetries are possible, and if I discuss it first it is mainly because it is the simplest, making the relevant types of asymmetry easiest to see. Those still moved by the objection should therefore suspend it temporarily, and let the pleasure-pain case illustrate, if only hypothetically, possibilities that arise in more complex ways for other, more readily comparable values.

Though it affirms a pairwise asymmetry, the view in Fig. 2 does not capture everything either Moore or Mayerfeld says. To begin with Mayerfeld, he supplements his first claim about happiness and suffering with the further claim that it is disproportionately more important to relieve more intense



**Figure 3.** Value of pleasure or pain as pleasure or pain.

pains, because they are disproportionately more evil. If we can reduce one person's suffering from 10 units to 9 or another's from 3 units to 1, he argues, we should do the former; though the resulting reduction in pain will be smaller, the reduction in hedonic evil will be greater.<sup>9</sup> This second claim of Mayerfeld's is also intuitively appealing, but it requires that the evil of increments of pain not be constant, as in Fig. 2, but increase. And this in turn requires the straight line below the horizontal axis to be replaced by a curve whose slope gets steeper as one moves to the left, as in Fig. 3. There is no hint of this second claim in Moore's discussion, but he implies a complementary claim about pleasure that Mayerfeld never mentions. Moore's view that pleasure "has at most some slight intrinsic value" could not be true if increments of pleasure had constant value, as in Fig. 2, since then a sufficiently intense pleasure could have as much value as one likes. His view requires there to be an upper bound on the value of pleasure, which in turn requires the goodness of increments of pleasure to diminish, as on Mayerfeld's view the evil of increments of pain increases. And in a pre-*Principia* article Moore explicitly mentions the possibility of diminishing marginal value for pleasure.<sup>10</sup> If it is added to Mayerfeld's second claim, the result is the view represented in Fig. 4, which has a smooth curve running from the bottom-left quadrant to



**Figure 4.** Value of pleasure or pain as pleasure or pain.

the top right, with the curve's slope getting progressively shallower as it rises to the right. This view still captures the pairwise asymmetry on a marginal-value basis, and in fact makes that asymmetry stronger. But in doing so it also expresses a version of what Derek Parfit calls the "priority view," which always gives some priority to improving the condition of the worse-off, interpreted here in terms of pleasure and pain. Whenever one person enjoys less pleasure or suffers more pain than another, it is better to give a fixed benefit to the first person, because the value of that benefit will be greater. The priority view is usually discussed in connection with egalitarian views about distributive justice, where it is contrasted with views that value the relation of equality as such. Here we have arrived at it by combining different attractive claims about pleasure-pain asymmetry.<sup>11</sup>

Though still affirming a pairwise asymmetry on a marginal-value basis, the priority view in Fig. 4 implies a second asymmetry thesis, which I call the *limit asymmetry* thesis:

There is some intensity  $n$  such that a pain of intensity  $n$  is more evil than any pleasure could be good.

Because it places an upper bound on the goodness of pleasure but none on the evil of pain, the priority view allows that an intense pain can exceed in disvalue any possible value in a pleasure. This is a second way in which pain can be a greater evil than pleasure is a good: not only are its instances always more evil in equal-intensity comparisons, but they can reach heights of evil greater than any goodness possible for instances of pleasure.<sup>12</sup>

The priority view in Fig. 4 supplements each of Moore's and Mayerfeld's views with a claim its author does not consider – in Moore's case about the increasing marginal evil of pain, in Mayerfeld's about the diminishing marginal goodness of pleasure – and it may give the best possible grounding of the pleasure-pain asymmetry both embrace. But the grounding of this asymmetry remains the marginal-value claim, and there is an alternative possibility: one can also generate a pairwise asymmetry while rejecting the marginal-value claim and holding that the values of increments of pleasure and pain are always the same. I will now explore this second strategy; though not plausible for pleasure and pain, it will prove attractive for other values.


Return to the symmetry view in Fig. 1. The second way to ground a pairwise (though not a limit) asymmetry is to shift the single line down the graph, so it cuts the vertical axis below the origin, as in Fig. 5. The resulting view treats increments of pleasure and pain as equal in value, thereby rejecting the marginal-value claim and treating the duties to relieve pain and promote pleasure as equal in strength. But it still generates a pairwise asymmetry. If the line is shifted down, say, 2 units, then a pain of 4 units of intensity has value  $-6$ , while a pleasure of 4 units of intensity has value  $+2$ . And the asymmetry's basis is now a different, *downshift* claim:

For any intensity  $n$ , the goodness of a pleasure or evil of a pain of intensity  $n$  is  $n - a$ , where  $a$  is some positive number.

An especially salient implication of this claim is that the point of hedonic neutrality, which had the value 0 in Fig. 1, now has the value  $-a$ , because the line cuts the vertical axis at  $-a$ . But the downshift makes a similar adjustment to the value of every hedonic state, with the result that a pain of a given intensity is always  $2a$  units more evil than an equally intense pleasure is good (at least when the pleasure is sufficiently intense to be good).<sup>13</sup>

This view, originally proposed by Gregory Kavka and generating what is now called “critical-level” utilitarianism,<sup>14</sup> is sometimes said to avoid Parfit's “repugnant conclusion” objection to total utilitarianism. If total utilitarianism is correct, then for any world in which billions of people enjoy ecstatic happiness, there is another world that would be better even though in it people's lives are barely above hedonic neutrality; if there are enough such lives, the sum of goodness they contain will be greater than in the first world.<sup>15</sup> The critical-level view avoids this implication by giving lives above neutrality but below the critical level  $a$  negative value, but its success here is surely

