https://doi.org/10.1017/S0265052509090128 Published online by Cambridge University Press

THE INTERPRETATION OF MAXIMIZING UTILITARIANISM*

By Jonathan Riley

I. Introduction

Pure maximizing utilitarianism prescribes that an agent ought to act in every situation so as to bring about a feasible outcome that maximizes general utility in the situation. General utility is a social norm that declares that each individual counts for exactly as much as another if each experiences an equal quantity of utility of the same kind or quality. As John Stuart Mill puts it: "Those conditions being supplied, [Jeremy] Bentham's dictum 'everybody to count for one, nobody for more than one', might be written under the principle of utility as an explanatory commentary."[1] This does not mean that an equal distribution of personal utilities is necessary to maximize general utility. Different individuals may experience different amounts and kinds of utility. It means, rather, that maximizing utilitarianism is perfectly impartial and unbiased. Given a conception of utility, utilitarianism counts equal quantities of utility of the same kind as equal in value, no matter who experiences the utility. Mill duly emphasizes the point in an important footnote: "[T]he principle of utility . . . suppos[es] that equal amounts of happiness are equally desirable, whether felt by the same or by different persons."[2]

Different conceptions of utility have been proposed. Classical hedonistic utilitarians such as Bentham and Mill conceive of utility as happiness in the sense of pleasure, including freedom from pain. The classical doctrine is usually associated with a crude way of determining general happiness which downplays the costs of acquiring and processing information about different persons' pleasures. The usual reading also ignores Mill's talk of different kinds of pleasures that differ in intrinsic quality or value.[3]

* For critical comments on earlier drafts, I am grateful to Molly Rothenberg, Ellen Frankel Paul, and the other contributors to this volume, especially Brad Hooker, Frances Kamm, Jan Narveson, and Alastair Norcross. I am also grateful to Harry Dolan for his careful copy-editing and suggested improvements in style. Responsibility for the views expressed remains mine alone.

[1] John Stuart Mill, *Utilitarianism,* in John M. Robson, ed., *Collected Works of John Stuart Mill* (hereafter *CW*), 33 vols. (London and Toronto: Routledge and University of Toronto Press, 1963–1991), X, 257.

[2] Ibid., 258.

[3] Mill's condition that "proper allowance" must be made for different kinds of utilities is typically ignored by utilitarians and their critics because of a received opinion that such

doi:10.1017/S0265052509090128

© 2008 Social Philosophy & Policy Foundation. Printed in the USA.

https://doi.org/10.1017/S0265052509090128 Published online by Cambridge University Press

General happiness is thus held to be simply a matter of adding up the known (expected) pleasant feelings of the members of a given population, taking for granted that pleasure is a single kind of feeling that has a common quality independently of its particular sources or objects. The object of the exercise is to select an outcome in which the sum total of these pleasures is maximized.[4]

Modern preference utilitarians such as Richard Hare and John Harsanyi reject hedonism as naive and instead conceive of utility in terms of rational preference-satisfaction. Their versions of utilitarianism are more sophisticated than the classical doctrine as usually interpreted, in that some allowance is made for limited human capacities to ascertain the consequences of particular acts and omissions for aggregate preference-satisfaction. But preference utilitarians continue to dismiss any suggestion of higher kinds of preferences that are intrinsically superior in value to inferior kinds. Moreover, they have different understandings of what rational preference-satisfaction means. Hare defines rational preferences as informed and prudent *desires* which are moral only if they also reflect the culturally-invariant logical properties of moral words such as the deontic 'ought' or 'must'. It emerges that moral desires are universal prescriptions that dictate how anyone in like circumstances ought to interact with his fellows to maximize aggregate desire-satisfaction. A rational agent forms such a moral desire by proposing a maxim or rule of action and continuing to affirm it as he imagines himself in the respective positions of every person in society with their own rational desires. Moreover, if he wishes to act in accordance with morality, then a rational agent's moral desires must generally override any of his competing desires, although Hare admits that it remains logically possible for a rational agent to opt out of being moral.[5]

---

qualitative distinctions are logically incompatible with utilitarianism. I shall argue later in the essay that this received view is mistaken. Nonetheless, it is worth noting at this point that the received opinion is not restricted to hedonistic conceptions of utility. Its status as conventional wisdom may be due in large part to its immediate acceptance by influential utilitarian contemporaries of Mill's such as Alexander Bain, Henry Sidgwick, W. Stanley Jevons, and Francis Y. Edgeworth, who (like Mill) also happened to be ethical hedonists. But the view that Millian qualitative distinctions are illegitimate has persisted even as hedonistic conceptions of utility have been replaced by nonhedonistic conceptions in recent sophisticated versions of utilitarianism such as those proposed by Richard Hare and John Harsanyi. Anti-utilitarians have not had much incentive to take Mill's side or challenge a view that has always been endorsed by leading utilitarian philosophers as well as by neoclassical economists.

[4] For a given population, maximizing the sum total is equivalent to maximizing the average pleasure per capita. But these two criteria diverge and neither is very appealing by itself if population is allowed to vary. The problems that arise under variable population must be faced by any plausible version of consequentialism, but I shall not address them here. For further discussion, see Charles Blackorby, Walter Bossert, and David Donaldson, *Population Issues in Social Choice Theory, Welfare Economics, and Ethics* (Cambridge: Cambridge University Press, 2005).

[5] Richard M. Hare, *Moral Thinking: Its Levels, Method, and Point* (Oxford: Clarendon Press, 1981); Hare, *Sorting Out Ethics* (Oxford: Clarendon Press, 1997); and Hare, *Objective Prescriptions and Other Essays* (Oxford: Clarendon Press, 1999).

Harsanyi equates rational preferences with informed and consistent *choices* independently of the desires or any other motivations underlying the choice behavior. Rational behavior must also satisfy certain axioms of morality to become moral. If a rational agent decides to behave morally, then he is required to make impersonal choices—and, in particular, to choose an optimal moral code to govern how agents are permitted to interact in any situation—by placing himself behind a veil of ignorance and imagining himself with equal probability in the positions of every person in society with their own choices while in their circumstances. While inhabiting any person's position, such an impersonal observer must count only self-oriented choices, that is, choices that depend solely on that person's own circumstances, including his property, job, lifestyle, and so forth. The observer must ignore what Ronald Dworkin calls "external" preferences, that is, other-oriented choices that affect others' circumstances and thereby augment or degrade others' self-oriented choices.[6] This prevents malevolent as well as benevolent other-oriented choices from being counted by the impersonal observer, which Harsanyi (agreeing with Dworkin) claims is necessary to ensure that each person's preference-satisfaction counts for exactly as much as another's in the determination of fair and impartial choices.[7] A moral agent must also have a disposition to comply with the rules of the optimal code at all times unless the code is justifiably suspended to enable society to survive some form of disaster or emergency situation.[8]

I shall refer indiscriminately to hedonistic and preference-based conceptions of utility, without denying that utility conceptions differ and that the differences may be important for some purposes. This strategy allows

[6] Ronald Dworkin, *Taking Rights Seriously* (London: Duckworth, 1977), 234.

[7] Dworkin rejects utilitarianism as unfair because he thinks that the self-oriented and external elements in any person's preference are sometimes "inextricably tied together" (ibid., 236). Prejudiced people may hold very intense preferences to avoid blacks or homosexuals, for example. In that case, "the personal preferences upon which a utilitarian argument must fix will be saturated with that prejudice," because it is impossible to isolate and ignore the underlying external preferences (ibid., 237). Harsanyi admits that the self-oriented and external elements may be inseparable, but he nevertheless suggests that utilitarianism can be suitably reconstituted so as to ignore unfair choices. His view seems to be that some personal choices can be classified as morally impermissible on the basis of their content and that a rational utilitarianism can be consistently required not to count such impermissible choices, including choices that exhibit undue bias for or against others' interests. See John C. Harsanyi, "Game and Decision Theoretic Models in Ethics," in Robert J. Aumann and Sergiu Hart, eds., *Handbook of Game Theory,* vol. 1 (Amsterdam: North Holland, 1992), 703–6. Hare is reluctant to ignore external preferences as he interprets them, although he struggles with the issue. See, for instance, Hare, *Objective Prescriptions,* 126–31, 158.

