

# Indirect, multidimensional consequentialism

*Abstract:*

*In this three part article I briefly present, in the first part, an outline of a revisionary reading of John Stuart Mill's normative theory in order to demonstrate that a deontic injunction to maximize overall value does not necessarily follow directly from a maximizing axiology. In the second part, I outline a different, and superior, axiology to Mill's: namely, a multidimensional one—in other words, one that recognizes a plurality of contributory values. And in the third part, I describe some features that one could expect the moral rules grounded by such a multidimensional axiology to display.*

**Words in main text: 6,890**

**Words in abstract: 97**

**Words in notes: 1,798**

## I

Consequentialist normative theories once dominated moral philosophy and once provided the often unquestioned background to debates within the domain of public policy. Today, however, consequentialism seems to be in full retreat. But the fundamental reason for its relative demise may be less to do with any inherent flaw within consequentialism *per se* and more to do with the one-dimensional axiologies at the heart of popular versions of consequentialism. Hence, if consequentialism is to regain its formerly ascendant position, it may well require a far richer axiology.

Now, there can be little doubt that John Stuart Mill bequeathed to the world the most famous version of consequentialism. But there is reason to think that his normative theory is both more complex and more interesting than is generally realized. And, as others have pointed out,<sup>1</sup> the key to understanding his theory may well lie in his *System of Logic*, where he insists that “the imperative mood is the characteristic of art, as distinguished from

---

<sup>1</sup>See, in particular, Alan Ryan, *J. S. Mill* (London: Routledge, 1974) and John Gray, *Mill on Liberty: A Defense*, 2<sup>nd</sup> ed. (London: Routledge, 1996).

science. Whatever speaks in rules or precepts, not in assertions respecting matters of fact, is art; and ethics or morality is properly a portion of the art corresponding to the sciences of human nature and society.”<sup>2</sup> This dense passage has proven to be a goldmine for those seeking to revise and improve our understanding of Mill.

One striking feature of Mill’s utilitarianism that this passage and the rest of the last chapter of the last book of his *System of Logic* seems to reveal is that morality is only a portion of what he calls “the art of life”.<sup>3</sup> It is that portion which corresponds to a particular science, which for Mill was “ethology”, or the science of how people are socialized. And it goes without saying that limitations on the ability to socialize individuals will have a bearing on what a workable morality will realistically demand of them. Not surprisingly, then, Mill had hoped to make greater headway in, what he took to be, the crucial empirical science of ethology than he succeeded in attaining.

But what is this art of life to which he refers? For Mill, it is that “body of doctrine” comprising three “departments,” which consist of “Morality, Prudence or Policy, and Aesthetics.” Or put in other terms, the art of life comprises “the Right, the Expedient, and the Beautiful or Noble, in human conduct and works.”<sup>4</sup> Moreover, art, as Mill has informed us, “speaks in rules or precepts,” which appears to imply that all of the art of life is rule-governed; and, he adds, “the general principle to which all rules of practice ought to conform, and the test by which they should be tried, is that of conduciveness to the happiness of mankind, or rather, of all sentient beings: in other words, that the promotion of happiness is the ultimate principle of Teleology.”<sup>5</sup> The principle of utility, then, applies not merely to one department—morality—as we might have expected; rather, it applies to the whole of the art of life.<sup>6</sup> (See Table 1.)

But this leads to what is perhaps the most striking feature of Mill’s normative theory, as apparently revealed by his *System of Logic*: for, as Alan Ryan and John Gray have both

---

<sup>2</sup>John Stuart Mill, *The Logic of the Moral Sciences* (London: Duckworth, 1987), p. 134, which comprises the Sixth Book of J. S. Mill, *A System of Logic*, first published in 1843.

<sup>3</sup>*Ibid.*, p. 140.

<sup>4</sup>*Ibid.*, pp. 140–1.

<sup>5</sup>*Ibid.*, pp. 142–3.

<sup>6</sup>Ryan: “the principle of utility is supposed to serve as the first principle not just of morality but also of the rest of the art of life.” Ryan, *J. S. Mill, op. cit.*, p. 104.

noticed, all of this seems to entail that, for Mill, *the greatest happiness principle is not a moral principle*, contrary to what is popularly assumed. Why can the principle of utility not be a moral principle? Because, as it governs the whole of the art of life, all three of its departments would fall within the bounds of morality if the greatest happiness principle were a moral principle. But Mill is explicit in claiming that morality is only a portion of the art of life. Hence the principle of utility cannot be a moral principle.

What, then, are moral principles? According to Mill, moral rules or precepts, unlike the rules that apply to policy or aesthetics, are distinctive in being obligatory: in short, one is liable to face certain forms of sanction if one does not abide by them. As he writes in his famous essay “Utilitarianism”:

We do not call anything wrong, unless we mean to imply that a person ought to be punished in some way or other for doing it; if not by law, by the opinion of his fellow-creatures; if not by opinion, by the reproaches of his own conscience. This seems the real turning point of the distinction between morality and simple expediency. It is a part of the notion of Duty in every one of its forms, that a person may rightfully be compelled to fulfill it. Duty is a thing which may be *exacted* from a person, as one exacts a debt. Unless we think that it may be exacted from him, we do not call it his duty.<sup>7</sup>

So, actions that are morally wrong render one liable to certain forms of sanction. But breaking prudential or aesthetic rules are not of this sort. Yet they, too, are governed by the principle of utility. And as surprising as this will no doubt appear to those who are only familiar with traditional interpretations of utilitarianism, what all of this seems to imply is that the failure to maximize utility is not always morally wrong.