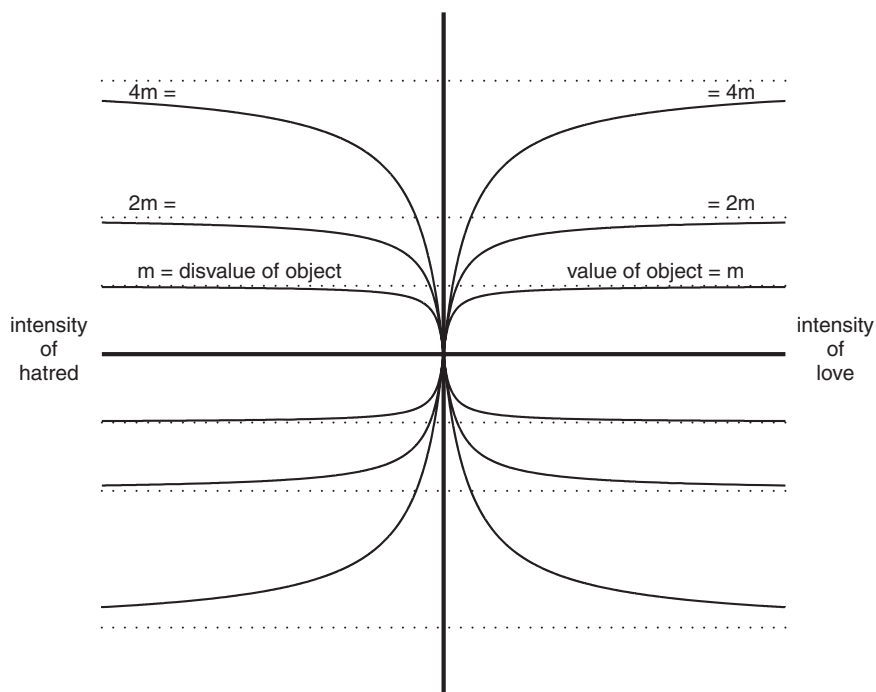




**Figure 5.** Value of pleasure or pain as pleasure or pain.

limited. Unless it sets the critical level implausibly high, it implies, if not quite the original repugnant conclusion, then one very similar to it, in which lives in the second world are only slightly above neutrality. And it faces an even more decisive objection. Imagine a hedonically horrible world, in which billions of people suffer excruciating pain. If the critical-level view is correct, there is another world that would be worse even though everyone in it enjoys positive happiness. If their lives are below the critical level  $a$ , their lives have negative value; and given enough such lives, the sum of evil in them is greater than in the horrible world. But surely it is absurd to say that a world containing only happiness could be worse than one in which everyone suffers excruciating pain.<sup>16</sup> And in fact these difficulties only reflect the deeper fact that the critical-level view has no philosophically credible rationale: specific objections aside, what reason is there to believe that positive happiness is other than positively good?

I conclude that the downshift is not plausible for the hedonic values of pleasure and pain, but it is so for other intrinsic values such as virtue and vice. My understanding of virtue differs from that found in ancient philosophy or contemporary virtue ethics, though it was shared by philosophers around the turn of the 20<sup>th</sup> century such as Hastings Rashdall, Franz Brentano,




**Figure 6.** Value of attitude as virtuous or vicious.

Moore, and W.D. Ross. It holds that virtue consists largely in morally appropriate attitudes to independently given goods and evils, and vice in inappropriate attitudes to them. If another person's pleasure is good, then the positive attitude of loving, or benevolently desiring, pursuing, and taking pleasure in, her pleasure for itself is virtuous and on that basis intrinsically good, while the negative attitude of enviously hating and wanting to destroy it is vicious and evil. Conversely, if another's pain is evil, the negative attitude of being compassionately pained by and wanting to relieve it is virtuous, while maliciously desiring or taking pleasure in it is evil. To discuss issues about value-symmetry we need to know which features of an attitude determine its degree of virtue or vice, and here I am guided by two ideas. One is that there is an upper bound on the value of any virtuous or vicious attitude, so the attitude is always less good or evil than its object. Thus, my compassion for your pain is good, but not as good as your pain is evil. The second idea is that the best division of virtuous concern between two or more objects is proportioned to their degrees of value, so if  $x$  is twice as good as  $y$ , it is most virtuous to be twice as pleased by  $x$  as by  $y$ . A view that captures these ideas is represented in Fig. 6, where the horizontal axis measures the intensity of love or hate for an object and each curve shows how, given a

fixed value in its object, the value of an attitude as virtuous or vicious varies with its intensity. (Some virtuous or vicious attitudes may also have other values. For example, compassionate pain at another's pain may be good as compassionate but evil as pain, with its value on balance depending on how the two weigh against each other. Fig. 6 abstracts from these other values and considers the values of attitudes only as instances of virtue or vice.) The curves for attitudes to good objects run from the bottom left to the top right, since hatred of these objects is evil and love of them good, while those for evil objects run from the top left to bottom right. And the shapes of the curves satisfy the two demands of boundedness and proportionality.<sup>17</sup>