[8] John C. Harsanyi, "Morality and the Theory of Rational Behavior," *Social Research* 44 (1977): 623–56; reprinted in Amartya K. Sen and Bernard A. O. Williams, eds., *Utilitarianism and Beyond* (Cambridge: Cambridge University Press, 1982), 39–62. See also Harsanyi, "Game and Decision Theoretic Models in Ethics," 669–707. Note that Hare's universalizability criterion is a more demanding condition for moral choice than Harsanyi's equiprobability criterion, because Hare assumes no veil of ignorance.

https://doi.org/10.1017/S0265052509090128 Published online by Cambridge University Press

https://doi.org/10.1017/S0265052509090128 Published online by Cambridge University Press

me to focus attention on two conditions which are commonly presumed to be essential to the very structure of maximizing utilitarianism, however utility is conceived. One condition concerns the richness of the utility information: more specifically, utility must be cardinally measurable and interpersonally comparable. The other condition relates to the intrinsic nature of utility: to wit, utility must be homogeneous in quality so that utility differences are always quantitative. Together, these conditions ensure that different utilities can invariably be compared in terms of a single scale of rational numbers, no matter how many persons are considered at once and no matter what the sources or objects of their utilities.

But both of these presumed structural requirements are contestable and should be rejected by utilitarians themselves. Remarkably, both *were* rejected by Mill. My main goal is to explain how he can consistently do so and remain a maximizing utilitarian. His version of utilitarianism is of considerable interest, not least because it is quite distinct from any of the crude or sophisticated doctrines which are often passed off as maximizing utilitarianism.

II. The Two Presumed Structural Conditions

To avoid confusion and appreciate what Mill rejects about the usual picture of maximizing utilitarianism, the two conditions of (respectively) rich utility information and homogeneous intrinsic utility value, which are typically built into the theory's very structure, must be clarified. With respect to the rich utility information which is presumed necessary, Hare reflects conventional wisdom when he says that he needs to discuss "the interpersonal comparison of *degrees or strengths of preference,* because that is the kind of comparison I need for my own argument."[9] Yet he is not as clear as he might be that cardinal measurability and interpersonal comparability are separate requirements. A method of cardinalization is a way of constructing meaningful measures of gains and losses of personal utility. A method of interpersonal comparison is a way of giving meaning to one person's utility in relation to another's. Moreover, it is one thing to state these separate informational requirements, and quite another to specify workable methods of acquiring the requisite utility information.

Hare never specifies a method of cardinalization. Given his interpretation of utility in terms of desire-satisfaction, he seems to suppose that a rational agent just knows the relative intensities of his own desires, which could be reported to others: the agent "must know," for instance, if he is suffering "to a certain degree or with a certain intensity."[10] But no method is specified which anyone can, in principle, use to verify the individual's intensities of desire or aversion. Hare's failure in this regard is typical of

[9] Hare, *Moral Thinking,* 117; emphasis in the original.
[10] Ibid., 93.

utilitarian philosophers. Harsanyi is exceptional because he does specify an operational method of cardinalization in line with his behaviorist interpretation of preference as choice behavior. He defends the Von Neumann–Morgenstern method of cardinalization, whereby a person's attitudes toward risk are used to measure the relative preference intensities revealed by his consistent choice behavior under uncertainty, as represented by a Von Neumann–Morgenstern expected utility function.[11] This is a plausible method, assuming that a rational agent behaves in accord with the expected utility axioms, and that he possesses attitudes toward risk which he employs solely to assess the relative importance of outcomes. Even if those assumptions are accepted, however, the method faces serious problems in practice.

The costs of implementing the method seem prohibitive in the context of moral and political conflicts involving large numbers of people. Gathering the cardinal utility information may not be economically feasible even if all of these people are rational in Harsanyi's sense. Moreover, some people may have motivations that are incompatible with rationality thus construed. Given such motivations, an agent's observed choice behavior under risk need not reveal anything about the intensities of his preferences for the possible outcomes. Perhaps he merely loves to gamble, for instance, in which case he may choose as he does in order to get that excitement, as opposed to choosing in order to attain one outcome rather than another. True, Harsanyi can reply that rational agents do not love gambling or risk per se. But the difficulty is that some people are not choosing rationally in that sense, although this may not be obvious from their behavior in the situation at hand. As a result, costly steps must be taken to distinguish between irrational gamblers and expected utility maximizers whose attitudes toward risk are employed solely to assess outcomes and are revealed by their observed behavior.

Even if the Von Neumann–Morgenstern method could be implemented at a reasonable cost to give reliable estimates of people's preference intensities, there remains a normative worry: namely, why should preference intensity matter for resolving disagreements, especially if preferences are not distinguished by kind or quality? No doubt it matters that an individual consistently ranks the outcomes one way rather than another, but why should an impartial decision procedure take account of the individual's relative intensities of preference? Intensity is not necessarily correlated with the individual's genuine needs or vital interests, let alone with others' legitimate concerns. This normative worry is tied up with the problem of making meaningful interpersonal comparisons of utility. Why

[11] See, e.g., Harsanyi, "Normative Validity and Meaning of Von Neumann and Morgenstern Utilities," in Brian Skyrms, ed., *Studies in Logic and the Foundations of Game Theory* (Dordrecht: Kluwer, 1992). For further discussion of Harsanyi's rational choice version of utilitarianism with examples of how the method of cardinalization works, see Ken Binmore, *Playing Fair* (Cambridge, MA: MIT Press, 1994), 259–315.

https://doi.org/10.1017/S0265052509090128 Published online by Cambridge University Press

https://doi.org/10.1017/S0265052509090128 Published online by Cambridge University Press

should it necessarily matter if a person intensely prefers this outcome instead of that, when other people are indifferent between the two or even mildly in favor of the outcome which he dislikes? Even if everyone can appreciate that the person's zealous preference is more intense than everyone else's preferences put together, for instance, why should we conclude that the only reasonable interpersonal comparison is one that gives overwhelming weight to him as compared to other people? After all, he might be a selfish expected-utility-maximizing fanatic.

As for a method of interpersonal comparison, Hare and Harsanyi both accept variants of the dominant approach in the literature, to wit, the method of extended sympathy whereby one person imagines himself in another's position, adopts that other person's utilities (pleasures, desire-satisfactions, or choices, as the case may be) as his own while in the other's position, and then compares those utilities with his own utilities at his own position. This is a valid description of how any rational agent might go about making interpersonal comparisons. But it does not even begin to address epistemological worries about the reliability of sympathy as a way to acquire knowledge of the other's utilities so that meaningful comparisons can be made. No explanation is provided of how sympathy works to enable one person to reasonably gain access to another's genuine feelings or desires, for instance. Nor is anything said about how one person learns what it means for another to choose as he does—merely imitating another's choice behavior without understanding why he makes those particular choices cannot suffice for normative purposes.

