

VALUE THEORY

Thomas Hurka

University of Toronto

The theory of value or of the good is one of the two main branches of ethical theory, alongside the theory of the right. Whereas the theory of the right specifies which actions are right and which are wrong, the theory of value says which states of affairs are intrinsically good and which intrinsically evil. The theory of the right may say that keeping promises is right and lying wrong; the theory of value can say that pleasure is good and pain evil, or that knowledge and virtue are good and vice evil. Since these states are not actions they cannot be right or wrong, but they can have positive or negative value.

The theory of value is important, first, because it gives content to some important claims about the right. Consequentialists about the right hold that one ought always to do what will result in the best outcome; to know what this implies we must know in particular what makes outcomes good. Even non-consequentialists usually recognize some moral duty to produce good outcomes, and that duty, too, needs content.¹ There is no point telling people to promote the good without telling them what the good is. Second, on some non-consequentialist views the duties that compete with promoting the good likewise presuppose claims about the good. These duties can make it wrong to do what will have the best overall outcome, for example, wrong intentionally to kill one innocent person even if this will save five innocent people's lives. But some say this is because, given an initial intrinsic value of life, there is not only a duty to promote and preserve it but also a separate and stronger duty not to destroy it; there can also be

separate and stronger duties not to destroy other goods such as knowledge and virtue.² On this view even the duties that constrain pursuit of the good concern the good, though as something to be respected rather than simply brought about. Finally, the theory of value is important in itself. Often things happen that do not result from anyone's choice and could not have been prevented by choice. They therefore cannot be right or wrong, but they can be intrinsically good or evil. Thus, it can be evil if someone suffers pain as a result of an entirely unforeseeable accident or good if she enjoys serendipitous pleasure; it can likewise be good if she stumbles onto valuable knowledge or is born with a virtuous character. Whereas the theory of the right judges only actions people voluntarily control, the theory of value can range over all the states of affairs the world contains.

Consistent with this point, there are several competing views about what value is. One holds that goodness is an unanalyzable property that can be had by states of affairs regardless of their connection to choice;³ others analyze the good as that the love of which is correct or as that which people have moral reason to desire and if possible pursue.⁴ But these views are less different than they seem. Those who treat goodness as unanalyzable usually agree that the good is what it is correct to love and what people have reason to desire; their only dispute with the other views concerns whether these latter claims are self-standing or derive from one that is more fundamental. There are also competing accounts of what it is for goodness to be intrinsic. A strict view says a state's intrinsic goodness can depend only on its intrinsic properties, those that do not involve relations to other states; it therefore tests for intrinsic value by asking whether a universe containing only a given state and no other would be good.⁵ A less strict view equates a state's intrinsic goodness with that portion of the overall goodness of the world that is located in or

attributable to it, whatever properties that goodness depends on. Both these views distinguish intrinsic from instrumental goodness, or goodness as a means to something else that is good. But they differ about what can be called conditional goodness. Consider the claim that pleasure is good only when it is the pleasure of a morally virtuous person.⁶ On the strict view, the goodness this claim ascribes is not intrinsic, since it depends on a relation between the pleasure that has it and virtue; on the looser view, it is intrinsic. As we will see, however, this difference has no substantive implications, since any claim that can be made using the one definition of “intrinsic” can also be made using the other. This entry will therefore adopt the looser view and allow that a state’s intrinsic goodness can in principle depend on its relations.

Assuming these conceptual issues settled, philosophers have defended very different views about which states are intrinsically good and evil. In the last part of the 20th century there was a tendency to prefer theories of value that are simple and austere, with only a few goods and only ones seen as making modest claims. But there is no persuasive rationale for these preferences. It is true that a theory should try other things equal to unify its values, and the more it can do so the greater its appeal. But the unification cannot be at the expense of intuitive credibility, and in particular cannot justify ignoring values that seem intuitively compelling. Nor is there any reason why the facts about value must fit some pre-conceived ideal of austerity. The more credible view is that there is an immense variety of at least initially plausible intrinsic values and of ways of combining them. Some of these values can be unified to some degree, and showing how is one task of theory. But it is hard to see them all being reduced to a single fundamental value; in addition, while some make relatively modest claims, others are more extravagant. The realm of value, in other words, is rich in possibilities and in subjects for debate.

This entry will survey a series of candidate intrinsic values, in rough sequence from the less to the more controversial.

1. Hedonism

The simplest theory of value is hedonism, which holds that only pleasure is intrinsically good and only pain intrinsically evil. Hedonism was defended in the ancient world by Epicurus and criticized by Plato and Aristotle; it was also defended by the classical utilitarians, notably Jeremy Bentham and Henry Sidgwick, and retains adherents today.⁷ It is a simple theory because it restricts good and evil to the one dimension of felt pleasure and pain, so there is only the one intrinsic good and one intrinsic evil.

Despite its simplicity, hedonism can be formulated in different ways, depending, first, on how the concept of pleasure is understood. One view identifies pleasures as sensations with an introspectible quality of pleasantness and pains as ones with the contrary quality of painfulness; this leads to a version of hedonism in which the only values are feelings with these introspectible qualities. Against this view it is sometimes objected that there are no such qualities; there is no feeling in common between, say, the pleasure of drinking beer and that of solving a crossword puzzle. But the view's defenders can reply that the quality of pleasantness is never experienced alone. Pleasurable sensations always have other introspectible qualities that make them as wholes very different, but they share the quality of pleasantness and can be ranked in pleasantness, just as we can rank the loudness of sounds that differ radically in pitch and timbre.⁸ A rival view identifies pleasures as those sensations people want to have and to continue having just for their qualities as sensations. It is not clear, however, that this view successfully picks out only

pleasures; can someone not want the sensation of redness just as that sensation? In addition, the view seems to point beyond hedonism to the more general theory that the good is whatever people desire, regardless of whether it is a sensation.⁹ Nonetheless, a second version of hedonism identifies its good as a sensation people want just for its qualities as a sensation.

However it understands pleasure, hedonism normally values both of what can be called simple and intentional pleasures. Simple pleasures are unstructured sensations with whatever feature makes them pleasures; they include, most notably, bodily pleasures such as those of taste and touch. Intentional pleasures, by contrast, are directed at an intentional object; one is pleased by something or that something is the case, for example, that one's friend got a promotion. Intentional pleasures are more complex than simple ones and raise more complex moral issues; we will discuss some of these below. But both types are pleasures and can be compared for their degrees of pleasantness.