The view in Fig. 6 treats virtue and vice symmetrically, so that for any object a virtuous attitude to it of intensity  $n$  is exactly as good as a vicious attitude of intensity  $n$  to it is evil. It does so in part because its curves pass through the origin, so the neutral attitude of indifference to a good or evil always has zero value. But this is morally questionable, and in a way that makes a downshift positively attractive. Intuitively, indifference to another's pain is not just not good but evil; it is callousness, and callousness is a vice rather than just the absence of a virtue. Similarly, having no desire for achievable goods is sloth or apathy, which is likewise a vice rather than merely not a virtue. So here there is a positive reason to shift the curves down so they cut the vertical axis below the origin, as in Fig. 7. More specifically, there is a reason to shift the curves for attitudes to greater goods and evils further down than those for attitudes to lesser goods and evils, so indifference to the former is a greater evil and the minimal intensities of concern for them needed for positive value are likewise greater. And making this downshift results in an asymmetry whereby vice is a greater evil than virtue is a good in both pairwise and limit ways. Go equal distances to the left and right of the origin, and the distance down to the vice portion of a given curve is always greater than the distance up to its virtue portion. This makes a malicious desire of intensity  $n$  to inflict a given pain more evil than a compassionate desire of intensity  $n$  to relieve the pain is good. Given paired instances of malice and compassion, the former is more vicious than the latter is virtuous, or more evil than the latter is good. In addition, the lower bound on the value of a vicious attitude to an object is always further below the horizontal axis than the upper bound on the value of a virtuous attitude to it is above the axis, so a malicious desire to inflict a given pain can be more evil than any compassionate desire to relieve it is good. Supplementing the pairwise asymmetry, then, is a further limit asymmetry.

If the downshift in Fig. 7 is uniform along each curve (as it must be to satisfy the proportionality condition), it makes not only indifference but also very weak appropriate attitudes, such as very mild compassion for great pain, intrinsically evil. This is also intuitively appealing; it seems right that feeling only mild distress at, say, the Holocaust is not just not good but evil. That too is a form, though a lesser one, of callousness. And this implication




**Figure 7.** Value of attitude as virtuous or vicious.

can be given a positive rationale: if what is evil is morally inappropriate attitudes, then one way an attitude can be evil is by being, while properly oriented, inappropriately weak for its object. For both these reasons, the downshift for virtue and vice does not invite a decisive objection like the one against critical-level utilitarianism. A world in which a huge number of people feel only mild distress at the Holocaust can indeed be worse than one in which a much smaller number take positive pleasure in it; the more extensive callousness in the first world can be more evil than the less extensive malice in the second.

In Fig. 7 the asymmetry depends entirely on the downshift and not at all on a marginal-value claim; because the curves retain their symmetrical shapes from Fig. 6, the value of equivalent increments of virtue and vice remains the same. For this reason, the view in Fig. 7 does not support an analogue of Mayerfeld's claim that the duty to relieve suffering is stronger than the duty to promote happiness. On the contrary, though it makes a vicious attitude to a given object more evil than the corresponding virtuous attitude is good, it holds that the duties to reduce vice and promote virtue are equally strong. The strengths of these duties depend only on the slopes of the curves above and below the point representing indifference; if those

slopes remain unchanged, as they do in Fig. 7, so do the duties' weights. Now one could supplement the view in Fig. 7 with a marginal-value claim, giving the curves a steeper slope below the point of indifference and pushing the lower bounds even further down. This would strengthen both the pairwise and limit asymmetries, and make the duty to reduce vice stronger than the duty to promote virtue. But though I do not have a firm view about this possibility, I do not find it that attractive. I see no reason why intensifying virtuous love for an object should be less good than intensifying vicious hatred of it is evil, and I will therefore retain the simpler view in Fig. 7. Doing so makes for an interesting contrast with the case of pleasure and pain. For those hedonic values the most attractive asymmetry rests only on the marginal-value claim and the alternative downshift basis is not plausible, while for virtue and vice the best asymmetry depends only on the downshift and not on marginal values. Of the two possible bases of asymmetry in value, only one suits the one pair and only the other suits the other. And this contrast can be extended, since there is a further pair of values for which both grounds of asymmetry are plausible.


This pair contains the good of desert, which I will understand more specifically as moral desert, or getting what one deserves on the basis of one's moral qualities, and its contrary undesert, or getting the opposite of what one deserves. Given these values, it is good if virtuous people enjoy pleasure or vicious people suffer pain, and bad if the virtuous suffer or the vicious are happy. These desert-values parallel virtue and vice in many ways, for example, by making a positive response to a positive value, this time the reward of happiness for the good of virtue, intrinsically good, and a negative response to that same value evil. They are also governed by similar ideas of boundedness and proportionality. The goodness of rewarding virtue or punishing vice is always less than the goodness of the virtue or evil of the vice; it is not better to have vice and its deserved punishment than to have no vice at all. And the best division of rewards or punishments among people is proportioned to their degrees of virtue or vice, so that, for example, those who are twice as virtuous enjoy twice as much pleasure. But there is an important difference between the two pairs. In Figs. 6 and 7 the slopes of the virtue curves, while diminishing, always remain positive, so a more intense love of a good is always intrinsically better. A more intense love may be instrumentally worse if it prevents one from having other, more valuable attitudes, but in itself it is always preferable. But it would not be plausible to take a similar view about desert. In particular, it would not be plausible to say that if a person is vicious, it is always better from the point of view of desert if he suffers more pain. Desert-values demand a different structure, whereby for any degree of virtue or vice there is a specific amount of pleasure or pain that it ideally deserves and that has most value, and where amounts above or below that ideal are less good and can even be evil. A desert view with this "peak" structure is represented in Fig. 8.