In any case, to remove the normative problem of disagreement about which particular interpersonal comparisons are reasonable, it must also be assumed that rational and impartial individuals will ultimately make identical interpersonal comparisons. This is an especially strong requirement. Nevertheless, utilitarians typically accept it, and Hare and Harsanyi are no exceptions. Hare insists that his perfectly rational moral thinkers or "archangels, at the end of their critical thinking, will all say the same thing, on all questions on which moral argument is possible; and so shall we, to the extent that we manage to think like archangels."[12] Harsanyi imposes a fundamental similarity postulate which ensures that his impersonal observers will also make identical interpersonal comparisons and moral choices. Both thinkers seem merely to assert that there is ultimately only one reasonable answer on these questions in any given social context, and that rational agents will all converge on the answer if they become fully informed about each person's position, sympathize with the other by imagining themselves with his utilities while in his position, and deliberate impartially about the best course of action for all.

[12] Hare, *Moral Thinking*, 46.

Utilitarians may pay lip service to the difficulties surrounding workable methods of cardinalization and interpersonal comparison of utility, but they typically maintain that the difficulties can be overcome, at least in principle. Ever since the late nineteenth century, they have followed Henry Sidgwick and early neoclassical economists such as Francis Y. Edgeworth, W. Stanley Jevons, and Alfred Marshall in their refusal to be disconcerted by the objection that meaningful interpersonally comparable cardinal utility information is impossible to obtain.[13] Their continual acknowledgment that their doctrine depends on meaningful interpersonally comparable cardinal utilities, however, has arguably created a widespread impression that maximizing utilitarianism is impossible in the absence of such rich utility information.[14] Yet even Bentham expressed doubts that such information could ever really be acquired.[15]

Besides the presumed need for interpersonally comparable cardinal utilities, conventional wisdom also has it that utility, whether interpreted in terms of pleasure or preference-satisfaction, cannot be taken to exhibit differences of intrinsic quality apart from the question of quantity. Again following the lead of Sidgwick and the neoclassical economists, even utilitarians have generally rejected Mill's claim that "[i]t is quite compatible with the principle of utility to recognize the fact, that some *kinds* of pleasure are more desirable and more valuable than others."[16] Most have accepted what Mill considered "absurd," namely, that "the estimation of

[13] Henry Sidgwick's *The Methods of Ethics* (London: Macmillan, 1874) came out when the so-called marginalist revolution was beginning to transform classical economic theory into its neoclassical successor. Among other things, the revolution introduced the use of mathematical calculus to provide more precise models of hedonistic utilitarianism, which was the dominant ethical doctrine within normative economics, as Sidgwick urged. Perhaps the most influential work in this regard was Francis Y. Edgeworth's *Mathematical Psychics: An Essay on the Application of Mathematics to the Moral Sciences* (London: Kegan Paul, 1881), which makes the assumption that the analyst can speak meaningfully of interpersonally comparable natural units of pleasure. For further discussion of the evolving interpretations of utilitarianism within economic theory, see Jonathan Riley, "Utilitarianism and Economic Theory," in Steven Durlauf and Lawrence Blume, eds., *New Palgrave Dictionary of Economics,* 2d ed., 4 vols. (London: Palgrave Macmillan, 2008).

[14] Modern economists have proposed a more flexible, ordinal version of utilitarianism that abandons cardinal measure but retains interpersonal comparisons of utility based on the extended sympathy method. Utility represents preference orderings revealed by consistent choice behavior, and it is assumed that different persons' levels of preference-satisfaction can be meaningfully compared through sympathy. No commitment is made to hedonism. See, e.g., Kenneth J. Arrow, "Some Ordinalist-Utilitarian Notes on Rawls's Theory of Justice," *Journal of Philosophy* 70 (1973): 245–63; Arrow, "Extended Sympathy and the Possibility of Social Choice," *American Economic Review* 67 (1977): 219–25; and Allan F. Gibbard, "Ordinal Utilitarianism," in George R. Feiwel, ed., *Arrow and the Foundations of the Theory of Economic Policy* (London: Macmillan, 1987), 135–53. But this ordinal variant is seen by some as a nonstandard version of the doctrine, which is perhaps better classified as something other than utilitarianism. See Amartya K. Sen, "Utilitarianism and Welfarism," *Journal of Philosophy* 76 (1979): 463–89.

[15] For evidence of Bentham's doubts about cardinal measurement and interpersonal comparison of utility in the sense of pleasure including freedom from pain, see Wesley C. Mitchell, "Bentham's Felicific Calculus," *Political Science Quarterly* 33 (1918): 161–83.

[16] Mill, *Utilitarianism,* 211; emphasis in the original.

https://doi.org/10.1017/S0265052509090128 Published online by Cambridge University Press

https://doi.org/10.1017/S0265052509090128 Published online by Cambridge University Press

pleasures [or preference-satisfactions] should . . . depend on quantity alone."[17] Indeed, most seem to have implicitly agreed with G. E. Moore that the charge of absurdity is better directed at Mill himself.[18]

The main objection pressed against the idea of plural kinds of utilities with different intrinsic qualities apart from the question of quantity seems to be that one or more intrinsic values other than utility must be smuggled in to serve as a basis for the qualitative distinctions. Sidgwick and Moore, for instance, both insisted that a coherent hedonism must be purely quantitative because values other than pleasure are needed to distinguish different qualities of pleasure. In their view, pleasure is a simple agreeable feeling that always displays a common quality, independently of its sources or objects. As Moore puts it, "'pleasant' must, if words are to have any meaning at all, denote some one quality common to all the things that are pleasant."[19] Many later thinkers, utilitarians and anti-utilitarians alike, have agreed, and have repeated the objection that a hedonist cannot coherently maintain that plural kinds of pleasant feeling exist such that higher kinds are intrinsically superior in value to lower kinds merely as pleasures, apart from the question of quantity.[20]

Even Amartya Sen appears to accept this sort of objection, generalized to cover other conceptions of utility. He seems to take for granted that utility, however conceived, must be homogeneous in quality. As a result, no version of utilitarianism can accommodate differences of intrinsic quality within its conception of utility or associate higher and lower kinds of utilities with different sources or features of possible outcomes. Rather, as an alternative to utilitarianism, Sen introduces the family of "utility-supported moralities," in which the goodness of an outcome still depends on utility but nonutility information is needed to

[17] Ibid.

[18] G. E. Moore, *Principia Ethica* (Cambridge: Cambridge University Press, 1959). Unlike Sidgwick and the economists, Moore was a sharp critic of ethical hedonism, although he seems to have remained a utilitarian for whom utility or good is "one definite thing, *absolutely indefinable,* some one thing that is the same in all the various degrees and in all the various kinds of it that there may be" (ibid., 13; emphasis added).

[19] Ibid., 78.

[20] See, e.g., Simon Blackburn, *Being Good* (Oxford: Oxford University Press, 2001), 82. John Dewey is another prominent figure in this tradition of criticizing Millian qualitative superiorities as incoherent. Indeed, like Moore and transcendental idealists such as Kant, Hegel, Thomas Hill Green, and Bernard Bosanquet, he rejects hedonism altogether. Nevertheless, he seems to endorse a nonstandard pragmatist version of rational preference utilitarianism. In an ideal Deweyan society, rational individuals would habitually resolve practical problems by cooperating in terms of a shared but evolving idea of the general welfare in which great importance is assigned to a code of justice that distributes equal rights and correlative duties. By cooperating successfully on these terms, reflective individuals would jointly satisfy their desires ("ends-in-view") in harmony as they resolved an indefinite series of problematical situations arising over time in their natural environment. Among his many works, see John Dewey and James Hayden Tufts, *Ethics,* in Jo Ann Boydston, ed., *John Dewey: The Later Works: 1925–1953,* vol. 7 (Carbondale: Southern Illinois University Press, 1985), 191–99, 240–52.

distinguish among different kinds of utility and assign them distinct intrinsic values.[21]

Yet the claim that utility may exhibit intrinsic differences of quality apart from the issue of quantity deserves far more attention than it has received in the literature. As Rem Edwards has argued, for example, those who insist that pleasure must be homogeneous in quality seem merely to beg the question against Mill, because they never investigate whether hedonism might consistently accommodate Millian qualitative superiorities.[22] Perhaps the term 'pleasure' may denote a family of enjoyable feelings some of which are intrinsically more valuable than others as pleasures. Or, if hedonism is rejected, perhaps 'rational preference satisfaction' may cover a family of reasonable satisfactions some of which are intrinsically more valuable than others as satisfactions. Moreover, perhaps a persuasive case can be made that some nonstandard version of maximizing utilitarianism that accommodates such qualitative superiorities has more normative appeal than any "utility-supported morality."