To yield determinate value-judgements, hedonism must be able to measure quantities of pleasure and pain. There are several dimensions to this measurement. If pleasures are discrete sensations, it is better to have more than fewer of them and also better to have ones that last for a longer time. In addition, it is better to have pleasures that are more intense, just as it is worse to have more intense pains.¹⁰ But there are different views about how the intensities of these two states compare. The most common view, held for example by Bentham and Sidgwick, treats pleasure and pain symmetrically, so a pain of a given intensity is always exactly as evil as a pleasure of the same intensity is good. But a different view holds that pain is a greater evil than pleasure is a good. Its most extreme version holds that pleasure is not good at all, but this implies that a life with many intense pleasures and only a few mild pains is on balance not worth living.

A more moderate version holds, more plausibly, only that pain of a given intensity is worse than pleasure of the same intensity is good, so it is more important to prevent the pain than to provide the pleasure. (This gives pain some priority over pleasure, but not infinite priority.) And this view can be extended to give disproportionate weight to more intense pains, so that given an intense pain for one person and two pains of just over half the intensity for two other people, it is more important to relieve the one intense pain. Within the general framework of hedonism, this view attaches the greatest ethical significance to very intense pains.¹¹

A final issue concerns the related concepts of happiness and suffering. Though happiness is a more inclusive concept than pleasure -- to call someone happy is to say more than that he is experiencing some pleasurable sensation now -- some philosophers define it in terms of pleasure, so a happy life is one with a clear preponderance of pleasures over pains. But others treat happiness as a distinct state, one involving a feeling of satisfaction with one's life as a whole, in at least most aspects and including the past and future as well as the present; an analogous view equates suffering or despair with dissatisfaction with one's life as a whole.¹² Some who take this view treat happiness as the central hedonic value, so what is to be promoted is not individual pleasurable feelings but this more general state of life-satisfaction. But within a framework that values sensations it is hard to see the rationale for this view. If happiness is good feeling about one's life as a whole, why should it count more than similar feelings with other intentional objects or with no objects at all? Happiness may be more stable than other good feelings, but that does not make it intrinsically more important. And the same is certainly true of bad feelings. Though despair about one's life as a whole is certainly an evil, no one would on that basis deny that intense bodily pain is comparably evil.

Hedonism is on strong ground when it says that pleasure is a good and pain an evil. But its stronger claim that these are the only intrinsic values has met with many objections. One is that hedonism can count as morally ideal a life containing only mindless pleasures and none of the higher achievements in art, science, and personal relations that are the distinctive prerogative of human beings. This objection has been raised in fiction, from the lotos-eaters of Homer's *Odyssey* to Aldous Huxley's *Brave New World*; it is also expressed in Robert Nozick's fantasy of an "experience machine" that, by electrically stimulating the brain, can give one the illusion and therefore the pleasure of any activity even though one is not actually engaged in it.¹³ While hedonism implies that a life spent entirely on this machine would be ideal, Nozick and others find it repellent. A second objection is that hedonism gives positive value to pleasures that are morally vicious. If a torturer takes sadistic pleasure in his victim's pain, hedonism says this makes the overall situation better than if the torturer were indifferent to the pain or, worse, pained by it. But surely it is compassion that is good and sadism that is bad.¹⁴ Those who are persuaded by these objections may adopt a rival "perfectionist" theory that values human excellences or perfections such as knowledge, difficult achievements, and moral virtue instead of or as well as pleasure. But before we examine perfectionism we should consider a second theory that shares important features with hedonism.

2. Desire Theories

Hedonism can be called democratic about the value of activities, since it holds that how good they are for a given person depends on how much he in particular enjoys them. As Bentham put it, if the pleasure they give is the same, pushpin (a game similar to tiddly-winks) is as good as

poetry. But hedonism is not completely democratic, since it requires people to prefer pleasure and the avoidance of pain to everything else. Not everyone does this. At the end of his life Freud refused painkillers in order to keep his mind clear for thinking, and this type of preference for a perfectionist good over pleasure is in fact quite common. Many people at least sometimes prefer knowing the truth or pursuing a difficult goal to an alternative that would be more pleasant.

According to hedonism, when they do this they are wrong.

Some who cannot accept this claim may move to a more fully democratic theory, one that equates the good in a person's life with his getting whatever he desires.¹⁵ If he wants pleasure, then pleasure is good for him, but if he prefers clear thinking, then that is better. Because people often do want pleasure, this theory's implications often coincide with those of hedonism; in addition, the known satisfaction of a desire, even for something other than pleasure, usually brings some pleasure. But the theories diverge whenever the pleasure a person will get from satisfying a desire for something other than pleasure is less than the alternative pleasure he forgoes, and they can also diverge in another type of case. Imagine that a person wants to be respected by his co-workers, believes he is respected, and derives pleasure from his belief, whereas in fact his co-workers ridicule him behind his back. Hedonism says their ridicule is not bad for him, since it does not affect his feelings, but a desire theory can say it makes his life worse, by frustrating an important desire.

Like hedonism, a desire theory can be formulated in different ways. One version says that what is good is the state of affairs in which a desire exists and is satisfied. But assuming that the satisfaction of more intense desires is better, this implies that people should form intense desires that are guaranteed to be satisfied, such as that grass be green and that two plus two equal four.

Apart from its counterintuitiveness, this is not what desire theorists typically say. Most of them hold that when a person desires some state of affairs, that state is what is good; the desire is a condition of value in something else rather than a part of what has value. This second version of the theory can tell people to form some desires if that will help them satisfy other, more fundamental desires, but it says nothing about what their fundamental desires should be. Rather than using values to guide desire, this theory waits for desires to create value, treating them as items outside the realm of good and evil that give value to items within it.

There is another point where versions of the desire theory can differ. Many philosophers present the theory as defining what is good *for* a person or constitutes her well-being, and some desires seem irrelevant to this issue. If I want some person I once met but have not heard of since to flourish and she does so, how does this make *my* life better? Or how does it increase my well-being if, unbeknownst to me, my desire that there be life on Mars is satisfied? One response to this difficulty is to count as good only the desire-satisfaction a person knows about and feels satisfaction in, but this brings the theory close to hedonism and prevents it from saying that being ridiculed behind one's back can make one's life worse. The more common response, therefore, is to place a content restriction on the desires relevant to a person's good: only desires about *his* life contribute to his good, whereas desires about distant people or planets do not.¹⁶ Deciding exactly which desires concern his life is difficult, but desires for internal states of himself obviously count, as do some desires for relations between himself and others, such as his desire that he not be ridiculed. But however the relevant boundary is drawn, only the satisfaction of desires about his life counts toward his good.