To make this clearer: If a composer, for example, were to spoil her musical composition by breaking an aesthetic rule, say a rule of harmony, she would thereby have failed to maximize utility; for in breaking that rule she would have composed a piece of music that gave everyone far less pleasure than they would have enjoyed had that rule not been infringed. But a composer is not morally required to adhere to the rules governing aesthetics. She is, however, obliged to obey moral rules. Now, it is true that our admiration for a

---

<sup>7</sup>John Stuart Mill, “Utilitarianism” in *Utilitarianism, On Liberty and Essay on Bentham*, ed. M. Warnock (London: Fontana, 1962), pp. 303–4.

composer may diminish if she keeps breaking aesthetic rules, and she might feel shame in having damaged her reputation within the world of music. But she would not be blamed for violating an aesthetic rule, and nor would we expect her to feel guilt for so doing. But if she breaks a moral rule, instead, then we would expect her to feel guilt, and she would be blamed for infringing it.

What is the explanation for this key difference between morality and the other two departments of the art of life? The answer is to be found in the importance of obedience to moral rules for the advancement of utility. Deterring violations of certain rules—such as the rule against killing innocent persons—is expedient for increasing the greatest happiness of the greatest number. And violations can be deterred by the threat of punishment, even if the punishment for breaking a rule is only the reproaches of one’s own conscience. So, the greatest happiness is advanced by threatening to punish violations of certain rules. But the greatest happiness is not advanced by threatening to punish violations of aesthetic or prudential rules. It is for this reason that only moral rules are obligatory, while aesthetic and prudential rules are not. In other words, as Gray explains, “[t]hat an act is maximally expedient is not, according to Mill, sufficient to show that it would be morally right or obligatory to do it: it must also be true that it is maximally expedient to punish that act’s non-performance.”<sup>8</sup> And recall: an action is morally wrong only if that action is liable to incur some form of punishment.

This is why it is not always immoral to fail to maximize utility. Composers do not act immorally simply by failing to produce the best piece of music they are capable of composing, even though their failing to do so deprives the world of great enjoyment. And this is because utility is not maximized by inculcating within composers a tendency to feel guilty whenever they break aesthetic rules. Nor is utility maximized by making composers liable to be blamed or to receive other forms of punishment whenever they violate an aesthetic rule. There are rewards enough for those successful within the arts to spur on composers to write sublime music. And in any case, composing music offers its own intrinsic delights. But as obedience to moral rules is far more important for the maximization of utility than obedience to aesthetic rules, other means are required to ensure their compliance. Clearly, we will not make the world a happier place if all composers live

---

<sup>8</sup>Gray, *Mill on Liberty*, *op. cit.*, p. 28.

in fear of their failing to produce the perfect symphony, say. But the world will be a happier place if potential murderers, for example, so fear the consequences of committing murders that they refrain from murderous acts. So, we deter potential murderers with the threat of punishment. When one breaks a rule that is such that breaking that rule makes one liable to incur some form of punishment one acts immorally. Thus, acts of murder are immoral, while producing a flawed symphony is not, even though there would be more happiness in the world if the musical composition had been better, and thus capable of giving more pleasure to its listeners. Consequently, as Gray summarizes: “[n]ot its maximal expediency alone, but only its maximal expediency plus the maximal expediency of making the failure to do it punishable shows an act to be morally right.”<sup>9</sup>

But if the greatest happiness principle is not actually a moral principle, for we are not always liable to be punished whenever we fail to act in conformity with it, then what, precisely, is it? Rather than being a principle of right action, which is how it is usually construed, it is, instead, a principle of general evaluation. In other words, the greatest happiness principle is not a deontic principle, for it does not tell us what to do;<sup>10</sup> rather, it is an axiological principle: it only tells us what is good.<sup>11</sup> Consequently, as Gray observes: “Because utility does not of itself impose moral requirements upon action, it is mistaken to think that a man must do wrong when he fails to maximize it.”<sup>12</sup>

Now, while Gray maintains that Mill is therefore an indirect utilitarian, he denies that Mill is a rule utilitarian.<sup>13</sup> Mill is an indirect utilitarian because the direct pursuit of happiness is self-defeating, in his view. However, because morality is a portion of the art of life, and because the art of life “speaks in rules or precepts,” then it seems mistaken to think that Mill is not, at the very least, some kind of rule utilitarian. In fact, Mill appears to be both

---

<sup>9</sup>*Ibid.*, p. 29.

<sup>10</sup>As Don Brown observes: “We can at least agree that the Principle of Utility is not a principle about what actions are wrong.” D. G. Brown, “Mill on liberty and morality,” *The Philosophical Review* 81, 2 (1972): 133–158, here at p. 157.

<sup>11</sup>Thus, at the very end of “Utilitarianism,” Mill insists that a distinction must be drawn between “the milder feeling which attaches to the mere idea of promoting human pleasure or convenience” and justice, which is “more absolute and imperative.” Mill, “Utilitarianism,” *op. cit.*, p. 321.

<sup>12</sup>Gray, *Mill on Liberty*, *op. cit.*, p. 27.

<sup>13</sup>*Ibid.*, ch. 2, § 2.

an indirect utilitarian in Gray's sense *and* a rule utilitarian. He is the former given that he holds the greatest happiness to be most effectively advanced by not prescribing or proscribing purely self-regarding actions.<sup>14</sup> But he is also the latter given that he holds certain forms of other-regarding actions to be subject to moral rules. As Mill writes in a passage that deserves to be far more widely known than it is:

First, [the general happiness] requires that each shall consider it his special business to look after himself.... The good of all can only be pursued with any success by each person's taking as his particular department the good of the only individual whose requirements he can thoroughly know; with due precautions to prevent these different persons, each cultivating a particular strip of the field, from hindering one another. Secondly, human happiness, even one's own, is in general more successfully pursued by acting on general rules, than by measuring the consequences of each act; and this is still more the case with general happiness, since any other plan would not only leave everybody uncertain what to accept, but would involve perpetual quarrelling and hence general rules must be laid down for people's conduct to one another, or in other words, rights and obligations must...be recognized; and people must, on the one hand, not be required to sacrifice even their own less good to another's greater, where no general rule has given the other a right to the sacrifice; while, when a right *has* been recognized, they must, in most cases, yield to that right even at the sacrifice, in the particular case, of their own greater good to another's less.<sup>15</sup>

In short, an action is only morally wrong for Mill when (a) it adversely affects another person, and (b) it violates a certain sort of rule—namely, one where utility is maximized by punishing its infractions.<sup>16</sup>

So, the principle of utility, as an axiological principle, tells us which is the best world: it is, in Mill's view, that world containing the greatest balance of pleasure over pain—or the greatest total quantity of happiness. But the principle of utility does not tell us to bring about

---

<sup>14</sup>See John Stuart Mill, "On Liberty" in *Utilitarianism, On Liberty and Essay on Bentham*, *op. cit.*

<sup>15</sup>John Stuart Mill, letter to George Grote, January 10<sup>th</sup>, 1862, quoted in John Skorupski, *John Stuart Mill* (London: Routledge, 1989), p. 318.

<sup>16</sup>See Alan Carter, "Is utilitarian morality necessarily too demanding?" in Timothy Chappell (ed.), *The Problem of Moral Demandingness: New Philosophical Essays* (Basingstoke: Palgrave, forthcoming).

that world directly.<sup>17</sup> What tells us directly what to do are the moral rules. And following those rules is necessary if the best world is to be brought about. Importantly, utilitarianism thus construed sidesteps many of the standard objections leveled against it.<sup>18</sup>

## II

While this form of consequentialism certainly constitutes an impressive normative theory, it contains at its heart an inadequate axiology. And its axiology is inadequate in two ways: (1) it fails to acknowledge important values; and (2) it fails to recognize that utility does not contribute value invariably.

In order to appreciate why the axiology at the heart of utilitarianism is inadequate, consider several possible worlds. The first of these— $W_1$ —has within it an extremely large number of people all with worthwhile lives, though at a very low level of average utility.<sup>19</sup> Clearly, this possible world would be better if, *ceteris paribus*, the level of average utility were to be raised. The level of average utility would, therefore, appear to be something that contributes to the value of a world.

But now consider a second possible world:  $W_2$ . This has within it an extremely small number of people all with worthwhile lives at a very high level of average utility. Surely, this possible world would be better if, *ceteris paribus*, the number of worthwhile lives were to be increased. The number of worthwhile lives would, therefore, appear to be

---

<sup>17</sup>Moreover, as Henry Sidgwick remarks: “[T]he doctrine that Universal Happiness is the ultimate *standard* must not be understood to imply that Universal Benevolence is the only right or always best *motive* of action. For...it is not necessary that the end which gives the criterion of rightness should always be the end at which we consciously aim: and if experience shows that the general happiness will be more satisfactorily attained if men frequently act from other motives than pure universal philanthropy, it is obvious that these other motives are reasonably to be preferred on Utilitarian principles.” Henry Sidgwick, *The Methods of Ethics*, 7<sup>th</sup> ed. (Indianapolis: Hackett, 1981), p. 413.

<sup>18</sup>See Carter, “Is utilitarian morality necessarily too demanding?” *op. cit.*

<sup>19</sup>This is, of course, the famous Repugnant Conclusion. See Derek Parfit, “Overpopulation and the quality of life” in Peter Singer (ed.), *Applied Ethics* (Oxford: Oxford University Press, 1986), p. 148–50. Also see Alan Carter, “Moral theory and global population,” *Proceedings of the Aristotelian Society* 99, 3 (1999): 289–313.

something that contributes to the value of a world.

However, many resist the claim that increasing the number of worthwhile lives positively affects the value of the world. But this may be due to an implicit association of population growth with overcrowding. However, an increase in the number of worthwhile lives can occur over time by humanity surviving longer. Every worthwhile life does not have to be crammed into a world at the same time. Indeed, if, like our world,  $W_2$  is subject to environmental constraints, then as more people can live over time than at any one time, trying to cram too many worthwhile lives into  $W_2$  at the same time could so damage its ecosystems that no generations follow the overcrowded one, and fewer worthwhile lives would then be enjoyed overall.<sup>20</sup> Hence, the number of worthwhile lives should be measured diachronically and not just synchronically.

Moreover, it is not the small total quantity of utility that seems to be the only reason for the “perception” of disvalue in  $W_2$ , for were there only one person—an unimaginably happy one—living within that world, there would be both a very high level of average utility and a very large total quantity of utility. But it would, surely, still be the case that  $W_2$  would be a better world if there were more people living within it; and, surely, this would be so even if each person were less happy than the one unimaginably happy person so that the total quantity of utility remained the same. In short, there is reason enough to think that the number of worthwhile lives contributes to the value of a world.<sup>21</sup>

Now consider a third possible world:  $W_3$ . Within  $W_3$  is an extremely large number of people all with worthwhile lives at a very high level of average utility (factors which together entail a very large total quantity of utility), but where the utility is distributed very unequally. So, there is within this possible world both a very high level of average utility and a very large total quantity of utility. But, surely,  $W_3$  would be a better place if, *ceteris paribus*, the inequality were to be reduced. Equality would, therefore, appear to be something that contributes to the value of a world.

So, we have reason to hold that three different things, at least, contribute to the value of a world: the level of average utility, the number of worthwhile lives, and equality. But

---

<sup>20</sup>See Robin Attfield, *The Ethics of Environmental Concern*, 2<sup>nd</sup> ed. (Athens, GA: University of Georgia Press, 1991), pp. 127–8.