**Figure 8.** Value of pleasure or pain as deserved or undeserved.

Here the horizontal axis measures the amount of pleasure or pain a person experiences while the vertical axis measures how good or evil this is just as a matter of desert. (The resulting desert-values must again be weighed against other values, such as hedonic ones, to determine the value on balance of his enjoying the pleasure or suffering the pain. Thus, the pain of a vicious person, while good as deserved, is bad as pain.) The topmost points on the different curves represent the ideal rewards and punishments for different degrees of virtue and vice, and the curves' shapes again incorporate the demands of proportionality.<sup>18</sup>

Fig. 8 already contains the first or pairwise asymmetry, on a marginal-value basis. It follows from the proportionality condition that the downward slopes on the inside of a peak, the side closer to the vertical axis, get steeper until they cut the vertical axis. Below that point they have to get shallower, but in Fig. 8 the rate of change in slope below the cut is much slower than the rate of change above it, so they stay close to the vertical longer. This implies that if we go equal distances to the left and right of the origin, the distance up to the desert portion of a given curve is always less than the distance down to its undesert portion. If a virtuous person ideally deserves, say, 4 units of happiness, his enjoying that happiness is less good from the point of view of desert than his suffering 4 units of pain is evil. And the view in



**Figure 9.** Value of pleasure or pain as deserved or undeserved.

Fig. 8 can be supplemented by a downshift, if we shift the curves so they cut the vertical axis below the origin. In the case of virtue, the neutral response of being indifferent to goods or evils was plausibly not just not good but evil. Here it is likewise plausible that a virtuous person's getting no reward but experiencing neither pleasure nor pain is positively evil in desert terms, as is a vicious person's suffering no punishment. Doing nothing to meet the demands of justice, or giving no response to merit or demerit, is a positive injustice. A graph that incorporates this downshift is given in Fig. 9; with two bases for the pairwise asymmetry, it makes that asymmetry even stronger. In a book on punishment A.C. Ewing writes, "When I look at the two, injustice in punishment seems to me a *very much greater intrinsic* evil than justice is a good, especially if the injustice consists in punishing somebody for an offence of which he is not guilty or in excessive severity."<sup>19</sup> Fig. 9 gives two different grounds for this intuitively appealing claim.

As in the case of virtue, the downshift-based asymmetry does not make the duty to prevent undesert stronger than the duty to promote desert, because it does not on its own change the curves' slopes. But the marginal-value-based asymmetry does have this effect, and in so doing connects with a familiar thesis about criminal punishment, namely that the state should be more

concerned about not punishing the innocent than about punishing every one of the guilty. As Sir William Blackstone put it, “Better that ten guilty persons escape than that one innocent suffer.”<sup>20</sup> This thesis can be grounded in deontological principles, if the state has a stronger duty not to *do* what punishes the innocent than not to *allow* the guilty to go free. But it can also be supported by claims about value if the incremental evil of punishing the innocent is greater than the incremental good of punishing the guilty. Then, even apart from deontological considerations, punishing the innocent does more to violate the demands of justice than punishing the guilty does to satisfy them. The criminal desert relevant to Blackstone’s thesis differs in several respects from moral desert, for example, by concerning specific criminal acts rather than overall moral virtue. But it is still governed by a proportionality condition, so the best division of punishments among crimes is proportioned to their degrees of seriousness, and it still supports a downshift, since it remains plausible that a criminal’s escaping punishment is not just not good but evil. An attractive axiological view of criminal desert therefore supports Blackstone’s thesis and the legal procedures it justifies, such as the presumption of innocence in criminal trials. It also supports a subtle consequentialist account of the morality of punishment proposed by Ewing. In this account considerations of desert play only a modest positive role. Though punishing wrongdoers is an intrinsic good, it is a comparatively minor one and does less to justify legal punishment or to fix its optimal severity than do considerations of deterrence and moral reform.<sup>21</sup> But desert is much more important on the account’s negative side, which says punishment may normally be inflicted only on wrongdoers and not on the innocent. Because punishing the innocent is a great intrinsic evil, retribution counts strongly against such punishment and will often forbid it when deterrence and reform do not. So the asymmetry between desert and undesert leads to an asymmetry of justifying roles: retribution counts only modestly in favour of punishing the guilty, which is primarily justified on other grounds, but does much more to forbid punishing the innocent.

Figs. 8 and 9 also contain the second or limit asymmetry. If there is an upper bound on the goodness of desert but no lower bound on the evil of undesert (as there cannot be given proportionality), then whatever a person’s degree of virtue or vice, his getting the opposite of what he deserves can be more evil than his getting what he deserves can be good. But in this case the limit asymmetry has a more specific basis. The value of desert is not just bounded but can be fully achieved, if a person gets exactly the happiness or suffering he merits. The relevant desert-goodness is then complete, in the sense that it cannot be improved. This was not the case with pleasure or virtue; since their curves never reached a peak, their instances could always in principle be made better. While those values are not fully achievable, desert is. And this provides a distinctive basis for asymmetry that is also found in a good like equality, understood as a relation between people’s



levels of, say, happiness that is valued as a relation, and is opposed to an evil of inequality. Like desert, equality can be fully achieved, when people's levels of happiness are exactly the same. But its contrary, inequality, can in principle increase without limit. Assuming a finite value for equality, therefore, there can be unequal distributions among people that are more evil than any equal distribution among them can be good, so, as in the case of desert, injustice is in the limit sense a greater evil than justice is a good.<sup>22</sup> These asymmetries may echo an ancient idea due to the Pythagoreans and discussed sympathetically by Plato and Aristotle, an idea that associates good with "limit" and evil with the "unlimited." As Aristotle says, "evil belongs to the class of the unlimited, . . . and good to that of the limited."<sup>23</sup> The thought here is not just that desert and equality are subject to limits, in that their value has an upper bound. It is that they themselves involve a limit, or a mathematical relation that can be completely achieved. On the Pythagorean view this makes desert and equality good, in contrast to opposites that can increase without limit. But in an implication the Pythagoreans may have found less welcome, it also makes their value less, in the limit sense, than their opposing intrinsic evils.