A further possible modification responds to the fact that people sometimes desire what is

not good for them, so getting it does not benefit them. In many cases they desire one thing only as a means to a second, secure the first, and find that it does not produce the second; for example, they desire money only because they think it will make them happy, get the money, and find they are not happy. These cases pose no special problem if the desire theory values the satisfaction only of desires for things as ends rather than just as means, which it is independently plausible for it to do. But there may be cases where a person's desire rests on a false belief that is not about means, and some theorists respond to them by equating a person's good with the satisfaction not of her actual desires but of those desires she would have if she were fully informed and rational.¹⁷ They identify her good not by looking at what she actually wants but at what she would want in some idealized circumstances. A sophisticated version of this theory equates a person's good with what her fully informed self would want for her uninformed self, taking the latter's uninformedness into account.¹⁸ These informed-desire theories are still democratic, since what different people would want if informed may be different. But they define the good in terms of hypothetical, idealized desires rather than actual ones.

Desire theories are popular, especially among economists but also among philosophers, and for several reasons. They seem to simplify the metaphysics of value, making it not a mysterious addition to the universe but the product of human desires. They are democratic and also make value comparatively easy to identify and measure. If we want to know what is good, we find out what people desire; if we want to know its degree of goodness, we find out how intensely they desire it, or how many other things they would risk to get it. Desire theories also partially answer objections to hedonism like those about the experience machine, since they say that if people do not want to plug in, as most of us do not, plugging in is not best for us. But they

do not completely answer the objection, since they also say that for people who do want to plug in doing so *is* best. Many philosophers reject this claim, holding that some features of life that are excluded by the machine, such as real knowledge, achievement and attachments to other people, are good in themselves and regardless of whether they are desired. Like hedonists, these philosophers deny that value is created by desire; on the contrary, they think value should guide desire, so good desires are those directed at what is independently good. But now the independent good is not just pleasure but, on what I have called a perfectionist view, includes other intrinsic excellences.

3. Perfectionism

Perfectionism has been prominent in the history of philosophy, defended in the ancient world by Plato and Aristotle and later by St. Thomas Aquinas, Georg Wilhelm Friedrich Hegel, Karl Marx, Friedrich Nietzsche, and others. Like hedonism and the desire theory, it comes in different versions, depending on which particular states of humans are deemed good. Some perfectionists give the highest value to knowledge, especially philosophical or religious knowledge; others prefer active goods such as political achievement and the creation of art; yet others emphasize moral virtue. While agreeing that the good is not just pleasure or desire-satisfaction, they place it in different human excellences. Some perfectionist theories are ineliminably pluralistic, recognizing a number of different goods with no unifying explanation of why they are goods. But the perfectionist tradition has also tried to unify its various values under some more abstract heads.

This unification is possible to some extent between the theoretical good of knowledge

and the practical good of achievement. Both these goods involve, first, a matching relation between a mind and the outside world. To count as knowledge a belief must be true, whereas achievement involves making the world match a goal one has formed. The direction of match in the two cases is different, with one's mind having to match the world in knowledge and the world coming to match one's mind in achievement. But both goods involve some matching and therefore capture a central part of what is missing on the experience machine: people who plug into the machine are disconnected from reality, having mostly false beliefs about their situation and never actually achieving any goals. Second, there are further similarities in the factors determining degrees of these values. Many writers on knowledge say its most valuable instances have contents that both stretch across space, times, and objects and also explain many other items of knowledge, which are therefore subordinate to it in an explanatory hierarchy of knowledge. This happens, for example, when a scientist knows an abstract physical law and uses it to explain many particular physical phenomena. Similarly, writers on achievement often value goals that stretch across times and persons and require many other subordinate goals to be achieved as means to them. This gives special value to political achievements and the carrying out of a plan for one's life as a whole, as well as to particular activities that are complex, intricate, and challenging. In their different domains, then, knowledge and achievement again instantiate a common formal structure.¹⁹

Perfectionism can also unify the moral virtues. Instead of just adding benevolence, courage, and the like to its list of goods, it can say they are all higher-level intrinsic goods that involve morally appropriate attitudes to other goods and evils. If knowledge and the pleasure of others are good, then caring positively about them, or desiring, pursuing, and taking pleasure in

them for their own sakes, is also good and constitutes virtue; by contrast, caring negatively about them, for example, by trying to destroy them, is evil and vicious. The first of these claims makes benevolence a virtue; the second makes malicious envy a vice. This account of the virtues, which was popular with early 20th century perfectionists such as Hastings Rashdall, G. E. Moore, and W. D. Ross, does not unify the virtues with other goods; on the contrary, it emphasizes their distinctness from goods that are not in the same way higher-level. But it does give the virtues and vices a common explanatory ground.²⁰

A theory can also see some intrinsic values as instantiating other, more fundamental values to a high degree. For example, it is plausible that an important element in a good human life is deep and loving relationships with a spouse, children, or friends. And though some have treated these as constituting a distinct good, they can also be seen as involving, alongside intense enjoyment, thorough knowledge of another's character, the pursuit of goals that extend beyond the self and over time form a complex hierarchy, and moral virtue, as one cares about the other's happiness, success, and good character for their own sakes. Loving relationships may be a very great good, but they may just instantiate other goods to a very high degree.²¹

Some more ambitious approaches try to unify all the perfectionist goods. One appeals to the concept of human nature, which in different formulations it takes to consist in those properties essential to humans, distinctive of them, or essential and distinctive.²² Its central idea is that the good in a human's life consists in the full development of whatever is fundamental to human nature; it is often generalized to hold that the good of any natural thing consists in developing its nature. This view can generate different particular values, depending on which properties it takes to constitute human nature. Aristotle, the first and best-known proponent of

this approach, believed that it is fundamental to humans to be rational and that the best life therefore is most rational.²³ Marx held, rather differently, that it is essential to humans to engage in productive, cooperative labour,²⁴ while Nietzsche held that humans fundamentally exercise a will to power.²⁵ But all these philosophers grounded their specific values in the same abstract ideal of developing human nature. The second approach, adopted by Idealists such as Hegel and F. H. Bradley, tries to unify the different perfections under the heading of “organic unity,” or unity-in-difference. On this view, intrinsic value is created whenever initially diverse elements are brought into an organized unity, or when a whole contains tightly related but also strongly differentiated parts. The degree of this whole’s value depends both on the degree of its parts’ differentiation and on the degree of their final unity, and all other values somehow instantiate this basic value of organic unity.²⁶

To be successful, these unifying approaches must have an intuitively plausible abstract principle, which each to some extent has. They must also yield an attractive list of concrete goods, and here, too, each has some success. It is plausible that humans are essentially rational, given the central role rationality plays in explaining human behaviour, and rationality is also realized to a high degree in the forms of knowledge and achievement described above. These forms likewise involve organic unity, both in the matching they require between a mind and the world and in the hierarchical structuring of beliefs and goals that gives them their greatest worth. But the approaches have a harder time accommodating moral virtue. They have certainly tried to: Aristotle held that developing the rationality that constitutes human nature entails being virtuous, while Bradley argued that a unified self is necessarily a virtuous self. But each of these claims is disputable. If we include rationality in human nature because it explains human behaviour, can it

not be as present in vice as in virtue? Is a person who carries out a complex plan aimed at hurting many people not exercising rationality as much as if he benefited them?²⁷ And can a self not be as unified around evil ends as around good ones? Though a perfectionist theory can unite some of its goods under more abstract heads, it is not clear that it can successfully unite them all.