But before examining what utilitarianism might look like when the two presumed structural conditions are dropped as unnecessary, it is useful to remind ourselves of how maximizing utilitarianism is commonly pictured. The common picture has its status because rich utility information of a homogeneous quality is presumed to be part of the very meaning of utilitarianism. A crude version of the common picture is perhaps more familiar to critics of utilitarianism. Nonetheless, even sophisticated utilitarians such as Hare and Harsanyi do not convincingly escape from the common picture altogether.

III. Crude Utilitarianism

Crude versions of maximizing utilitarianism assume that the requisite rich and homogeneous utility information is somehow available, and that rational and impartial individuals can always make use of this information at negligible cost to directly calculate acts or omissions which they can perform to bring about outcomes that maximize general happiness. This leads immediately to a dismaying conception of morality. Given that utility is the same kind of thing whatever its sources or objects, so that utility differences are purely quantitative in nature, and given that every person experiences utility, it follows that any rational and impartial individual must be willing to give up his own happiness for the general good.

[21] Amartya K. Sen, "Plural Utility," *Proceedings of the Aristotelian Society* 81 (1980): 193–215.

[22] Rem B. Edwards, *Pleasures and Pains: A Theory of Qualitative Hedonism* (Ithaca, NY: Cornell University Press, 1979). As will become obvious, I do not entirely agree with his view of the matter. But his discussion provides an excellent corrective to the conventional prejudice that utility (in this case, hedonistic utility) must be seen as homogeneous in quality independently of its sources or objects.

https://doi.org/10.1017/S0265052509090128 Published online by Cambridge University Press

https://doi.org/10.1017/S0265052509090128 Published online by Cambridge University Press

More specifically, the individual must be prepared to recognize myriad moral obligations to sacrifice his vital interests, even his life, to promote the happiness of other people. However large the quantity of utility which he may experience at some outcome, his utility can often be outweighed by a larger quantity of the same kind of thing experienced by some number of others in combination at a different outcome. His pleasure or preference-satisfaction, including freedom from pain or dissatisfaction, can be overridden by the combined pleasures or satisfactions of enough other people on so many occasions that he is obliged, as a matter of course, to make sacrifices, including extreme sacrifices, to bring about outcomes that maximize general utility.

It is child's play to invent examples that illustrate the crude utilitarian duty to ignore personal well-being to an alarming extent for the purpose of promoting others' happiness. In this connection, any human is plainly capable of experiencing only a finite amount of pleasure or preference-satisfaction. No individual or group of individuals can really experience an actual infinity of pleasure (including freedom from pain).[23] Even the satisfaction of being saved from death must be assumed to be of some finite magnitude. Any individual's rescue from death frees him from experiencing a feeling of pain which may be very large but not actually infinite. Thus, he may be obliged to give up his life to promote the greater happiness of others. Given that utilities are of the same quality whatever their sources or objects, the obligation may be triggered because he would thereby allow, say, a billion other people to go without suffering a bad headache. Moreover, if the situation is repeated, many individuals may be obliged to die to keep the rest free of headaches.[24]

Critics of utilitarianism are fond of such examples, and there is little doubt that the critics have a strong case if utilitarianism is genuinely committed to moral duties of this sort. Even so, some of the examples posed by critics are not nearly as persuasive as others. Indeed, some examples are properly viewed as hard cases for any moral theory, including any nonutilitarian theory which the critics may have in mind. The hard cases include situations in which an individual can save the lives of many other people only by committing suicide and/or killing a few innocent victims. Unless he sacrifices his own life and/or the lives of the few, the many will die because, by assumption, there is no other possible

[23] The term 'infinite' does not connote a real number or quantity. Rather, it properly refers to a hypothetical procedure of endless or unlimited expansion. True, the procedure might converge on a finite limit, but the procedure itself is endless and, strictly speaking, never actually reaches the limit. This is not incompatible with the invention of ideal objects such as 'infinite sets' for mathematical purposes, keeping in mind that nobody could ever observe the complete contents of an infinite set and, indeed, nobody can even conceive what such a completed set would look like in practice. For relevant discussion, see John Stuart Mill, *An Examination of Sir William Hamilton's Philosophy,* in Mill, *CW,* IX, 37–49, 82–86, 425–29.

[24] See, e.g., Alastair Norcross, "Comparing Harms: Headaches and Human Lives," *Philosophy and Public Affairs* 26 (1997): 135–67.

https://doi.org/10.1017/S0265052509090128 Published online by Cambridge University Press

way to rescue them. In such cases, the claim that the individual has a moral obligation to sacrifice his own life and/or the lives of the few does not appear implausible and may be made to seem compelling by making the number of people who would be saved sufficiently large. Consider the following instances suggested by hard cases discussed in the literature.[25]

Suppose that an extremely large individual can save five innocent bystanders from certain death by throwing himself in front of a runaway trolley car. Unless he kills himself by blocking the tracks with his massive body, the five will be crushed to death. Assume for convenience that these six people are very similar, so that each one apparently has the same amount of lifetime happiness at stake and each one's death would have the same utility consequences for other people. Crude utilitarianism says that the large man has a moral duty to sacrifice his life, given that the only possible outcomes involve his death or theirs: the combined happiness of the five seems to be about five times as large as his similar kind of happiness, and their deaths would apparently cause about five times as much dissatisfaction to others as his death would. Yet this crude utilitarian conclusion might not be unreasonable in this case. If there is any doubt, then assume that a hundred or a thousand others could be saved by the large man's act of sacrifice. The duty may remain even if it is impossible to be certain beforehand that the rescue attempt will succeed: some probability of success may be sufficient.

Similarly, perhaps a person has a moral duty to shoot an innocent man if that is the only way to prevent the deaths of nineteen others at the hands of an unjust kidnapper. A sheriff might also have a duty to hang an innocent stranger if there is no other way to stave off widespread violence and mayhem by unruly mobs who are upset over an unsolved crime in their community. And perhaps an outstanding military leader has a moral duty to allow himself to be taken captive and executed by a powerful foreign state whose leaders fear his prowess and threaten to annihilate his country and fellow citizens unless he gives himself up, as the Roman senate apparently demanded of Hannibal and his fellow Carthaginians a few years after the battle of Zama that ended the Second Punic War. Anti-utilitarians may well be prepared to endorse, however reluctantly, the crude utilitarian conclusions in these cases.

Nevertheless, the objections to crude utilitarianism are decisive. Its fatal weakness is shown clearly by the examples in which an individual has a duty to sacrifice his vital interests merely to promote some com-

[25] See, among many others, Philippa Foot, "The Problem of Abortion and the Doctrine of Double Effect," *Oxford Review* 5 (1967): 5–15; Bernard Williams, "A Critique of Utilitarianism," in J. J. C. Smart and Bernard Williams, *Utilitarianism: For and Against* (Cambridge: Cambridge University Press, 1973), 77–150; Judith Jarvis Thomson, "Killing, Letting Die, and the Trolley Problem," *The Monist* 59 (1976): 204–17; and Frances Kamm, *Intricate Ethics* (Oxford: Oxford University Press, 2007). Mill discusses the case of Hannibal in a letter to William Thomas Thornton dated April 17, 1863, in Mill, *CW,* XV, 853–54. See also Mill's letter to E. W. Young dated November 10, 1867, in *CW,* XVI, 1327–28.

paratively trivial concerns of a sufficient number of other people, such as saving them from suffering bad headaches or upset stomachs. Thus, Samuel Scheffler seems to be on solid ground when he objects to "the stringent normative principle that constitutes the heart of [utilitarianism]: namely, that the right act in any given situation is the one that produces the best overall outcome, where outcomes are ranked from best to worst by reference to a principle that gives equal weight to the interests of everyone."[26] Crude utilitarianism *does* yield a highly unattractive conception of morality.