Let me summarize. I have examined three main pairs of values and in each found either or both of a pairwise and a limit asymmetry. The bases of the asymmetries have been different: for pleasure and pain just the marginal-value claim, for virtue and vice just the downshift, and for desert and undesert both. But there is an obvious pattern to the asymmetries. In each pair it is the evil that is greater than the good, or the evil that is morally more potent. Contrary good-favouring asymmetries can certainly be formulated, but they do not have much intuitive appeal. Surely no one would say it is more important to increase the happiness of the very happy than to decrease the suffering of the miserable, or that the neutral attitude of indifference to another's pain is positively good. In all these pairs the intuitive pressure is to accentuate the negative. Nor should this be surprising, since there are other parts of morality where negative considerations are more potent.

Consider a deontological morality that sometimes forbids acts that have the best outcome, and imagine that, like consequentialism, it formulates its principles by reference to good and evil states of affairs rather than to some alleged Kantian value in persons. This morality may use either or both of the distinctions between doing and allowing and between intending and merely foreseeing, but these distinctions are engaged only by the production of evils and not by the production of goods. If I can relieve five people's pain by directly causing another person pain, a deontological morality may say it is wrong for me to do so. But if I can either directly cause one person's pleasure or allow someone else to cause pleasure for five, the same morality will say I should prefer the pleasure of the five. While it is more objectionable to actively cause evil, it is not usually more creditable to actively cause good.

A similar point applies to intention and foresight. It may be deontologically wrong to intend one person's pain as a means to relieving five other people's, but right to prefer an act that merely foresees pleasure for five to an act that intends it for one.<sup>24</sup> So in a deontological context evil is morally more potent than good because it engages the central deontological distinctions where good does not. The asymmetry here is not exactly analogous to the ones we have discussed, because it is not between evil and good as such; it is between changes that make a situation worse and changes that make it better. If to prevent five people from having their happiness reduced I must directly reduce the happiness of one, a deontological morality may forbid this act even though it involves only the good of happiness, and likewise if I intend the reduction. So what differentially engages the deontological distinctions is not good and evil as such, but changes in the values of states of affairs for the worse and for the better. Nonetheless, there is a fundamental asymmetry in these moralities that makes a kind of negative effect more potent than its corresponding good one.

The pattern whereby in three pairs of values the evil is greater than the good has implications for several traditional philosophical issues. Consider the 19<sup>th</sup>-century debate between optimists and pessimists about whether the world is on balance good or on balance evil. The asymmetries we have identified strengthen the case for pessimism; since it now takes more pleasure or virtue to outweigh a given quantity of pain or vice, the result of the weighing, given a fixed set of facts, is more likely to be negative. The asymmetries also affect the theological problem of how an all-powerful, perfectly good God could create the evil we see around us. It again takes more pleasure to outweigh a given quantity of pain in God's creation, and if God gave humans the possibility of either benevolence or malice, then, assuming an initially equal probability of each, he did something whose expected moral value was negative. The obvious question, however, is whether this pattern of evil-favouring asymmetries is universal or whether in some pairs a good is greater than its corresponding evil. I will discuss this question in the last part of this paper, after first considering some different asymmetries between positive and negative forms of the same good or evil.

Return to Figs. 6 and 7. Though they differ in how they relate virtue and vice, they are both symmetrical around the vertical axis, and the views they express therefore treat positive and negative forms of virtue and vice as equal in value. Imagine that a given pleasure is exactly as good as a given pain is evil. (For Bentham and Sidgwick this will mean the two are equally intense, for Moore and Mayerfeld that the pleasure is to a specified degree more intense.) Fig. 7 makes a benevolent pleasure of intensity  $n$  in the pleasure exactly as good as a compassionate pain of intensity  $n$  at the pain: the  $n$ -point on the curve in the top right quadrant is exactly as high as the  $n$ -point on the comparable curve in the top left. It likewise makes a malicious pleasure of intensity  $n$  in the pain exactly as evil as an envious pain of intensity  $n$  at

the pleasure. Should we retain this feature of the graphs or supplement our up-down asymmetry with a further, left-right asymmetry?

The answer depends on whether positive and negative forms of virtue should be proportioned to each other, and in my view they should. If it is disproportionate, and more specifically selfish, to prefer a minor good for oneself to great goods for other people, surely it is equally disproportionate to prefer avoiding a minor evil for oneself to securing great goods for others, and likewise disproportionate to prefer a minor good for oneself to preventing great evils for others. The demands of proportionality apply not only within the categories of virtuous love and hatred but also across them, so an ideally virtuous person divides all his concerns, both positive and negative, in proportion to their objects' values. And this requires symmetry around the vertical axis: if virtuous love and hate are to be balanced proportionally, the functions determining their values must mirror each other, as in Fig. 7.

A similar issue arises about desert. While Figs. 8 and 9 make undesert a greater evil, they too are symmetrical around the vertical axis and so make positive and negative desert equally good. If one person's virtue is exactly as good as another's vice is evil, then the first's getting  $n$  units of happiness is exactly as good as the second's suffering  $n$  units of pain. Should we retain this symmetry, as we retained one for virtue?