It may be, then, that the most plausible perfectionist theories are pluralistic, and there is another factor that pushes in the same direction. Most value-theorists grant that pleasure is to some degree good and, even more commonly, that pain is evil, and it is hard to see how these values can be unified with perfectionist ones. Can a person who suffers pain not still develop his human nature and still instantiate organic unity? (Pain often signals disruptions in unified human functioning and can cause further disruptions, but the question is whether it is contrary to such functioning in itself.) If he can, the resulting theory will be pluralistic in another respect, combining perfectionist values with hedonic or desire-based ones that are independent of them.

The simplest such theory treats the two types of good as independent, simply adding them to determine the overall value in a person's life at a time. But a more complex view says they only have significant value in combination: pleasure on its own has little worth, as do perfectionist activities on their own, and significant goodness results only when a person engages in a perfectionist activity that he also wants and takes pleasure in for its own sake. John Stuart Mill's theory of "higher pleasures" has this form. By holding that all goods are pleasures, Mill agreed with Bentham that a perfectionist activity such as reading poetry has no value if it is not accompanied by pleasure. But he denied that when it is accompanied by pleasure, its value is determined entirely by the intensity or quantity of that pleasure. There are considerations of quality that make a pleasurable reading of poetry better than an equally pleasurable playing of

pushpin, and in fact Mill held that the resulting higher pleasure is much better than lower pleasures such as those of the body. For him, both pleasure and perfectionist activities have little value on their own; what is significantly good is only the combination of the two.²⁸

4. Comparison and Aggregation

The values we have identified so far are instantiated in states of individual persons at individual times. A person can feel a pleasure at a given time, or know one or more truths, or be pursuing a goal that he will eventually achieve. But a complete theory of value must be able to combine these values into measures of the goodness or evil of larger states of affairs and, ultimately, of the whole universe. This process has two parts. First, the theory must be able to compare different values to determine the overall value in each person's life at each time. Then it must be able to aggregate these measures across times and persons to arrive at measures of the total value in each person's life through time and then in a whole population. Let us consider these tasks in turn.

The topic of comparison raises the question of whether some values are higher or greater than others. The strongest claim here is that one good is infinitely or lexically greater than another, so even the smallest quantity of the former outweighs any amount of the latter. This type of claim is sometimes made about virtue, which is said to have infinite value compared to pleasure and even to all non-moral goods.²⁹ But high-minded though it sounds, this claim is hard to accept intuitively. It implies that if a virtuous person suffers unremitting and agonizing pain, then despite that pain his life is overwhelmingly good with only a tiny admixture of evil; it also implies that if this person were slightly less virtuous but enjoyed ecstasy, his life would be worse.

And lexical claims are in general dubious. Whatever considerations make a state such as pleasure good to some degree seem also to suggest that it is good enough to at least sometimes outweigh instances of other goods and evils.

This leaves non-lexical comparative claims, which say that one good is only finitely greater than one or more others. This type of claim has again been made about virtue in comparison with non-moral goods, about knowledge in comparison with practical achievement, and about numerous other values. But it cannot always be formulated strictly. Imagine that we have measured each of two goods on its own cardinal scale and, given those scales, have decided that one unit of the first good equals two units of the second. This seems to make the first good finitely greater than the second, but its doing so depends on our choice of units on the scales, which is arbitrary. If we had measured the first good using units that are four times larger, the 2:1 ratio would have run in the opposite direction. So we cannot say strictly that one of these goods is greater than the other, but we can say this in a less formal way. If the result of our comparative judgements is that most people should spend most of their time pursuing the first good rather than the second, or should usually prefer its instances in conflicts with the second's, then for practical purposes the first good is finitely greater.

Some claims of this sort are intuitively attractive, for example, the claim that pain is a greater evil than pleasure is a good. But others are harder to accept. Thus, the claim of Aristotle, Aquinas, and others that knowledge is a greater good than practical achievement is hard to reconcile with the many parallels between the two. If both these goods involve a matching relation between a mind and the world as well as hierarchically ordered mental states, should their values not also be roughly similar? As for virtue, the most plausible view is that it is a lesser

good in the following sense: that its instances always have less value than their particular intentional objects. Although compassion for another's pain is good, for example, it cannot be as good as the pain is evil; it cannot be better for there to be pain and a feeling of compassion for it than no pain and no compassion. Similarly, it cannot be better for a person to have a vicious impulse and feel shame about it than to have no such impulse and no shame. This is not to say that virtue can never outweigh non-moral goods. If an attitude can have as much as one half of its object's value, then virtuously pursuing another person's ten units of pleasure can have more value than one unit of pleasure for oneself, so preferring the other person's ten units makes one's own life better. But virtue remains a lesser good in the sense that it always has less value than its specific intentional object.³⁰

There is another, different issue about comparison. It concerns not how the individual values weigh against each other, but what pattern an ideal combination of them has. Imagine that a number of goods are of roughly equal value, so it is in that respect a matter of indifference how a person chooses among them. We may nonetheless have views about what mixture of them he should seek in his life. One view says he should aim at a well-rounded achievement of all the goods, so his life embodies a variety of values rather than specializing narrowly on any one.³¹ A contrary view says he should choose one good -- it does not matter which -- and concentrate on it, so his greatest achievement of a value is as great as possible. The practical implications of these views depend on how pursuing a variety of goods affects one's achievement of them individually. If seeking variety hampers this achievement, resulting in a mediocre diletantism, even the first view favours some specialization; if it enhances one's achievement of the individual goods, say by cross-fertilization, even the second view favours some well-

roundedness. But any position on this issue presupposes an abstract view about which pattern of goods in itself makes a life best, and here the two contrary views just described are at least initially plausible. Parallel views are possible about other topics. Thus, one can hold that a group of people such as a society do best when they achieve excellence in a wide variety of domains; perhaps this is part of what is admirable about ancient Greece. Or one can say that, although certain animal species are intrinsically of equal value, the world is better when it contains a broad diversity of species rather than only a few.