<sup>21</sup>See Alan Carter, “A plurality of values,” forthcoming.

all versions of hedonistic utilitarianism—including varieties of sophisticated indirect, rule utilitarianism, such as Mill’s—mistakenly hold that only one thing contributes to the value of a world: either the total quantity or the average level of utility. But this completely fails to explain our considered response to  $W_3$ .

Moreover, it is more important to raise the level of average utility when it is at a very low level than when it is at a very high level. And it is more important to increase the number of worthwhile lives when there are very few people than when there are very many. In addition, it is more important to reduce inequality when distributions are very unequal than when they are not. But this seems to imply that an increase in the level of average utility contributes more value to a world when it is at a low level than when it is at a high level, that an increase in the number of worthwhile lives contributes more value to a world when that number is small than when it is already large, and that an increase in equality contributes more value to a world when there is great inequality than when there is not.<sup>22</sup> How might we combine all of this together into a coherent axiology?

Begin by calling the level of average utility, the number of worthwhile lives and equality “*contributory values*”. And call the kind of value, all things considered, that they contribute towards “*overall value*”. As contributory values contribute more overall value the less they are realized, such values contribute *diminishing marginal overall value*. Hence, what they contribute is variable; they are *variable* values. (Contrast this with utilitarianism, which mistakenly holds not only that one thing alone—either the total quantity or average level of utility—contributes to the value of a world but also that the value of a world invariably increases in direct proportion to increases in that value.)

Given that contributory values are variable, what if we were unable to raise the level of average utility in  $W_1$  with so many people living there? And what if we could choose, without killing anyone, to bring into existence either  $W_1$  or a possible world with a higher level of average utility than  $W_1$  but with fewer people living within it? We would need to know which would be the better world. In other words, once we recognize that there are a plurality of contributory values, we immediately face the problem of trading them off.<sup>23</sup>

---

<sup>22</sup>See Alan Carter, “Some groundwork for an alternative axiology,” forthcoming.

<sup>23</sup>“Here I set aside the question of how we are able to trade off one value against another, for one’s answer will, at least in part, depend upon one’s metaethic, and our task here lies purely within normative ethics.

To gain a better purchase on what this involves, imagine that we could plot the graph in Figure 1.<sup>24</sup> The horizontal  $x$ -axis measures the number of worthwhile lives, while the vertical  $y$ -axis measures the level of average utility. Because each contributory value contributes diminishing marginal overall value, then the isovalue-curve  $AB$ , where all of the possible worlds represented by points falling upon that curve possess the same overall value, is convex when viewed from the origin  $O$ .<sup>25</sup> The isovalue-curve  $DE$ , where all of the possible worlds represented by points falling upon  $DE$  also possess the same overall value, is similarly convex when viewed from the origin. But each possible world represented by a point falling upon  $DE$  possesses more overall value than any world represented by a point falling upon  $AB$ , for  $DE$  is further to the north-east from the origin than  $AB$ .<sup>26</sup>

How could we identify the best world that it is within our power to bring about—namely, the one containing the most overall value? Figure 1 merely tells us how valuable possible worlds would be if they existed. It tells us nothing whatsoever about which worlds we could bring into being. So, imagine that we could plot the graph in Figure 2, where, as in Figure 1, the  $x$ -axis measures the number of worthwhile lives, while the  $y$ -axis measures the level of average utility. The curve  $PQ$  is the frontier of the points representing all of the worlds that we could practicably bring about. Call  $PQ$  “the

---

This notwithstanding, the lack of an answer to the question of how we are able to trade off values should not be allowed to stand in the way of our developing a normative theory, for it is a simple fact that we *are* prepared to trade off values when deciding which would be the best world to bring about.... Once it is granted that we do engage in such trade offs, then there is nothing preventing us from incorporating that feature into our axiology; and we should certainly not exclude it merely because, as yet, we lack an agreed answer to the puzzle of how we manage to trade off values.” *Ibid.* For critiques of arguments against the comparability or, alternatively, commensurability of values, see Ruth Chang, “Against constitutive incommensurability or buying and selling friends,” *Philosophical Issues* 11 (2001): 33–60, especially pp. 51–2, and Donald Regan, “Value, comparability, and choice” in Ruth Chang (ed.), *Incommensurability, Incomparability, and Practical Reason* (Cambridge, MA: Harvard University Press, 1997).

<sup>24</sup>Nb.: I am not assuming that Figure 1 could ever be plotted. What follows in the remainder of this section forms a kind of thought-experiment.

<sup>25</sup>For a full explanation of why the isovalue-curve  $AB$  is convex when viewed from the origin, see Carter, “Some groundwork for an alternative axiology,” *op. cit.*

<sup>26</sup>If it is not immediately apparent why this is so, see Carter, “A plurality of values,” *op. cit.*

practicability frontier". All of the worlds that we could practicably bring about, then, are represented by points falling either on  $PQ$  or between  $PQ$  and the origin  $O$ . (In other words, they all fall upon the area  $PQO$ .) But we cannot tell from Figure 2 which is the best world.

We can identify the best practicable world, however, if we superimpose Figure 2 onto Figure 1. In so doing, we thereby obtain Figure 3. Point  $T$  in Figure 3 falls upon the practicability frontier  $PQ$ . But it also falls upon an isovalue-curve—namely,  $DE$ —that is further to the north-east from the origin  $O$  than any isovalue-curve (either drawn, such as  $AB$ , or not drawn) upon which falls any other point also falling upon the area  $PQO$ . In other words, of all the worlds that we could practicably bring about, the world represented by  $T$  contains the most overall value. It is, therefore, the best practicable world.