In Scheffler's terms, crude utilitarian morality is *pervasive, overriding,* and *stringent*. It is pervasive in the sense that its scope is unlimited, so that every act or omission is subject to moral assessment, from which it seems to follow that any conduct is morally blameworthy or deserving of punishment to the extent that the conduct fails to produce an outcome that maximizes general utility.[27] It is overriding in the sense that it claims that moral considerations override all competing considerations. By implication, any rational agent must accept that, all things considered, his own interests are promoted by his doing the right things: prudence coincides with the satisfaction of moral duty.[28] And crude utilitarian morality is stringent because it places extremely demanding constraints on conduct. Indeed, the constraints are far too demanding for humans. An individual cannot reasonably be expected to satisfy a moral obligation to give up his own life even if his death is the only way to keep many others from suffering headaches. Perhaps heroes or saints could satisfy such an obligation, although it is far from obvious that they should. There does not appear to be anything heroic or saintly in the sacrifice even though it supposedly promotes the general happiness.

But perhaps the objections to utilitarianism lose force if a more sophisticated version of the doctrine replaces the crude version. How have

[26] Samuel Scheffler, *Human Morality* (Oxford: Oxford University Press, 1992), 103. See also Peter Vallentyne, "Against Maximizing Act-Consequentialism," in James Dreier, ed., *Contemporary Debates in Moral Theory* (Oxford: Basil Blackwell, 2005), 21–37.

[27] Scheffler apparently does not identify morality with deservingness of punishment for failure to do what we think is right. But I shall adopt Mill's view in *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" (184).

[28] Scheffler notes that his notion of overridingness is stronger than Hare's linguistic thesis that an individual who competently uses moral language to make *sincere* moral judgments treats moral principles as overriding (Scheffler, *Human Morality,* 56 n. 4). Hare allows that a rational agent may know how to use moral language yet opt out of morality, whereas the claim of overridingness as Scheffler understands it implies that a rational agent must, in his own interest, give absolute priority to morality over competing considerations. Although Scheffler largely ignores the point, a society may take various steps to promote harmony between prudence and morality by threatening to duly punish by law or opinion any individual who fails to meet his moral obligations. Society need not take for granted that all of its members are rational and impartial agents who are always willing to do what is right even in the absence of external incentives.

https://doi.org/10.1017/S0265052509090128 Published online by Cambridge University Press

sophisticated utilitarians such as Hare and Harsanyi attempted to blunt the force of the objections, and how persuasive are their arguments?[29]

IV. Sophisticated Utilitarianisms

Hare and Harsanyi both explicitly reject crude act-utilitarianism. Both may be understood to endorse maximizing utilitarianism as a criterion of right action but to reject the crude way of directly calculating the particular acts that have the best effects in terms of general utility. Both argue, in effect, that maximizing utilitarianism itself prescribes the replacement of the crude method of direct calculation by a more sophisticated process of deliberation for selecting right acts. A relatively sophisticated deliberative process involving a reliance on rules is said to be optimal because rules are useful guides to practical decision-making for most if not all humans, and general compliance with them promises to bring greater amounts of preference-satisfaction for all concerned. Even if a sophisticated process of deliberation has practical utility, however, it is difficult to see how that contention alone can allow maximizing utilitarians to deal convincingly with the serious criticisms leveled against the structure of their theory. The problems seem to lie ultimately with the utilitarian criterion and its implications for morality.[30]

Consider Hare's sophisticated utilitarian theory, for instance, which explicitly seeks to capture the virtues of rule-utilitarianism on an "intuitive level" and unite them with a "careful act-utilitarianism" on a distinct "critical level" of moral thinking. In his view, humans should typically forgo the difficult task of directly calculating right acts and instead learn the "relatively simple reaction-patterns" prescribed by common-sense morality for dealing with new contingencies that resemble in their main features contingencies already faced. The common moral intuitions, including principles of justice and rights, are *prima facie* moral principles because they are not always binding and may be overridden. But they arguably coincide with the prescriptions of maximizing utilitarianism in ordinary types of situations that people are likely to encounter repeatedly. Given the body of prima facie principles, direct calculation of right acts becomes necessary only if these principles come into conflict or if extraordinary situations emerge which are unlike any faced previously and thus are not covered by the common principles.

[29] Scheffler himself defends a conception of morality as pervasive but moderate and perhaps not overriding. I shall come back to his view of morality briefly in Section VII.

[30] Critics of utilitarianism recognize that the utilitarian account of right action can be distinguished from the utilitarian account of optimal deliberation, but this does not lead them to endorse utilitarian ethics. See, e.g., Scheffler, *Human Morality,* 38–39. For criticisms of both Hare's and Harsanyi's versions of utilitarianism, see Amartya Sen, *Collective Choice and Social Welfare* (San Francisco, CA: Freeman, 1970), 131–46; Bernard Williams, *Ethics and the Limits of Philosophy* (Cambridge, MA: Harvard University Press, 1985), 71–119; and D. Seanor and N. Fotion, eds., *Hare and Critics* (Oxford: Oxford University Press, 1988).

https://doi.org/10.1017/S0265052509090128 Published online by Cambridge University Press

https://doi.org/10.1017/S0265052509090128 Published online by Cambridge University Press

Nevertheless, critical moral thinking is needed to endorse any prima facie principles in the first place, to adjust them in order to remove conflicts, and to supplement them with additional principles to cover extraordinary situations. According to Hare, critical thinking involves a purely formal process of moral reasoning in which the thinker is committed by the logical properties of moral words to universalize his judgments about how to act in a given situation, all things considered. Such a thinker must follow Kant, Hare says, and express his judgment in the form of a universal rule that prescribes to anyone what he must do when interacting with other people in situations of the relevant sort. However, if it is to yield any competent concrete moral judgments, this purely formal reasoning process must be combined with actual personal desires informed by factual information about the natural environment. As Hare understands it, the critical thinker must imagine himself in the position of every individual with that person's preferences over the possible outcomes, compare the relative intensities of the personal preferences, and choose the best maxim of action for everyone to follow as that which brings about an outcome that maximizes aggregate preference-satisfaction. The upshot is that critical thinking, despite its formal neo-Kantian aspect, settles on utilitarian outcomes: "the logical apparatus of universal prescriptivism . . . will lead us in critical thinking . . . to make judgements which are the same as a careful act-utilitarian would make."[31] Moreover, Hare emphasizes that "discrimination . . . between pleasures or preferences because of their content" has "no place" in critical thinking: "We are not allowed to bring in intuitions that such and such pleasures are more worthwhile [than others]."[32] Thus, aggregate preference-satisfaction is understood in the usual way, as the sum total of personal utilities that are homogeneous in nature.[33]

Strictly speaking, critical moral thinking as Hare conceives it depends on meaningful interpersonally comparable and cardinal utility information of homogeneous quality. He admits, however, that it is difficult if not impossible to acquire such rich utility information in practice. Even so, he assumes that a superhuman "archangel" could acquire it in principle and infallibly use it to directly calculate right acts: "To archangels, who can do [critical thinking] perfectly, I have ascribed superhuman powers and superhuman knowledge, and there is no reason why this should not include knowledge of other people's states of mind."[34] Still, humans are not archangels and can only roughly approximate what the archangel does: "The most that human beings can ask for, when they are trying to do the

[31] Hare, *Moral Thinking,* 42–43.
[32] Ibid., 145.
[33] Hare's claim that Kant might be read as a sophisticated utilitarian is no doubt controversial. Cf. Hare, *Sorting Out Ethics,* 147–65; and Williams, *Ethics and the Limits of Philosophy,* 54–70, 82–92, 189–96.
[34] Hare, *Moral Thinking,* 122.

best critical thinking they can, is some way of approximating, perhaps not at all fully, to the thought-processes of an archangel."[35] This raises doubts that utilitarians or anyone else can ever really calculate right acts in accordance with the utilitarian criterion or ever really know that common moral intuitions coincide with utilitarian prescriptions in ordinary situations. For the sake of argument, however, I shall assume with Hare that humans can acquire something like the requisite rich utility information and use it to do roughly what the archangel would do perfectly.