Having arrived at a measure of each person's overall achievement of all values at each time, a theory must aggregate these measures across times and persons. It must decide how a person's values at different times combine to determine the value in his life as a whole, and how the values in different people's lives combine to make up the value in a population or in all humanity. These are parallel issues, and they are also difficult. About each several different views are possible, but none is entirely free of difficulties.

The simplest aggregative view involves addition, so the value in a life equals the sum of the values at its individual moments and the value in a population equals the sum of the values in its members. This view was taken by the classical utilitarians, especially Bentham and Mill, and also by some perfectionists. But it has several troubling implications. Across persons, it implies that we would make the world much better if we created a large number of new human beings, even if their lives had much less of what gives life value than those of existing people. Against this, some philosophers deny that adding new people does anything to make the world better;³² others hold that population additions can make the world only a little better. But both groups deny the strong duty to procreate that follows from an additive view. Relatedly, this view implies

what have been called certain “repugnant” conclusions. Try to imagine an ideal life, one that lasts a long time and scores very highly on whatever dimensions make life good -- it is intensely happy, involves deep understanding and important achievements, and is thoroughly virtuous. On an additive view there is another, longer life that would be better even though its value at any particular time is negligible. If this second life is long enough, the sum of its values across times will be greater than in the supposedly ideal life. Or try to imagine an ideal population, one in which a large number of people all enjoy wonderfully valuable lives. On an additive view there is another, larger population that is better even though its members all lead lives that are barely worth living.³³

These implications are avoided by a different approach, which equates the aggregate value in a collection of states with the average value per member. It denies that a very large population of people with lives barely worth living has great value; on the contrary, it holds that because of its low average this population has minimal value, and it makes a similar claim about a long life with a low average value per time. But this averaging view makes other problematic claims. It holds that, no matter how high the current average well-being per person, adding extra people even slightly below that average makes the population worse, so there is a duty not to procreate, and it holds this even if the additions will have no effects on existing people. It likewise holds that if additional years in a life will be below the average for a person’s life to now, those years should not be lived -- it would be better if the person died.³⁴ Claims like these are at least intelligible given perfectionist values; many believe that a career in, say, sports or art is better if it ends near its peak rather than continuing through a long period of decline. But the claims are hard to accept when applied to other values such as pain -- we would not say that

adding pain of less than the current average intensity makes things better -- or to all values taken together. And in some versions the averaging view has even more horrific implications, for example, that it is best to kill off any below-average members of a population.

Some other aggregative views echo the two pattern-views about comparison, those valuing, respectively, well-roundedness and specialization. One such view values equality in the distribution of other goods, holding that a society in which, say, happiness is equally distributed can be better than an unequal society even though the sum of happiness in the second society is greater.³⁵ A contrary view, defended by Nietzsche, holds that the value of a society depends primarily or solely on the achievements of its few most excellent members, so everyone else should dedicate themselves to improving the lives of those few;³⁶ a parallel view about times says the value of a life depends primarily on its achievements at its few best moments. These aggregative views do not address the issues about collection size that distinguish the additive and averaging views; those must be settled some other way. But taking a collection's size as given, they are like the parallel views about comparison in preferring certain patterns of distribution within it.

There can also be issues about where a view aggregates first: across times or across persons. Consider a view that values equality in people's happiness but aggregates across times by adding. If it first adds across times and then equalizes across persons, it will end up caring about equality in the total happiness in people's lives as wholes, or in their total life-happiness. This means that if a person who is a little better off than another person now was much worse off in the past, the view will not mind the present inequality and may even want it increased, to further reduce the inequality between their lives as wholes. It will also not object to practices that

treat people of different ages differently, such as mandatory retirement, so long as these practices treat them all the same way through time. But the view has very different implications if it first equalizes across the states of different persons at each time and then adds across times. In this formulation it does object to mandatory retirement, for treating the old at a time worse than the young at that time, and cares about equalizing only in the present and the future. If someone who is a little better off now was much worse off in the past, the view still says the present inequality is evil and should be removed.³⁷

Claims about aggregation, especially when they concern whole populations or the entire universe, can seem remote from everyday moral thinking. How are such grand issues relevant to the ordinary moral agent? But when interacting with another person, we do not want to give him a short-term benefit that will make his life as a whole worse, and whether that happens depends crucially on how his life's overall value is determined. Similarly, we do not want to benefit some people in a way that makes our society as a whole worse, which presupposes a view about social aggregation. Unfortunately, we often do not know the other facts needed to determine these aggregate values and therefore cannot judge our actions decisively. But for two reasons this need not be a debilitating disadvantage. First, sometimes our current action will have no effect on those other facts. If it will increase the aggregate in a large collection it will do so only by changing some local facts, so we need attend only to those facts. Second, even when it may affect remote facts, it can often do so in either of two contrary ways that are equally probable; here we can let these probabilities cancel each other out and again attend only to local facts. Even in these cases, however, if an action of ours is right because of its effects, it is so ultimately because of its effects on value as completely aggregated across times and persons.

5. The Principle of Organic Unities

Though most of the values we have examined so far have involved discrete states of individual persons at individual times, some views make the value these states contribute to the world depend on their relations to other states. These views presuppose an important principle that Moore called the principle of organic unities.

Stated generically, this principle says that the value of a whole need not equal the sum of the values its parts would have on their own; if two or more states enter into the relations that constitute a given whole, the resulting value may be either more or less than those states had when apart.³⁸ This principle is illustrated, first, by views that value organic unity itself. If two organic unities are combined into a single, larger unity, the result is more unity and therefore more value than existed before. But the principle is also implicit in Mill's theory of higher pleasures, which holds that a perfectionist activity such as reading poetry adds value to a life only when it is accompanied by pleasure in it for itself, as well in views that value knowledge and achievement, well-roundedness, and distributive equality.

There are, however, two different ways the principle can be interpreted, with the choice between them depending in part on what we think it is for value to be intrinsic.³⁹ Assume first the strict view on which a state's intrinsic goodness can depend only on its intrinsic properties. Given this view, it cannot be that when parts enter a whole the parts' values themselves change; the parts must be exactly as good or evil inside the whole as outside it. Instead, any additional value that results from their combination must be located in the whole as a whole, or in the whole considered as a separate entity comprising those parts in those relations, with this value of the whole as a whole to be added to the values in its parts to determine its value overall. But the

looser view that allows intrinsic value to depend on relations involves no such condition. When parts enter a whole their own values can change, so if the whole is better or worse it can be because those parts are better or worse.