Moreover, point  $R$ , in falling upon the isovalue-curve  $AB$ , which is an isovalue-curve closer to the origin  $O$  than the isovalue-curve  $DE$ , represents a suboptimal world given that point  $T$ , which falls upon  $DE$ , represents a practicable world. Equally, point  $S$ , in also falling upon the isovalue-curve  $AB$ , represents a suboptimal world given that  $T$  represents a practicable world.  $S$  represents  $W_1$ , which has within it an extremely large number of people all with worthwhile lives, though at a very low level of average utility, while  $R$  represents  $W_2$ , which has within it an extremely small number of people all with worthwhile lives at a very high level of average utility. We can thus see why we feel that both of these possible worlds could be improved, even if that involved bringing about a world with fewer worthwhile lives than are currently are to be found within  $W_1$  or with a lower level of average utility than is currently to be found within  $W_2$ . Such a world is represented by point  $T$ .

But what about  $W_3$ , within which, it will be recalled, is to be found an extremely large number of people all with worthwhile lives at a very high level of average utility, but where the utility is very unequally distributed? If, in Figure 3, instead of the  $x$ -axis measuring the number of worthwhile lives and the  $y$ -axis measuring the level of average utility, we were to measure the total quantity of utility—the product of the level of average utility and the number of worthwhile lives—along the  $y$ -axis and equality along the  $x$ -axis, then  $W_3$  would be represented by point  $R$ . But in falling upon the isovalue-

curve  $AB$ ,  $R$  represents a suboptimal world given that point  $T$ , which falls upon the isovalue-curve  $DE$ , represents a practicable world. We can thus see why we feel that this possible world, too, could be improved, even if that involved bringing about a world with fewer worthwhile lives and a lower level of average utility than are currently to be found within  $W_3$ .

But now consider  $W_4$ , which has within it an extremely large number of people all with worthwhile lives at a very high level of average utility, and where the utility is distributed equally, but where there is very little freedom. Because there is, within this possible world, a very high level of average utility enjoyed by an extremely large number of people all with worthwhile lives, there is a very large total quantity of utility within it, and this utility is distributed equally. But  $W_4$  would be a better place if, *ceteris paribus*, freedom were to be increased. Freedom would, therefore, appear to be something that contributes to the overall value of a world.

In order to accommodate this case, imagine that we could plot the graph in Figure 4.<sup>27</sup> The horizontal  $x$ -axis measures the degree of equality, the vertical  $y$ -axis measures the total quantity of utility, while the  $z$ -axis, which should be viewed as coming out of the page, measures the degree of freedom. Because each contributory value contributes diminishing marginal overall value, then the isovalue-plane  $DEF$ , where all of the possible worlds represented by points falling upon that plane possess the same overall value, is convex when viewed from the origin  $O$ . The plane  $PQR$  is the frontier of the points representing all of the worlds that we could practicably bring about. Hence  $PQR$  is the practicability frontier within Figure 4. All of the worlds that we could practicably bring about, then, are represented by points falling either on  $PQR$  or between  $PQR$  and the origin. (Put another way, they all fall within the volume  $PQRO$ .) Point  $T$  falls upon the practicability frontier  $PQR$ . But it also falls upon an isovalue-plane—namely,  $DEF$ —that is further to the north-east and out of the page from the origin than any imaginable isovalue-plane (such as  $ABC$  in Figure 5) upon which falls any other point falling within  $PQRO$ . In other words, of all the worlds that we could practicably bring about, the world represented by  $T$  in Figure 4 contains the most overall value.<sup>28</sup> It is the best practicable

---

<sup>27</sup>Again, note that I am not assuming that Figure 4 could ever be plotted.

<sup>28</sup>However, bear in mind that it would not be if the large total quantity of utility resulted from an enormous

world.<sup>29</sup> But  $T$  does not represent  $W_4$ , for  $W_4$  has very little freedom (as measured along the  $z$ -axis), while freedom is not lacking in the world represented by  $T$ .  $W_4$  is, therefore, a suboptimal world, even though it performs well in terms of the number of worthwhile lives, the level of average utility, and equality. Indeed, we can now see that this possible world could be improved, even if that involved bringing about a world with fewer worthwhile lives, a lower level of average utility and less equality than is currently to be found within  $W_4$ .

If we put all of this together, then we have reason for holding that the best practicable world will be deficient neither in freedom, nor in equality, nor in the level of average utility nor in the number of worthwhile lives. And this means that a multidimensional axiology can be fully compatible with our intuitive responses to  $W_1$ ,  $W_2$ ,  $W_3$  and  $W_4$  if it recognizes as variable contributory values the number of worthwhile lives, the level of average utility, equality and freedom.

---

population at a low level of average utility, as in the Repugnant Conclusion. Recall that when, in Figure 3, the horizontal  $x$ -axis measures the number of worthwhile lives and the vertical  $y$ -axis measures the level of average utility, point  $S$  represents a suboptimal world. But that is precisely a world in which the large total quantity of utility results from an enormous population at a low level of average utility. Consequently, while treating the product of the number of worthwhile lives and the level of average utility—namely, the total quantity of utility—as if it were a single contributory value can be expedient in illustrating certain evaluations, the failure to resolve the total quantity of utility into its component contributory values can be highly problematic at times. Perhaps total utility could best be described as a second-order contributory value that resolves into first-order contributory values—namely, the number of worthwhile lives and the level of average utility. If so, then second-order contributory values should always be treated with caution.

<sup>29</sup>Interestingly, Mark Timmons argues that certain seemingly monistic theories are considerably more plausible when interpreted as pluralistic. See Mark Timmons, *Moral Theory* (Lanham, MD: Rowman & Littlefield, 2002), *passim*. However, he concludes from this that the most plausible normative theories will therefore necessarily display a measure of indeterminacy. The normative technology mooted here implies that this need not necessarily be the case, for the point where the practicability frontier is tangential to an isovalue-plane (as with point  $T$  in Figure 4) represents a determinate answer to the question of which is the best world to bring into existence.