Hare argues that his two-level theory allows us to appreciate that critics of utilitarianism such as Bernard Williams are unconvincing because the examples they use to condemn utilitarian thinking are extraordinary cases to which our common moral intuitions do not apply. It is hardly surprising that those who stick with their intuitions are outraged by the suggestion that an innocent victim ought to be shot to save the others in the kidnapper case, for instance, or that the innocent stranger ought to be hanged by the sheriff to prevent widespread mob violence. Although Hare often asserts that these extraordinary cases are unlikely to ever actually occur, he also suggests that, if they do, critical thinking shows that the counterintuitive acts of murdering the innocent may be the right ones for maximizing general utility. He points out that even the critics appear to admit this possibility, however reluctantly: "even Williams seems at least to contemplate the possibility of its being right to shoot the innocent man to save the nineteen other innocent men."[36] Similarly, perhaps it would have been right for the Carthaginians to surrender Hannibal to the Romans to save Carthage from annihilation.[37] At the same time, utilitarians and their critics can agree that such right acts are so repugnant in terms of common moral intuitions that individuals may fail to do their utilitarian duties.

Hare's sophisticated theory seems merely to reiterate the crude utilitarian prescriptions in the hard cases listed earlier, even as it tries to account for the way in which such prescriptions clash with our common moral intuitions. Hare does attempt to allow for supererogation, however, by introducing the idea that different classes of prima facie moral principles may be tailored for different types of humans: rare noble types of people could be covered by principles that oblige them to sacrifice their lives for the greater good of others, for instance, whereas most of us would be covered by different principles that excuse us from such sacrifices and place them beyond the call of duty.[38] But such complexity may ruin the simplicity which is supposed to be a feature of intuitive moral principles, and the question arises as to how many distinct types of humans

[35] Ibid.

[36] Ibid., 49. See also Hare, *Objective Prescriptions,* 161–63.

[37] Hannibal apparently saved his fellow citizens from having to make a choice by fleeing into exile in Asia Minor.

[38] Hare, *Moral Thinking,* 198–205.

https://doi.org/10.1017/S0265052509090128 Published online by Cambridge University Press

https://doi.org/10.1017/S0265052509090128 Published online by Cambridge University Press

may be recognized. Even more importantly, it remains unclear why the archangel would not impose instead a universal moral duty to develop into a noble individual capable of the requisite self-sacrifice for the general good. Hare might retort that humans exhibit more or less *fixed* natures of the different types, but that seems dangerously close to an excuse for most to engage in immorality.

Indeed, it seems difficult for Hare to avoid endorsing the crude utilitarian prescription that an individual may have a moral duty to give up his life to save billions from suffering bad headaches. True, critical thinking may reject such a duty by supposing that an archangel would agree that any human's intensity of preference for his life is so strong in comparison to the preference intensities of everyone else put together to be relieved of their headaches that the desire for life always wins out. But such fanatically intense personal preferences appear implausible given that everyone's preference-satisfaction is homogeneous in quality, and Hare tends to dismiss the possibility that fanaticism of this sort will be observed in actual situations.[39] At the same time, as Bernard Williams points out, fanatical utilitarian judgments that prescribe the annihilation of some minority group such as Jews, although perhaps implausible, cannot be entirely ruled out by critical thinking as Hare conceives it. The utility sums still may not come out right. "Moreover, on the utilitarian argument it emerges as *a* consideration—though, if the sums come out right, not a decisive or winning consideration—that racists get some satisfaction out of the sufferings of the Jews; but this cannot be a consideration *at all*."[40]

I now turn briefly to Harsanyi's sophisticated rule-utilitarianism. Harsanyi argues that the utilitarian criterion should be restricted to the selection of optimal codes because the crude method of directly calculating right acts in every situation cannot take account of social benefits that are associated exclusively with codes of rules. By distributing various individual rights and duties, codes are able to create significant freedom of choice, stability of expectations, and incentives for productivity, which are important social benefits in light of our more or less fixed human nature that makes us biased in favor of ourselves, family, and close friends. Thus, maximizing utilitarians ought to endorse a two-stage process of deliberation. At the first stage, an impersonal observer directly calculates a code that maximizes aggregate expected preference-satisfaction, assuming that rational individuals would assess and reveal preferences for alternative codes in terms of the freedom and other benefits made pos-

[39] Ibid., 169–82.

[40] Williams, *Ethics and the Limits of Philosophy,* 86–87; emphasis in the original. As I have already indicated, Hare is reluctant to deal with this problem more directly by endorsing Harsanyi's move of excluding "external" preferences in Dworkin's sense. Utilitarians do seem unable to consistently make such a move, given that pleasures or satisfactions are homogeneous in quality.

sible by them. At the second stage, the observer indirectly calculates right acts by prescribing that any individual should do whatever best serves his own biased interests provided he complies with the rules of the code selected at the first stage.[41]

Even so, although he defines utilitarianism so that compliance with some such code always maximizes general utility, Harsanyi does not really provide any justification for this beyond his assertions that humans are biased by nature in favor of family and friends. Unless firm evidence can be presented that humans are incapable of significant self-development, however, such assertions are vulnerable to the objection that they merely attempt to excuse the individual from performing his genuine act-utilitarian duties. Harsanyi also sounds at times (perhaps unintentionally) as if he thinks that acts of deviation from an optimal code may temporarily increase general utility, even if compliance is invariably best in the long run. But maximizing utilitarianism cannot consistently proscribe such deviations, it seems, and must revise the rules of the code accordingly to allow for the exceptions. This raises the issue of how fine-grained the rules must be to admit all of the deviations that may increase the general good. Even Harsanyi accepts that it is logically conceivable that an optimal code might be extensionally equivalent to crude act-utilitarianism, although he finds this possibility highly implausible given our empirical knowledge of human nature.

Even supposing that an optimal rule-utilitarianism does not reduce to crude act-utilitarianism, we can only guess at the nature of the particular rules and rights that would be selected by any perfectly rational impersonal observer with the requisite interpersonally comparable and cardinal utility information of a homogeneous quality. In this regard, Harsanyi provides no justification for his substantive assumption that the ideal observer will refuse to count any gains or losses of utility that can be traced to benevolent or malevolent intentions toward others, even if such intentions could somehow be identified and distinguished from purely self-oriented ones. Without further argument, his move seems merely an attempt to smuggle in common moral intuitions as preordained conclusions of the utilitarian calculus. He also dismisses any suggestion that a set of basic rights should have absolute priority over any competing considerations.[42] He thus seems to admit that any right, including a right to life, may be overridden to promote the general good of others, possibly in a wide variety of situations. Moreover, as Peter Diamond first pointed out, an impersonal observer in Harsanyi's theory must be indifferent

[41] Harsanyi presents his version of rule-utilitarianism as a two-stage game in which participants mutually cooperate to choose an optimal code at the first stage and then compete with each other by playing their permissible strategies to maximize their own advantage at the second stage. See Harsanyi, "Game and Decision Theoretic Models in Ethics," 692–94.