The difference between these interpretations does not affect the overall values of states of affairs, which can always come out exactly the same. Whatever final value the first or holistic interpretation arrives at by finding additional value in a whole as a whole, the second or variability interpretation can arrive at by changing the values in its parts. And sometimes the difference seems not to matter at all. This is true, for example, of Mill's theory of higher pleasures. The holistic interpretation says that when a person who reads poetry takes pleasure in doing so, his reading still has zero value but there is additional value in the reading-plus-pleasure-in-reading, while the variability interpretation says the reading itself now has value. There seems nothing significant to choose between these claims. The same is true of claims about well-roundedness, which can either value variety as a holistic property of collections or say the value of increases in a given good becomes greater the less that good has been achieved compared to other goods.

In other cases, however, one interpretation seems to fit a given view better than the other. Thus, the variability interpretation seems best for claims about the value of knowledge and achievement, which are usually seen as states of an individual. If someone's long-time pursuit of an important goal turns out successful, that is normally seen as making her life and activities better rather than creating new value in a whole combining those activities, their external result, and the relation between them. But the rival holistic view seems best for claims about equality. If two people enjoy equal happiness, that does not make either person's happiness itself better;

instead, the additional egalitarian value is located in the relation between their levels of happiness as a relation.

Other values raise more complex issues. One we have not yet discussed is desert, which makes it good if people get what they deserve and evil if they get the opposite.⁴⁰ People can be said to deserve different things on different bases: medals for acting bravely, punishments for committing a crime, or income for making economic contributions. But a prominent view holds that people deserve happiness if they are virtuous and suffering if they are vicious. In terms of organic unities, this means the combination of virtue and happiness in the same life has positive value over and above the values the virtue and happiness would have alone, whereas the combination of virtue and suffering has negative value. But there are different ways in which these additional values can be analyzed.

Some of these values clearly call for a holistic treatment. Consider the retributive claim that a vicious person deserves pain, so it is good if he is punished. A variability interpretation will say that, when combined with his vice, his pain is transformed in value from purely evil to purely good.⁴¹ But this implies that our emotional response to his pain should be pure and simple pleasure, which is counterintuitive. The morally best response to deserved punishment is sombre, mixing pleasure that justice is being done with pain at the infliction of pain. And this can only be so if the pain, while good as deserved, remains evil as pain. But the contrary case of undeserved pleasure may demand a different treatment. Some theorists hold that when a vicious person enjoys pleasure, his pleasure is not good as pleasure and therefore is not good in any respect. Its being combined with vice destroys the pleasure's value, -- a variability claim. Other theorists reject this claim, holding that undeserved pleasure is still good as pleasure, though less good than

it is evil as undeserved and therefore evil on balance. This view allows a consistent holistic treatment of desert-values, but the competing view that undeserved pleasure is not good as pleasure requires mixing a variability treatment of one such value with a purely holistic treatment of the others.

Similar issues arise for some intentional pleasures and pains, those involving virtue and vice. (This is why these hedonic states raise more complex issues than simple ones.) Consider compassionate pain, or pain at another's pain. It is not strictly speaking an organic unity, since instead of combining two separate states by a relation it is a single state with two aspects. But it is like deserved pain in that these two aspects retain when together the same value they would have if apart, so while good as compassionate it is also evil as pain. This is why we sometimes do not share our bad news with friends, to spare them the pain of sympathizing with us. But some take a contrary view of sadistic pleasure, holding that like undeserved pleasure it has no value as pleasure. For them the presence of vice destroys the goodness of the pleasure, leaving the sadism purely and simply evil. Their view again requires combining a variability treatment of one value in a family with a non-variability treatment of the others.

Other variability claims emerge if a theory treats some values as agent-relative, so they are greater from some people's point of view than from others'. So far we have tacitly assumed that all values are agent-neutral, so they all make the same claim on all agents. If one person's pleasure is good, it is equally good from all persons' points of view and gives them all equal reason to pursue it. But this view has what some find counterintuitive implications. Imagine that a father's child is in pain, but he knows that at the same time some other child in a distant country is experiencing equal pain. If pain is agent-neutrally evil, then each child's pain is from

his point of view equally evil and he should be equally moved by each. If he cares more about his child's pain, feeling more upset about it and trying harder to relieve it, his combination of attitudes is out of proportion to its objects' values; in a theory of virtue this means his attitudes involve at least a failing in virtue and maybe a vice.⁴² Many reject this claim, saying that a father should care more about his child's pain and is seriously failing in virtue if he does not prefer it to a stranger's. A value theory can capture their view if it makes the agent-relative claim that from a parent's point of view his child's pain is a greater evil than a stranger's. Perhaps all pain has some agent-neutral value, so everyone has some reason to relieve it, but a child's pain makes greater demands on her parent.

Whether a theory can make this claim, however, depends on how it understands the properties of goodness and evil. If these are simple, unanalyzable properties, it is hard to see how a state of affairs can have one of them "from one point of view" but not "from" another. Surely it must either have the property or not. But there is no such difficulty if the good is analyzed as that which it is correct to love or which people have reason to desire and pursue, since it can be correct for a person to care more about his child's pain or he can have stronger reason to relieve it. In fact, given either of these analyses many agent-relativities about value are possible. From each person's point of view, the goods of people who stand in many relations to him, including those of spouse, friend, and fellow-citizen, can be greater than the otherwise similar goods of people who do not. Of course, the strength of the agent-relativities generated by these different relations may be different; thus, the degree of preference one should show a fellow-citizen may be less than for one's child. But they will all make some goods count more from a given person's point of view than others. And although this is not a standard case of organic unities, it does

involve a variability claim. If the value for me of a given person's pain depends on that person's relation to me, then the pain's value does not depend just on its intrinsic properties but can vary as its relations do.

The principle of organic unities complicates the theory of value, allowing the goodness or evil contributed by a state of affairs to be affected by its relations to other states or by the relations between its parts. But it also enriches the theory, since those relations often do have intuitive moral significance.

6. Environmental Values

There remains a final possible extension of the realm of value, to non-human animals and even non-animal parts of the environment. Some philosophers reject these extensions, saying that intrinsic values are restricted to humans or, more plausibly, to beings with conscious minds, so there can be nothing good or evil in a world without such minds. But if this is true it can only be so substantively rather than as following from the nature or definition of value; there is nothing self-contradictory about ascribing value to non-mental states. And if it is substantive the claim can also be denied, as it is by views that posit distinctively environmental values.