### III

So, there is reason to think that the multidimensional axiology here outlined constitutes a better axiology than that at the heart of Mill's indirect, rule utilitarianism, for it explains what is bad about  $W_1$ ,  $W_2$ ,  $W_3$  and  $W_4$ , which seems beyond the capacity of any (one-dimensional) utilitarian axiology. And, of course, it also explains why the best practicable world will be deficient neither in the number of worthwhile lives, in the level of average utility, in equality or in freedom. But this is not to say that we must abandon everything in Mill. Rather, if his indirect, rule consequentialism is preferable to direct forms of consequentialism, when both have at their core the same utilitarian axiology, then even more preferable might be an indirect, rule consequentialism that has at its core a multidimensional axiology. Moreover, for the same reasons that support Mill's indirect, rule utilitarianism in preference to direct forms of utilitarianism, an indirect, rule consequentialism that has at its heart a multidimensional axiology might well be preferable to a direct consequentialism that has at its heart the same axiology.

What are some of the reasons for preferring an indirect form of consequentialism? For one thing, obeying moral rules rather than directly trying to bring about the greatest good seems closer to our commonsense intuitions regarding the nature of morality, which appears to regard moral behavior primarily as the avoidance of rule-breaking. For another, this explains the seeming plausibility in deontological approaches. For a third, obeying agreed rules rather than trying to work out for oneself what to do makes it far more likely that everyone's actions will be coordinated rather than accidentally frustrating each other's well-intentioned efforts. And for a fourth, working out the consequences of every possible action is far less efficient than obeying rules that experience has shown work well with respect to bringing about the greatest good. Indeed, very little good will result from each person having to spend most of his or her time trying to calculate what to do. And all of this is so regardless of whether the greatest good is thought to consist only in the maximization of utility or in the maximization of a balance of plural contributory values.

In what remains of this concluding section, allow me to make some preliminary remarks regarding the moral rules that a multidimensional axiology of the sort adumbrated here

might ground. First, it might be objected that if a multidimensional axiology succeeds in identifying the best practicable world, then, surely, we ought directly to try to bring it about. But arguably, for the third and fourth reasons just mentioned for preferring an indirect form of consequentialism, no axiology at the core of any respectable consequentialist theory would actually ground a direct form of consequentialism, and that is because a direct consequentialism would, almost certainly, fail in practice to bring about the greatest overall value. Moreover, the deontic sense of ought, as used in this objection, assumes that morality requires us to bring about the greatest good. But, as we saw in §1, what we ought to do in this sense is obey the rules of morality. With respect to the greatest good, it will be found in the world where the degree of realization and the balance of contributory values contribute the greatest overall value, as illustrated by point *T* in Figure 4. If we are fully rational, then we will desire to live in the world containing the greatest overall value. If that requires that we are all subject to moral rules, then, if we are fully rational, we will desire that we all be subject to those rules—rules that are such that the greatest overall value will result from infractions of those rules being liable to some form of sanction (even if that sanction is merely a bad conscience). Obeying such rules is what we morally ought to do. We would not be *morally* required to try to bring about directly the greatest overall value.

However, there might be rules that, ideally, if everyone were to obey them would result in a world with greater overall value than the world represented by point *T* in Figure 4, but which are, for example, too demanding for many people to obey. Hence, such a world would not be a practicable one. Call such rules “the best ideal rules”. If many more people would obey slightly less demanding rules, say, then a world containing more overall value might well result from their promulgation than from the promulgation of the best ideal rules that, as a matter of fact, few would obey. Call the rules that many more people would obey and whose promulgation would, in fact, bring about the best practicable world “the best practicable rules”.<sup>30</sup> It is the best practicable rules rather than the best ideal rules that the multidimensional axiology here outlined would ground, for they are the rules that would bring about the world represented by point *T* in Figure 4—that practicable world which contains more overall value than any

---

<sup>30</sup>See Carter, “Is utilitarian morality necessarily too demanding?” *op. cit.*

other practicable world. In short, there is reason for advocating “practicable rule consequentialism” in preference “ideal rule consequentialism”.<sup>31</sup>

But those who are sufficiently motivated to do more than is required by the best practicable rules may legitimately try to bring about more overall value than would result from their doing no more than the moral rules require. For one would surely expect the best practicable rules to allow one to bring about more good than they prescribe, if one were so inclined. An indirect, multidimensional consequentialism can, therefore, easily make sense of supererogatory actions.

Second, we saw in §1 that Mill seems to consider an action to be morally wrong (a) only when it adversely affects another person, and (b) only when it violates a certain sort of rule—namely, one where the greatest happiness is advanced by punishing its infractions. The limitation (a) appears to be justified in Mill’s eye because a measure of personal freedom seems necessary in order to maximize utility. An indirect, multidimensional consequentialism that acknowledges that total utility is a contributory value, or that it resolves into two different contributory values (namely, the number of worthwhile lives and the level of average utility), can easily agree with Mill, here. But an indirect, multidimensional consequentialism that also acknowledges that freedom is a contributory value in its own right has even stronger grounds for choosing moral rules that leave space for personal projects. The best practicable rules grounded by such an axiology can thus be expected to include something like agent-centered prerogatives,<sup>32</sup> which offer one answer to famous objections concerning the supposed inherently alienating aspects of consequentialism.<sup>33</sup> And this further enables an indirect, multidimensional

---

<sup>31</sup>And Sidgwick seems to have drawn a parallel distinction within utilitarianism, for he observes that “any particular existing moral rule, though not the ideally best even for such beings as existing men under the existing circumstances, may yet be the best that they can be got to obey...”. Sidgwick, *The Methods of Ethics*, *op. cit.*, p. 469.