[42] For further discussion of this point, see Jonathan Riley, "Rule Utilitarianism and Liberal Priorities," in M. Fleurbaey, M. Salles, and J. Weymark, eds., *Justice, Political Liberalism, and Utilitarianism* (Cambridge: Cambridge University Press, 2008).

https://doi.org/10.1017/S0265052509090128 Published online by Cambridge University Press

about whether an individual is given a "fair shake" in the process of calculating a utilitarian outcome: the individual may be denied any opportunity to earn a decent living, for instance, yet this does not affect the observer's judgment about which outcomes are best.[43] Indeed, Ken Binmore argues that Harsanyi's neo-Kantian approach, which supposes that an ideal observer can transcend any particular cultural context, lacks any conceivable basis for meaningful interpersonal comparisons of utility.[44]

It seems fair to conclude that the sophisticated utilitarianisms of Hare and Harsanyi are unable to convincingly escape from the serious objections presented by critics of the utilitarian criterion. Their commitment to rich utility information of a homogeneous quality ultimately ties utilitarians to intuitively unappealing prescriptions in various situations. Upon reflection, the hostile intuitions appear to have more normative appeal than the utilitarian prescriptions.[45] If utilitarianism is to be made a credible doctrine, therefore, a very different approach is called for.

V. Plural Ordinal Utilitarianism

Mill rejects both of the presumed structural conditions of maximizing utilitarianism. He does not assume the feasibility of cardinal and interpersonally comparable utilities but rather supposes that only purely ordinal information is available about any person's pleasures or satisfactions. At the same time, he argues that there are plural kinds of pleasant or satisfying feelings, some of which are incommensurably more valuable than others. This combination of pure ordinalism and qualitative pluralism gives rise to what I shall call a plural ordinal version of utilitarianism. Utilitarianism thus understood aggregates over plural kinds of individual preference or judgment orderings, where an ordering of a given kind ranks possible outcomes in terms of estimated quantities of the relevant kind of pleasure or satisfaction to be expected from them, and thereby

[43] Peter A. Diamond, "Cardinal Welfare, Individualistic Ethics, and Interpersonal Comparisons of Utility: Comment," *Journal of Political Economy* 75 (1967): 765–66.

[44] Ken Binmore, *Natural Justice* (Oxford: Oxford University Press, 2005). In Binmore's view, cultural evolution supplies standards of interpersonal comparison which participants in that culture typically believe are reasonable. For a critical discussion of Binmore's contractarian theory of justice, see Jonathan Riley, "Genes, Memes, and Justice," *Analyse & Kritik* 28 (2006): 32–56.

[45] The considered appeal of various common moral intuitions has led many moral thinkers, including those who assess acts and omissions in terms of their utility consequences for other persons or sentient creatures, to reject even sophisticated versions of maximizing utilitarianism. Brad Hooker, for instance, defends a version of rule-consequentialism that admits some reductions in the sum total when needed to achieve a fair distribution of utilities: extra weight is given to the utilities of the worst-off members of society so that fairness may be attained. In his view, such a rule-consequentialism provides a systematic impartial justification for shared moral convictions that have their own independent credibility. See Brad Hooker, *Ideal Code, Real World* (Oxford: Clarendon Press, 2000). One serious worry, however, among others, is that Hooker's consequentialism continues to rely, in principle, on rich cardinal and interpersonally comparable utility information.

https://doi.org/10.1017/S0265052509090128 Published online by Cambridge University Press

provides an ordinal measure of the quantities, that is, a rough measure that merely orders them in terms of more or less. The plural kinds of orderings are arranged into a hierarchy that reflects a rational judgment that higher kinds are superior in quality to lower kinds. The aggregation procedure starts with the highest kind at the top of the hierarchy, counts each person's ordering of that type for exactly as much as another's assuming that the people concerned are competent to form that kind of preference, and then proceeds down the hierarchy to similarly count the lower types of orderings. Ultimately, an overall judgment is produced about which outcomes ought to be achieved to maximize the general happiness, although general happiness is conceived so that there is no conflict between it and personal happiness. A best outcome is one in which each member of the community simultaneously maximizes his *personal* happiness in point of quantity and quality. As Mill puts it, "the ultimate end, with reference to and for the sake of which all other things are desirable (*whether we are considering our own good or that of other people*), is an existence exempt as far as possible from pain, and as rich as possible in enjoyments, both in point of quantity and quality."[46]

Plural ordinal utilitarianism is an extraordinary doctrine which is not easily understood. Moreover, as I have already indicated, the large claims made for it by Mill have provoked expressions of bewilderment and outright contempt from many scholars, even those who are otherwise sympathetic to utilitarian ethics. As a result, it is important to clarify what he apparently proposes to substitute for the presumed structural conditions. What does it really mean, for instance, to replace rich cardinal and interpersonally comparable utility information with relatively impoverished ordinal utility information? How can utilitarianism work with the relatively poor information about pleasure or satisfaction that is contained in interpersonally noncomparable preference orderings?

Mill occasionally remarks that different sensibilities, upbringings, or cultural backgrounds defeat any easy presumption that one person can appreciate either the intensity or the kind of another's actual feelings of pleasure and pain.[47] He implies that utilitarianism as he conceives it does

[46] Mill, *Utilitarianism,* 214; emphasis added.

[47] In his discussion of the "desire to be in unity" that he considers "a natural basis of sentiment for utilitarian morality," Mill admits that even in fairly advanced societies "a person cannot indeed feel that entireness of sympathy with all others, which would make any real discordance in the general direction of their conduct in life impossible" (*Utilitarianism,* 233). He says that "differences of opinion and of mental culture [may] make it impossible for him to share many of their *actual* feelings—perhaps make him denounce and defy those feelings" (ibid., emphasis added). Even so, he suggests that at least some individuals have some sympathy for "promoting" other people's "own good," that is, their personal happiness correctly perceived in point of quantity and quality, as opposed to what they may actually desire and take pleasure in. The implication seems to be that sympathy for others' actual feelings is not necessary for an individual to develop the desire for social unity. *A fortiori,* utilitarianism, as Mill conceives it, does not depend on interpersonal comparisons based upon sympathy.

https://doi.org/10.1017/S0265052509090128 Published online by Cambridge University Press

not require interpersonally comparable cardinal utilities. Rather, the final judgment as to how much of given kind of pleasure is to be expected in this situation or that must be left to the majority among those who are competently acquainted with the relevant feeling:

> What means are there of determining which is the acutest of two pains, or the intensest of two pleasurable sensations, except the general suffrage of those who are familiar with both? Neither pains nor pleasures are homogeneous, and pain is always heterogeneous with pleasure. What is there to decide whether a particular pleasure is worth purchasing at the cost of a particular pain, except the feelings and judgement of the experienced?[48]

Given his hedonistic conception of utility, Mill apparently supposes that the only feasible utility information is contained in fallible people's possibly mistaken estimates of the various amounts and kinds of pleasures and pains to be expected with respect to any given domain of possible outcomes, where it is taken for granted that the persons are at least minimally competent to estimate the quantities and kinds of pleasures and pains to be expected. The estimates are embodied in personal judgments or preferences defined over the relevant domain. In accordance with hedonism, any rational individual may be assumed to rank or order the possible outcomes to reflect what he considers the most enlightened estimates available of the happiness which can be expected from them. If individuals disagree, then the judgment of the majority among them must be accepted as final.[49]

[48] Mill, *Utilitarianism,* 213.