The answer again depends on whether there are proportionality demands between positive and negative forms of goodness, and here I am less persuaded. Do we think society should balance its rewards and punishments, so that, say, a Nobel Peace Prize makes a person exactly as happy as life imprisonment makes one suffer? I do not think we have any such concern. We care that rewards be proportioned to rewards and punishments to punishments, but seem not to care about proportionality between the two categories. This leaves room for positive-negative asymmetries about desert, and here two seem attractive. I think many who value desert will say there is a stronger demand to punish the vicious than to reward the virtuous, so failing to do the former is a greater failing in justice, or involves a greater loss of value.<sup>25</sup> To reflect this view, we can make the peaks on the curves to the left of the vertical axis higher, so they represent greater positive value, and also make those curves cut the vertical axis further below the origin, so ignoring that value is worse. (This will be another case where a broadly negative value is more potent: returning evil for evil will be more important than matching good with good.) Second, compare the slopes of the curves on either side of a peak. In Figs. 8 and 9 the slopes are always steeper outside the peak, so getting more happiness or suffering than one deserves involves a greater loss of value than getting less. This is certainly attractive for negative desert. Just as the state should be more concerned not to punish the innocent than to punish all the guilty, so it should be more concerned not to punish the guilty too severely than not to punish them enough. But the parallel view is much less plausible for positive desert, where giving an excessive reward seems less

bad than giving an insufficient one.<sup>26</sup> If so, the shapes of the curves should differ on the two sides of the vertical axis, being steeper on the outside on the left, for negative desert, and on the inside on the right.<sup>27</sup>

Our main question, however, is whether our initial pattern of evil-favouring asymmetries is universal, or whether there are cases where a good is greater than its corresponding evil. This is not, it will turn out, an easy question.

To pursue it, consider Robert Nozick's fantasy of an experience machine that, by electrically stimulating the brain, can give one the experience and therefore the pleasure of any activity one likes. And assume with Nozick that life on this machine would not be intrinsically best and, in particular, would not be as good as the lives we currently lead. A question that has not to my knowledge been discussed is whether life on the experience machine merely lacks goods found in ordinary life or, while containing some goods such as pleasure, also contains evils that weigh against those goods. Does machine life merely lack positive value or does it also contain some negative value? Like Nozick, I take what is problematic about life on the machine to be its disconnect from reality. People who plug in have false beliefs about the world and their place in it; they think they are, say, climbing Mt. Everest when they are not. Nor do they actually accomplish any goals; while they intend to climb Everest, they do not do so. So there are two pairs of values highlighted by Nozick's example. One is knowledge and its contrary, false belief, surrounding the neutral state of not having any opinion about a subject. The other is achievement and its contrary, failure in the pursuit of a goal, with the neutral state of not pursuing the goal at all. Do both these pairs contain an intrinsic good and an opposing intrinsic evil? If so, do the good and evil relate symmetrically or is one greater than the other?

Let me begin with knowledge, and consider first the kind that involves knowing one's relation to the external world, and in particular to one's immediate environment. The absence of this knowledge seems a large part of what is troubling about the experience machine: that people on it believe falsely that they are doing things rather than being electrically stimulated plays a large role in making their condition less good. But I think it involves more the presence of an evil than merely the absence of a good. If having true beliefs about one's current environment is good – which some may dispute given the extreme particularity of its subject matter – it is surely not a great good. If someone with a painful terminal illness believes correctly that he is lying in a hospital bed, the goodness that knowledge involves does not weigh heavily against the evil of his pain. So if the delusions about one's place in the world generated by the machine do weigh heavily against its pleasures, they must be a positive evil, and that seems intuitively right. Being systematically mistaken about where one is is not just not good but evil, and more evil than its contrary is good. It may be objected that having false beliefs about one's place in the world is not much worse than having no beliefs about it, or not

being in a position to have beliefs about it. But that only shows that the asymmetry may have a downshift basis, where a neutral state has negative value. So we seem here to have an evil-favouring asymmetry like the one for virtue and vice: knowing one's relation to one's immediate environment has at best modest positive value, but being mistaken about it has greater negative value, in part because not having an opinion also has negative value.

But now consider a different kind of knowledge, of purely external facts about the world such as scientific laws. And think of some scientist of the past who had mostly false beliefs on this topic, as Aristotle did about the basic laws of physics and biology. Aristotle's errors about these laws do not seem significantly evil, especially in comparison with his non-scientific contemporaries who had no beliefs at all about them. Nor does those contemporaries' lack of beliefs seem evil – it is just the absence of a good – so there is here no downshift. A case like Aristotle's is complicated because there were other intrinsic goods associated with his scientific activities. His beliefs about physics and biology were arguably justified by his evidence, and there may be positive value in having justified beliefs even when they are false. Moreover, his active pursuit of true scientific beliefs showed a love of scientific knowledge that is a form of virtue. But I think we can abstract from these goods and consider his false beliefs on their own, and when we do, I do not think we find them significantly evil. But having true beliefs or knowledge about scientific laws does seem significantly good. Most philosophers who have discussed the intrinsic value of knowledge have held that the best knowledge is of the most general and explanatory principles, including pre-eminently the knowledge of scientific laws.<sup>28</sup> So here we seem to have an opposite, good-favouring asymmetry: knowing scientific laws is significantly good, but the contrary state of being mistaken about them is not significantly evil.

Though there are other kinds of knowledge, this value has already proved complex. For one subject-matter, that of one's relation to one's immediate environment, the evil of false belief seems greater than the good of knowledge, while for another, concerning scientific laws, the good of knowledge seems greater.<sup>29</sup> But there is here at least a partial break from our earlier pattern of evil-favouring asymmetries, in that sometimes knowledge is a greater good than false belief is evil. And at least one philosopher has affirmed this view. Making a contrast with hedonic values, though not distinguishing among subjects of knowledge, E.F. Carritt says, "Pain seems more obviously bad than pleasure is good, but knowledge more plausibly good than either ignorance or error is evil."<sup>30</sup>

When we turn to the second relational good, achievement, there is an even sharper break, since the contrary state of failure in pursuit of a goal seems not evil at all. Of romance Tennyson said, "Tis better to have loved and lost/Than never to have loved at all." Of practical endeavours we may say, similarly, "Tis better to have sought and failed than never to have sought at all," so failing in pursuit of a goal is if anything better than not pursuing it.