<sup>32</sup>See Samuel Scheffler, *The Rejection of Consequentialism: A Philosophical Investigation of the Considerations Underlying Rival Moral Conceptions* (Oxford: Oxford University Press, 1994).

<sup>33</sup>See Bernard Williams, “A critique of utilitarianism” in J. J. C. Smart and Bernard Williams, *Utilitarianism: For and Against* (Cambridge: Cambridge University Press, 1973). But also see Peter Railton, “Alienation, consequentialism, and the demands of morality” in James Rachels (ed.), *Ethical Theory* (Oxford: Oxford University Press, 1998). Williams famously objects that utilitarianism would

consequentialist to make sense of supererogation.

Third, an indirect, multidimensional consequentialism that acknowledges equality as a contributory value<sup>34</sup> is unlikely to ground rules that are applied unequally. For if people are to be treated equally as far as is practicable, and as far as is consistent with the other contributory values, then it seems that moral rules will have to apply to everyone unless there is some consideration that would justify a difference in their application. The independent contributory value of equality thus suggests that the best practicable rules would be imperatives that were universalizable in form.<sup>35</sup>

Fourth, if an indirect, multidimensional consequentialism acknowledges that the number of worthwhile lives is an independent contributory value, then that seems to imply that each life that is worthwhile is valuable. This also seems to suggest that each worthwhile life merits respect. But if so, the best practicable rules grounded by a multidimensional axiology that included the number of *worthwhile* lives within the plurality of independent contributory values that it recognized are unlikely to sanction the sacrifice of a worthwhile life for a trivial gain in the level of average utility. The moral rules that the multidimensional axiology here adumbrated would ground would be unlikely, therefore, to allow a person to be tortured to death in order that millions may avoid having to suffer

---

undermine our integrity as persons, while also charging Kantianism with the same defect. He also famously objects that the utilitarian has “one thought too many” in choosing to save his or her spouse, rather than another person, in accordance with utilitarian injunctions. See Bernard Williams, “Persons, character and morality” in *idem.*, *Moral Luck* (Cambridge: Cambridge University Press, 1981). However, Barbara Herman has provided Kantianism with a compelling defense against both charges. See Barbara Herman, “Integrity and impartiality” in *idem.*, *The Practice of Moral Judgment* (Cambridge, MA: Harvard University Press, 1993). Interestingly, Herman’s defense of Kantianism seems eminently co-optable by a rule consequentialist—especially a practicable rule consequentialist, who is more likely than an ideal rule consequentialist to construe an agent’s usual motive of duty as a limiting condition.

<sup>34</sup>Just as total utility can be argued to resolve into more than one contributory value, so can equality. See Alan Carter, “Value-pluralist egalitarianism,” *Journal of Philosophy* 99, 11 (2002): 577–99. It, too, can therefore be viewed as a second-order contributory value. (See note 28, above.) Moreover, as what makes a life worthwhile might involve multiple considerations, the number of worthwhile lives might also, conceivably, be viewed as a second-order contributory value.

<sup>35</sup>For the claim that moral rules are imperatives that possess the property of universalizability, see R. M. Hare, *The Language of Morals* (Oxford: Oxford University Press, 1952).

a mild headache, say. The moral rules grounded upon such an axiology are therefore more likely to be closer to commonsense morality than critics of consequentialism presume could ever be justified by a consequentialist approach.

Fifth, if an indirect, multidimensional consequentialism acknowledges that the level of average utility is also an independent contributory value, then, within the constraints set by paying due deference to the other contributory values (freedom, equality, and the number of worthwhile lives), the best practicable rules will be such as to help make people's lives, on average, happier. Much of what indirect utilitarians have cogently argued will thus be readily co-optable by an indirect, multidimensional consequentialism that recognizes the level of average utility as a contributory value. However, with the exception of disaster-avoidance clauses exclusively with respect to the greatest happiness of the greatest number, utilitarianism does not seem to explain all of the exceptions that reasonable moral rules might incorporate. But these exceptions do make sense from the standpoint of a rich, multidimensional axiology, for they serve to prevent the contributory values of freedom, equality and the number of worthwhile lives from being rode roughshod over by those who single-mindedly pursue the maximization of utility.

Sixth, when one is equipped with a multidimensional axiology that recognizes as contributory values freedom, equality, the number of worthwhile lives and the level of average utility, both the plausibility and limitations of other normative theories are readily apparent. Libertarianism perhaps has something going for it, for it recognizes the value of freedom. But it is deficient in failing to recognize as contributory values equality,<sup>36</sup> average utility and worthwhile lives. Kantianism has something going for it, too, for, depending upon the variety in question, it recognizes either the value of worthwhile lives or the value of freedom. But it is deficient in failing to recognize as contributory values equality and average utility. Pure egalitarianism<sup>37</sup> also has something going for it, for it recognizes the value of equality. But it is deficient in failing to recognize as contributory values freedom, average utility and worthwhile lives. Utilitarianism, too, has something going for it, for it

---

<sup>36</sup>Left libertarians value equality, but like all so-called libertarians, they are insufficiently critical of rights in private property. For a critique of justifications of such mooted rights, see Alan Carter, *The Philosophical Foundations of Property Rights* (Hemel Hempstead: Harvester Wheatsheaf, 1989).