The least contentious such values are the pleasures and pains of animals that can feel them or the satisfaction of their impulses or desires. These values have been recognized by classical hedonists; thus, Bentham said the important question about animals is not "Can they reason?" but "Can they suffer?"⁴³ Animals may not be capable of as varied or as intense pleasures as humans, but when their pleasures and especially pains are similarly intense they should have similar value. Perhaps there is some agent-relativity here, so from a human's point

of view the pain of other humans counts somewhat more, but the pains of other species are still to some significant degree evil, and their prevention can justify the sacrifice of some human good.

There can also be perfectionist values in animals and even in plants. If there are human excellences such as knowledge and achievement, there can be animal excellences such as speed, robust health, and skill in hunting prey. More generally, if it is good for humans to develop their nature, understood as involving the properties essential to and/or distinctive of humans, it should likewise be good for other species to develop their natures.⁴⁴ This has again been recognized, with many classical perfectionists situating their claims about human nature within a broader scheme in which the good of all living things consists in developing their natures. This scheme raises the question of how the perfections of different species compare with each other, and here different views are possible. A radical species-egalitarianism says the development of each species's nature is exactly as good as the development of any other's, so any preference among species is immoral.⁴⁵ But the more traditional view, expressed in the idea of a "great chain of being," is that some species' development is greater or higher than others'. If their natures involve more complex or sophisticated capacities, their perfections have more value.⁴⁶ This view again implies that human goods count more than non-human ones, but it still grants some value to the latter. A world in which even lower animal and plant species flourish is better than one in which they do not, and the achievement of this good can justify some sacrifices by humans.

There can also be value in the variety of non-human species. Many environmentalists hold that a world with many different animal and plant species is better than one that is biologically less diverse, so the preservation of diversity is an important goal. This value can be

seen to arise in the aggregation of species-values, so the preservation of a species counts for more the fewer similar species exist, but it is a further addition to the list of non-human goods.

This last value of variety shades into more holistic ones. Some environmentalists find a prime intrinsic value not in individual plants or animals but in integrated ecosystems such as the tropical rain forest. Here the symbiotic relationships between living things, or the ways their different activities support each other and the life of the whole, are the prime ground of value and individual organisms only means to this larger good. The classic expression of this holistic view is Aldo Leopold's claim that something is right if it tends to preserve "the integrity, stability, and beauty of the biotic community."⁴⁷ This environmental holism is often regarded as an extravagant view, since it locates value far from individual states of individual minds. But if the good is just what we ought to desire and pursue, there is no reason in principle why it cannot attach to environmental wholes. In fact, the value of these wholes can be seen as just another instance of the organic unity or unity-in-difference that has been valued within states of mind such as knowledge and achievement. Of course a view that values ecological wholes must weigh their goods against the goods, assuming it recognizes them, of individual non-human organisms and of humans and their mental lives. But though this cannot be done precisely, it surely can be done in some cases. Thus, preserving an ecosystem can sometimes justify culling a particular species whose population growth threatens the ecological balance, and can also justify the human efforts this culling requires.

The importance of non-human values can be illustrated by two thought-experiments. Imagine that you are the last human, the last member of the species before it goes extinct, and can either detonate a bomb that will destroy all life forever or refrain and leave behind you the

existing rich variety of plants and animals. If all values were located in human minds there would be no reason to prefer one of these choices to the other. Yet many will hold, as the more radical environmental views do, that there is a strong moral duty not to detonate the bomb.⁴⁸ Or imagine that humans can, at no cost to their own well-being, implant life on a previously barren planet such as Mars and then let it develop by its own process of natural selection. On the narrower view there is no reason to do this; on the broader environmentalism, anything that increases the life, variety, and vibrancy in the universe makes the universe better.⁴⁹

7. Conclusion

The theory of value is important for several reasons. It gives content to our duty to promote good and prevent evil and perhaps also to stronger duties not to destroy good directly. It also matters in itself, determining apart from any issues about duty which states of affairs are intrinsically desirable or make the world better. But the topic of value is also one where many different views are at least initially attractive. Some of these views value competing states of human minds, such as pleasure, knowledge, and virtue; others value patterns of distribution across these states, such as equality or the proportioning of happiness to virtue; yet others compare or aggregate goods differently, while a final group value states of the non-human environment. The debate between these views is not easily resolved, but its sharpness, and the way the competing positions all make plausible claims, only underscores the importance and fascination of issues about intrinsic value.

Endnotes

1. See, e.g., W. D. Ross, *The Right and the Good* (Oxford: Clarendon Press, 1930), Ch. 2.
2. John Finnis, *Natural Law and Natural Rights* (Oxford: Clarendon Press, 1980), Ch. 5.
3. G.E. Moore, *Principia Ethica* (Cambridge: Cambridge University Press, 1903), Ch. 1; Ross, *The Right and the Good*, Ch. 4.
4. Franz Brentano, *The Origin of Our Knowledge of Right and Wrong*, trans. Roderick M. Chisholm and Elizabeth Schneewind (London: Routledge & Kegan Paul, 1969), p. 18; Henry Sidgwick, *The Methods of Ethics*, 7th ed. (London: Macmillan, 1907), p. 112.
5. G. E. Moore, “The Conception of Intrinsic Value,” in his *Philosophical Studies* (London: Routledge & Kegan Paul, 1922); and *Principia Ethica*, pp. 93, 95, 187.
6. Immanuel Kant, *Groundwork of the Metaphysics of Morals*, trans. Mary Gregor (Cambridge: Cambridge University Press, 1997), p. 7.
7. Jeremy Bentham, *Introduction to the Principles of Morals and Legislation*, ed. J. H. Burns and H. L.A. Hart (London: Methuen, 1970), Ch. 1; Sidgwick, *The Methods of Ethics*, Bk. III, Ch. 14.
8. Shelly Kagan, “The Limits of Well-Being,” *Social Philosophy and Policy* 9 (1992), pp. 172-73.
9. James Griffin, *Well-Being: Its Meaning, Measurement, and Moral Importance* (Oxford: Clarendon Press, 1986), pp. 7-10.
10. Bentham, *Introduction to the Principles of Morals and Legislation*, Ch. 4.
11. Jamie Mayerfeld, *Suffering and Moral Responsibility* (New York: Oxford University Press, 1999), Ch. 6.