This conclusion is overdetermined when the goal in question is independently good, since then pursuing it, even without success, manifests the virtue of loving the good, while declining to pursue it may be viciously indifferent. But the conclusion seems also to hold for failures to achieve intrinsically neutral goals, as in games or business. While successfully achieving a neutral goal such as a low golf score or large profits for one's company can be significantly good, failing to achieve it does not seem significantly evil, and in particular does not seem worse than not pursuing it. Nor does this case involve a downshift, since not pursuing a neutral goal seems neutral in value rather than evil. So for this pair, and without the distinctions we found for knowledge, there seems to be a broad-based good-favouring asymmetry whereby achievement is a greater good than failure is an evil, if failure is evil at all. Together with the more restricted good-favouring asymmetry for knowledge, this provides a counterweight to the contrary asymmetries in discussions of optimism vs. pessimism and the problem of evil. Whereas the pleasure-pain and virtue-vice asymmetries make the world more likely to be on balance evil and God's creation of it harder to reconcile with his goodness, those for knowledge and achievement make it more likely to be good.<sup>31</sup> It is not always evil that is more potent; in some cases good has greater weight.

This paper does not have a grand conclusion. I have explored several possible asymmetries within pairs of intrinsic values, and while it would be exciting to report a single pattern in these asymmetries, I have not found one. There was an initial pattern of evil-favouring asymmetries in four pairs of values, but even there desert and undesert did, but virtue and vice did not, support a further left-right asymmetry between positive and negative forms of good or evil. And my later discussion, though more tentative, seemed to find some contrary good-favouring asymmetries in at least two other pairs of values. But even without a single pattern I hope the exploration has been illuminating. Our first impulse may be to relate the elements of a pair of values symmetrically, and this is certainly the simplest view. But there is no necessity for value-symmetry, and often an asymmetrical view is more attractive. Moreover, there are several different possible asymmetries and several different bases they can have, with different implications for claims about the right. As Moore said, ethics would be far simpler if values always related to each other symmetrically, but the truth seems rather more complicated.

### Notes

<sup>1</sup> Henry Sidgwick, *The Methods of Ethics*, 7<sup>th</sup> ed. (London: Macmillan, 1907), p. 413.

<sup>2</sup> G.E. Moore, *Principia Ethica* (Cambridge: Cambridge University Press, 1903), pp. 212, 222.

<sup>3</sup> Jamie Mayerfeld, *Suffering and Moral Responsibility* (New York: Oxford University Press, 1999), p. 136.

<sup>4</sup> The Moore-Mayerfeld view may seem more problematic when applied to choices within a single life. If a person prefers to trade off pain for himself for a slightly greater amount of pleasure for himself, are we to say he is wrong? And will we interfere to prevent him from making that choice? First, even if his preference is wrong, respect for his autonomy may make paternalistic interference with the preference wrong. Second, Mayerfeld suggests that the pleasure-pain asymmetry may be less strong within lives than across them, so a quantity of pleasure that would not outweigh a given pain for another can outweigh the same pain for oneself (*Suffering and Moral Responsibility*, pp. 150–51). Finally, and subject to this last qualification, some self-regarding trade-offs of pain for pleasure may indeed be wrong. (The symmetry view certainly thinks so.) In this paper I assume that value does not depend entirely on people's desires, so someone who prefers mindless pleasure to knowledge and achievement has a mistaken preference. If so, someone who trades off pain for only slightly more intense pleasure may likewise be mistaken.

<sup>5</sup> See e.g. James Griffin, "Is Unhappiness Morally More Important Than Happiness?" *Philosophical Quarterly* 29 (1979): 47–55.

<sup>6</sup> On the multiple role of evaluations in measures of inequality, see Larry S. Temkin, *Inequality* (New York: Oxford University Press, 1993).

<sup>7</sup> Mayerfeld defends the view in *Suffering and Moral Responsibility*, pp. 74–83.

<sup>8</sup> It is no part of this view that the psychological facts about pleasure and pain mirror each other exactly. In particular, it is possible and even plausible that the most intense pain humans can experience is more intense than their most intense pleasure. If so, there is, asymmetries aside, a stronger reason to prevent the intensest possible pain than to promote the intensest possible pleasure. What the pairwise asymmetry grounds is a stronger reason to relieve a pain of a given intensity than to promote a pleasure of the same intensity.

<sup>9</sup> Mayerfeld, *Suffering and Moral Responsibility*, pp. 134–36.

<sup>10</sup> G.E. Moore, "Mr. McTaggart's Ethics," *International Journal of Ethics* 13 (1903): 341–70, p. 358.

<sup>11</sup> The priority view is often formulated in whole-life terms, so it tells us to prefer benefits to those whose level of well-being has been lower through time. This can mean, contra the view in Fig. 4, that it is better to relieve a lesser present pain of someone who has suffered more in the past than a greater present pain of someone who has been happier. But the priority view can also be applied at single times. Then it says that in so far as our acts will have effects at time  $t$ , our priority should be improving the condition at  $t$  of those whose condition at  $t$  would otherwise be worst. And this single-time priority view comes close to a Fig. 4 asymmetry view that considers not discrete pleasures and pains, of which a person may have several simultaneously, but his overall hedonic condition at a time, which is in effect what Mayerfeld's conception of happiness and suffering does (*Suffering and Moral Responsibility*, Ch. 2). So while not close to every priority view, the Fig. 4 asymmetry view is close to one.

<sup>12</sup> Note that, unlike the pairwise asymmetry discussed above, this limit asymmetry does not require non-evaluative comparisons between pleasures and pains. Even if such comparisons are impossible, some pains can be more evil than any pleasure is good.