<sup>37</sup>See Derek Parfit, *Equality or Priority?: The Lindley Lecture* (Kansas: University of Kansas, 1995).

recognizes the value of utility. But it is deficient in failing to recognize as contributory values freedom and equality. Prioritarianism<sup>38</sup> has something going for it, for it recognizes the greater value of helping the worst off. But it is deficient in failing to recognize as a contributory value equality, which it tends to dismiss because of the disvalue in an equal, but impoverished, world. Moreover, a multidimensional axiology that recognizes as contributory values both equality and the level of average utility will agree that there is disvalue in such a world; but the reason for that disvalue is that world's low level of average utility. And a multidimensional axiology that regards contributory values as variable values agrees that the worst off should be prioritized; but that is because greater overall value is added by improving greater rather than smaller deficiencies in a contributory value: a feature of diminishing marginal value. This also suggest that there might be something going for a sufficiency theory.<sup>39</sup> But sufficientarianism is deficient in failing to recognize equality as a contributory value. And some varieties of sufficiency theory may also be deficient to the extent that they are content with low levels of a contributory value when more overall value would be gained by increasing in degree any contributory value when that is not at the expense of another.<sup>40</sup>

Seventh, it should also be noted that an indirect, multidimensional consequentialism does not need to be confined to the promulgation of rules. There are dispositions that, were they cultivated, would help bring about the practicable world that contains the greatest overall value. And the cultivation of one disposition, amongst others, that a multidimensional axiology could be expected to ground is adherence to moral rules: those rules that are such that the greatest overall value would be advanced by their violators being liable to incur punishment, to incite in others blame and to arouse in themselves feelings of guilt. And just as there is reason for choosing the best practicable rules over

---

<sup>38</sup>See *ibid.*

<sup>39</sup>For one form of sufficiency theory, see Michael Slote, "Satisficing consequentialism," *Proceeding of the Aristotelian Society, Supplementary Volume* 58 (1984): 139–63, and for another, see Harry Frankfurt, "Equality as a moral ideal," *Ethics* 98 (1987): 21–43.

<sup>40</sup>Were there a limit, at least in principle, to the amount of moral value that the world could ever contain, then that might initially seem to add support to sufficientarianism. For a rejoinder, see Alan Carter, "A preemptive response to some possible objections to a multidimensional axiology with variable contributory values," forthcoming.

the best ideal ones, it might be argued that there is reason for expecting only the best practicable dispositions to be generally displayed. The best ideal dispositions—which perhaps only the strong-willed and highly motivated could manifest—might, however, provide the exemplars for the weak-willed to aim for, although, for them, the best practicable dispositions would suffice. A multidimensional axiology can, therefore, ground its own form of virtue ethics.

Finally, there are a number of other contributory values that would need to be recognized if an indirect, multidimensional consequentialism were to constitute an adequate normative theory in today's world. Thus far I have been arguing as if it were the case that only human beings are morally considerable. But there is reason to think that (at least some) nonhuman animals are morally considerable, too.<sup>41</sup> Moreover, a world that was devoid of natural environments would be a worse world than ours, even if all the humans living within such a world were extremely happy.

So, amalgamate all contributory values that pertain primarily to humans under the rubric “anthropocentric contributory value”,<sup>42</sup> and measure it along the *x*-axis of Figure 4. Amalgamate all contributory values that pertain primarily to morally considerable sentient nonhumans under the rubric “zoocentric contributory value”, and measure it along the *y*-axis. Finally, amalgamate all contributory values that pertain primarily to ecosystems and other non-sentient natural entities under the rubric “ecocentric contributory value”, and measure it along the *z*-axis. The world represented by point *T* is the world containing the greatest overall value. But it is not the world that exclusively anthropocentric, zoocentric or ecocentric normative theories would succeed in identifying

---

<sup>41</sup>See Peter Singer, *Animal Liberation: A New Ethics for Our Treatment of Animals*, 2nd ed. (London: Pimlico, 1995) and Tom Regan, *The Case for Animal Rights* (Berkeley: University of California Press, 1983). For a summary, see Alan Carter, “Animals” in John Skorupski (ed.), *The Routledge Companion to Ethics* (London: Routledge, forthcoming).

<sup>42</sup>The anthropocentric contributory value would resolve into second-order contributory values, such as the total utility enjoyed by humans. We might, therefore, think of it as a third-order contributory value. But just as treating second-order contributory values as if they were first-order, while often being expedient, could at times be highly problematic, treating third-order contributory values as if they were either second-order or first-order could be highly problematic, too, and for the same reason. Third-order contributory values should, therefore, also be treated with caution. (See note 28, above.)

as the best practicable world.

In other words, a multidimensional axiology can also ground an environmental ethic that is different, and less partial, than any proposed by anthropocentrists, zoocentrists or ecocentrists.<sup>43</sup> Such an environmental ethic would affect which moral rules are best promulgated. This means that the multidimensional axiology at its heart could well require some revision to our moral rules. It would also ground environmental virtues. Indeed, the multidimensional axiology at the heart of such an environmental ethic could well require considerable revision to what we currently regard as virtuous behavior. But it would also provide the yardstick for judging which public policies are best. And the multidimensional axiology at its heart could well require very different public policies or political activities<sup>44</sup> to those that are now commonly espoused.

---

<sup>43</sup>See Alan Carter, "Towards a multidimensional, environmentalist ethic" (forthcoming) and Alan Carter, "Inegalitarian biocentric consequentialism, the minimax implication, and multidimensional value theory: A brief proposal for a new direction in environmental ethics," *Utilitas* 17, 1 (2005): 62–84.

<sup>44</sup>For some indication of just how radical those activities might need to be, see Alan Carter, *A Radical Green Political Theory* (London and New York: Routledge, 1999).

**The Principle of Utility**

serves as the First Principle of


***The Art of Life***

(which deals in rules or precepts)


{ *Aesthetics*  
*Prudence or Policy*  
*Morality* (which is obligatory,  
for infractions of its rules are  
liable to incur punishment)

{ corresponds to a particular  
science—*ethology*  
(which provides the relevant  
empirical information)

**Table 1**



**Figure 1**



**Figure 2**




Figure 3





Figure 4



**Figure 5**