12. Robert Nozick, *The Examined Life: Philosophical Meditations* (New York: Simon and Schuster, 1989), Ch. 10; L. W. Sumner, *Well-Being, Happiness, and Ethics* (Oxford: Clarendon Press, 1996), Ch. 6.
13. Robert Nozick, *Anarchy, State, and Utopia* (New York: Basic Books), pp. 42-45.
14. C. D. Broad, *Five Types of Ethical Theory* (London: Routledge & Kegan Paul, 1930), pp. 234-35; R. B. Brandt, *Ethical Theory* (Englewood Cliffs, NJ: Prentice-Hall, 1959), pp. 315-18.
15. John Rawls, *A Theory of Justice* (Cambridge, MA: Harvard University Press, 1971), pp. 395-433; Griffin, *Well-Being*, Ch. 1-2; Derek Parfit, *Reasons and Persons* (Oxford: Clarendon Press, 1984), pp. 494-99.
16. Parfit, *Reasons and Persons*, p. 494.
17. R. B. Brandt, *A Theory of the Good and the Right* (Oxford: Clarendon Press, 1979), Ch. 6; Griffin, *Well-Being*, Ch. 1-2.
18. Peter Railton, "Facts and Values," *Philosophical Topics* 14 (1986), p. 16.
19. Thomas Hurka, *Perfectionism* (New York: Oxford University Press, 1993), Ch. 8-10.
20. Hastings Rashdall, *The Theory of Good and Evil*, (London: Oxford University Press, 1907), vol. 1, p. 59, vol. 2, pp. 41-42; Moore, *Principia Ethica*, pp. 204, 208-09, 211, 217; Ross, *The Right and the Good*, pp. 134, 160; Thomas Hurka, *Virtue, Vice, and Value* (New York: Oxford University Press, 2001).
21. Ross, *The Right and the Good*, p. 141; Hurka, *Virtue, Vice, and Value*, pp. 35-36.
22. Hurka, *Perfectionism*, Ch. 2.
23. Aristotle, *Nicomachean Ethics*, trans. W. D. Ross and J. O. Urmson, in *The Complete Works of Aristotle*, ed. Jonathan Barnes (Princeton, NJ: Princeton University Press, 1984), Bk. I,

Ch. 7.

24. Karl Marx, *Economic and Philosophical Manuscripts* (excerpt), in *Karl Marx: Selected Writings*, ed. David McLellan (Oxford: Oxford University Press, 1977).

25. Friedrich Nietzsche, *Beyond Good and Evil*, trans. Walter Kaufmann (New York: Vintage, 1966), sec. 36, 186, 259; *The Will to Power*, trans. Walter Kaufmann and R. J. Hollingdale (New York: Vintage, 1968), sec. 55, 693.

26. F. H. Bradley, *Ethical Studies*, 2nd ed. (Oxford: Clarendon Press, 1927), pp. 74-81, 247-49, and *Appearance and Reality*, 2nd ed. (Oxford: Clarendon Press, 1897), Ch. 24-25; Robert Nozick, *Philosophical Explanations* (Cambridge, MA: Harvard University Press, 1981), pp. 415-50.

27. To say that rationality involves knowing about and pursuing the good is to give up the ambition to ground all goods in human nature; on the contrary, it is to assume there are goods independent of this nature and constraining its content.

28. John Stuart Mill, *Utilitarianism*, ed. George Sher (Indianapolis: Hackett, 1979), pp. 7-11. See also William K. Frankena, *Ethics*, 2nd ed. (Englewood Cliffs, NJ: Prentice-Hall, 1973), pp. 89-92; and Parfit, *Reasons and Persons*, pp. 501-02.

29. Ross, *The Right and the Good*, pp. 150, 152.

30. Hurka, *Virtue, Vice, and Value*, Ch. 5.

31. Hurka, *Perfectionism*, Ch. 7.

32. Jan Narveson, "Utilitarianism and New Generations," *Mind* 76 (1967).

33. J. M. E. McTaggart, *The Nature of Existence* (Cambridge: Cambridge University Press, 1921), vol. 1, pp. 452-53; Parfit, *Reasons and Persons*, pp. 387-90.

34. Parfit, *Reasons and Persons*, p. 420.

35. Derek Parfit, "Equality or Priority?", *The Lindley Lecture* (Lawrence, KS: University of Kansas Press, 1995); Larry S. Temkin, *Inequality* (New York: Oxford University Press, 1993).
36. Nietzsche, *Beyond Good and Evil*, sec. 126, 199, 257, 258, 260, 265; *The Will to Power*, sec. 246, 252, 373, 660, 681, 766, 877, 881, 987, 997.
37. Dennis McKerlie, "Equality and Time," *Ethics* 99 (1989).
38. Moore, *Principia Ethica*, pp. 27-36, 92-96, 183-225.
39. Thomas Hurka, "Two Kinds of Organic Unity," *The Journal of Ethics* 2 (1998).
40. Joel Feinberg, "Justice and Personal Desert," in his *Doing and Deserving* (Princeton, NJ: Princeton University Press, 1970); George Sher, *Desert* (Princeton, NJ: Princeton University Press, 1987); Shelly Kagan, "Equality and Desert," in Louis P. Pojman and Owen McLeod, eds., *What Do We Deserve? Readings on Justice and Desert* (New York: Oxford University Press, 1999).
41. Fred Feldman, "Adjusting Utility for Justice: A Consequentialist Reply to the Objection from Justice," in his *Utilitarianism, Hedonism, and Desert* (Cambridge: Cambridge University Press, 1997).
42. Michael Stocker, *Valuing Emotions* (Cambridge: Cambridge University Press, 1996), p. 319.
43. Bentham, *Introduction to the Principles of Morals and Legislation*, pp. 282-83n; Peter Singer, *Animal Liberation* (New York: New York Review of Books, 1975).
44. Robin Attfield, *A Theory of Value and Obligation* (London: Croom Helm, 1987), Ch. 3.
45. Paul Taylor, *Respect for Nature: A Theory of Environmental Ethics* (Oxford: Clarendon Press, 1996), Ch. 6.

46. Arthur O. Lovejoy, *The Great Chain of Being: A Study of the History of the Idea* (Cambridge, MA: Harvard University Press, 1936); Attfield, *A Theory of Value and Obligation*, Ch. 5.

47. Aldo Leopold, *A Sand County Almanac* (London: Oxford University Press, 1949), p. 224-25; J. Baird Callicott, *In Defense of the Land Ethic: Essays in Environmental Philosophy* (Albany, NY: State University of New York Press, 1989).

48. Richard Routley, "Is There a Need for a New, an Environmental Ethic?", in Bulgarian Organizing Committee, ed., *Proceedings of the Fifteenth World Congress of Philosophy I* (Sofia, Bulgaria: Sophia-Press, 1973), p. 207.

49. Christopher P. McKay and Robert H. Haynes, "Should We Implant Life on Mars?", *Scientific American* (Dec., 1990), p. 144.