<sup>13</sup> This downshift too is possible without non-evaluative comparisons. So long as there is a non-arbitrary point of hedonic neutrality, as few who question such comparisons deny, both that point and pleasures only slightly above it can have negative value.

<sup>14</sup> Gregory S. Kavka, "The Paradox of Future Individuals," *Philosophy and Public Affairs* 11 (1982): 93–112. Critical-level utilitarianism has been defended more recently in a series of articles by Charles Blackorby, Walter Bossert, and David Donaldson; see, e.g., "Critical-Level Utilitarianism and the Population-Ethics Dilemma," *Economics and Philosophy* 13 (1997): 197–230.

<sup>15</sup> Derek Parfit, *Reasons and Persons* (Oxford: Clarendon Press, 1984), pp. 387–90.

<sup>16</sup> See Gustaf Arrhenius, "An Impossibility Theorem for Welfarist Axiologies," *Economics and Philosophy* 16 (2000): 247–66, pp. 255–56.

<sup>17</sup> I give a fuller account of these issues in *Virtue, Vice, and Value* (New York: Oxford University Press, 2001), Ch. 3.

<sup>18</sup> For a similar peak structure (though without proportionality) see Shelly Kagan, "Equality and Desert," in Louis P. Pojman and Owen McLeod, eds., *What Do We Deserve? A Reader on Justice and Desert* (New York: Oxford University Press, 1999), pp. 298–314. In Fig. 8 the "degree of virtue/vice = 0" line represents the desert-value of happiness and suffering for people who are neither virtuous nor vicious, or whose characters are at the neutral point between the two. This line has a peak at the origin, implying that what is best is these people's experiencing neither happiness nor suffering, though their doing so has neutral value, and that their experiencing either happiness or suffering has negative value. This may seem harsh; should their experiencing happiness or suffering not also have neutral value? But it is required by the peak structure plus proportionality.

<sup>19</sup> A.C. Ewing, *The Morality of Punishment* (London: Kegan Paul, Trench, Tubner & Co., 1929), p. 109.

<sup>20</sup> Sir William Blackstone, *Commentaries on the Laws of England* (Chicago: University of Chicago Press, 2002), Bk. 4, Ch. 27, p. 352; for a survey of variants on Blackstone's formula, see Alexander Volokh, "n Guilty Men," *University of Pennsylvania Law Review* 146 (1997): 173–216.

<sup>21</sup> The last claim is not true for excessive punishment, which on Ewing's view can be a great evil. It is when deterrence and reform favour milder punishments than retribution that they outweigh it; when they favour harsher ones, they do not.

<sup>22</sup> Note that, as before, a limit asymmetry does not require non-evaluative comparisons, which I have said are not possible for equality and inequality.

<sup>23</sup> Aristotle, *Nicomachean Ethics*, 1106b28–30. The remark comes when Aristotle is introducing his doctrine of the mean.

<sup>24</sup> That the intention-foresight distinction is engaged only by the production of evils and not by that of goods is noted in Thomas Nagel, *The View From Nowhere* (New York: Oxford University Press, 1986), p. 181.

<sup>25</sup> Ewing denies this, suggesting that rewarding the virtuous is just as good as punishing the vicious; see *The Morality of Punishment*, p. 147.

<sup>26</sup> This is certainly Kagan's view, and is captured in the "bell-motion" feature of his desert graphs; see "Equality and Desert," p. 301.

<sup>27</sup> This asymmetry could in principle be taken further. Some may say that, while a peak structure is necessary for negative desert, it is never less good in desert terms if a virtuous person is happier; on the contrary, it is always better. To capture this view, the desert-curves would have to have entirely different shapes on either side of the vertical axis, turning down from a peak to the left but continuing up toward a limit on the right. But this more radical asymmetry is problematic. It requires the right half of the "degree of virtue/vice = 0" line, which is below the horizontal axis in Figs. 8 and 9, to lie along that axis, implying that if a person is neither virtuous nor vicious his enjoying happiness has zero rather than negative desert-value. And while this may seem attractive in itself, it introduces a serious discontinuity into the account of desert-value. Imagine that a slightly vicious person enjoys great happiness. Given the more radically asymmetrical view, improving his virtue to the neutral point between virtue and vice makes for a very large increase in desert-value, from a possibly large negative value to zero. More specifically, it can make it better in desert terms, and perhaps even on balance, to make a small improvement in virtue if that takes one to or across the neutral point than to make larger increases either below or above that point, for example, to improve from 1 unit of vice to neutrality than to improve from 5 to 1 units of vice or from 0 to 4 units of virtue. If this discontinuity is counterintuitive, that is a reason to resist the more radical left-right asymmetry and adopt at best the milder one that retains peaks on both sides of the vertical axis.

<sup>28</sup> See e.g. F.H. Bradley, *The Principles of Logic*, 2<sup>nd</sup> ed. (London: Oxford University Press, 1922), pp. 684–88; W.D. Ross, *The Right and the Good* (Oxford: Clarendon Press, 1930),



pp. 139, 147–49; and Thomas Hurka, *Perfectionism* (New York: Oxford University Press, 1993), pp. 115–20.

<sup>29</sup> A third kind of knowledge is of one's inner states, not only one's experiences but also one's character traits and dispositions. I have no clear intuitions about whether either knowledge of or delusion about these states is a greater intrinsic value.

<sup>30</sup> E.F. Carritt, *Ethical and Political Thinking* (Oxford: Clarendon Press, 1947), p. 94.

<sup>31</sup> Perhaps there are other values with the same effect, such as beauty and creativity. They too may be more good than their contraries are evil, if those contraries are evil at all.