[49] On the link between utilitarianism and majority rule in the context of interpersonally noncomparable ordinal utility information, see Jonathan Riley, "Utilitarian Ethics and Democratic Government," *Ethics* 100 (1990): 335–48. Majority rule is subject to inconsistencies, unfortunately, as Arrow's impossibility theorem implies. See Kenneth J. Arrow, *Social Choice and Individual Values,* 2d ed. (New Haven, CT: Yale University Press, 1963). One way to remove the inconsistencies is to employ a positional rule or scoring function such as Borda count. Under Borda count, each person ranks the m feasible outcomes from best to worst, and $m-1$, $m-2$,…,1,0 points are assigned to the best, next-best, . . . , second-worst, and worst outcomes, respectively. The procedure then selects an outcome with the greatest aggregate point total. More generally, to escape from the Borda method's insistence on *equal* spacing of points awarded for first choice, second choice, and so forth, a generalized scoring rule could be employed whose point assignments are unique only up to a positive monotonic transformation so long as such transformations are applied to every person's ordering. Although it has the formal appearance of an interpersonally comparable ordinal utility function, this generalized scoring function is really just a way to consistently count one person's preference ordering for exactly as much as another's. There is no attempt to justify the scoring function in terms of interpersonal comparisons of actual pleasures or satisfactions. It does not reflect the results of the extended-sympathy method applied to different people's enjoyments as estimated in their respective preference orderings. Despite its possible attractions as a way of achieving the utilitarian aim of impartially counting everyone for one and only one, however, this scoring approach does occasionally fail to choose Condorcet winners when they exist; that is, it may fail to select an outcome that is majority-

https://doi.org/10.1017/S0265052509090128 Published online by Cambridge University Press

By itself, any such attempt to rely on the purely ordinal utility information contained in interpersonally noncomparable preferences may seem very unpromising. In fact, insurmountable problems appear to exist if we accept the usual assumption that pleasure or satisfaction is a single kind of feeling whatever its sources or objects. In that case, any rational individual may be assumed to form an ordering in which the top-ranked outcome is one that he estimates will give him more personal pleasure net of pain than any alternative. But utilitarianism requires impartial estimates of the general happiness, and it is not obvious that such estimates can be provided by a majority decision among individuals with conflicting estimates of the amounts of personal happiness which they may experience at different outcomes. Mill emphasizes that there are no ideal observers or archangels available to provide certain knowledge of the true amounts of happiness to be expected by anyone from this or that outcome.[50] Rather, to get an idea of the happiness expected by other members of the community in matters of joint concern, a rational individual must participate with his fellows in an open political procedure such as majority voting that impartially counts each person's preference ordering after free discussion and debate. Any outcome selected as best by the majority might be justified as the result of a procedure that impartially takes into account each person's ranking based on his estimates of his own happiness. Nevertheless, it can be objected that although the majority outcome may give most people their greatest feasible personal happiness, this is not necessarily the same thing as the general happiness. Obviously, the majority might achieve personal happiness by oppressing the minority. The problem is that a rational individual is not required to show any concern for others when estimating the personal happiness he expects from the possible outcomes.[51]

To remedy this problem, it might be assumed that any rational agent directly and impartially estimates the general happiness instead of his own happiness to be expected from possible outcomes. The individual would now form preference orderings in which outcomes are ranked in

preferred to every other outcome in a series of binary contests. In any case, I shall henceforth assume that such a generalized scoring procedure is used to remove any inconsistencies otherwise associated with majority voting. (For recent critical surveys of the vast literature relating to scoring rules, including Borda count, see Steven J. Brams and Peter C. Fishburn, "Voting Procedures," in K. J. Arrow, A. K. Sen, and K. Suzumura, eds., *Handbook of Social Choice and Welfare,* vol. 1 [Amsterdam: Elsevier, 2002], 173–236; and Prasanta K. Pattanaik, "Positional Rules of Collective Decision-Making," in ibid., 361–94.) Needless to say, Bentham and Mill never discuss such a scoring procedure, but it is arguably in the spirit of utilitarianism as they conceive it. For further discussion, see Jonathan Riley, "Classical Ordinal Utilitarianism," unpublished working paper.

[50] Mill, *On Liberty,* in *CW,* XVIII, 252–57. See also Jonathan Riley, *Mill on Liberty* (London: Routledge, 1998), 67–68.

[51] Bentham's awkward phrase "the greatest happiness of the greatest number" makes more sense if he is interpreted as referring to possibly distinct notions of personal happiness held by the majority rather than to some independent notion of collective happiness based on rich utility information.

https://doi.org/10.1017/S0265052509090128 Published online by Cambridge University Press

https://doi.org/10.1017/S0265052509090128 Published online by Cambridge University Press

line with his estimates of the total amount of pleasure net of pain which the members of the community are likely to experience at those outcomes. There would no longer be any need to rely on a political process such as majority voting to aggregate over different people's estimates of their own happiness. Nevertheless, even the most rational human is still a fallible creature whose estimates of general happiness may be badly mistaken. Although the estimates developed by one person may look quite precise, different people may disagree in their estimates so that the reliability of any individual's estimates is open to doubt. Since they cannot appeal to an infallible umpire to provide incontestable knowledge of the true quantities, people require some other way to resolve their disagreements about what ought to be done to maximize the general happiness in matters of joint concern. Mill implies that they must again rely on the majority's judgment "since there is no other tribunal to be referred to."[52] As is well known, majority voting may be justified under certain conditions as a way to obtain maximum-likelihood estimates of the true quantities of general happiness.[53] In effect, it now serves as an aggregation procedure in which everybody's personal estimates of general happiness, as contained in their respective moral preference orderings, count for one and only one.

Even so, this way of interpreting the purely ordinal approach has its own severe problems. After all, it apparently substitutes rough and vague personal estimates of general happiness for the bold calculations based on rich utility information which are so frequently taken to be distinctive of utilitarianism. It also seems to ignore any legitimate concern for personal happiness. Nothing is said to assuage the familiar worry that the individual may be called upon to sacrifice his own happiness to promote that of other people, for instance, or to explain why a rational individual would ever feel motivated to act in the impartial manner demanded by utilitarianism. All that seems to have happened is that the individual is now expected to neglect his vital personal interests in favor of some fallible majority's opinion of the general happiness.

On its own, the turn to interpersonally noncomparable preference orderings seems problematic whether the preferences are shaped by estimates of personal happiness or general happiness to be expected from the feasible outcomes. Remarkably, however, these problems disappear once due consideration is paid to Mill's argument that there are different kinds of pleasures, some of which are intrinsically more valuable than others apart from the issue of quantity. In his view, a qualitatively superior feeling of pleasure is so much more valuable than a qualitatively inferior

[52] Mill, *Utilitarianism,* 213.

[53] On the use of majoritarian voting procedures as epistemic devices for providing maximum-likelihood estimates of some objective good, such as aggregate net pleasure, see H. Peyton Young, "Condorcet's Theory of Voting," *American Political Science Review* 82 (1988): 1231–44.

one that people who are competently acquainted with both kinds would not give up the higher pleasure "for any quantity" of the lower "which their nature is capable of":

> If I am asked, what I mean by a difference of quality in pleasures, or what makes one pleasure more valuable than another, merely as a pleasure, except its being greater in amount, there is but one possible answer. . . . If one of the two is, by those who are competently acquainted with both, placed so far above the other that they prefer it, even though knowing it to be attended with a greater amount of discontent, and would not resign it for any quantity of the other pleasure which their nature is capable of, we are justified in ascribing to the preferred enjoyment a superiority in quality, so far outweighing quantity as to render it, in comparison, of small account.[54]

In short, Millian qualitative superiority is *infinite* superiority in the sense that even a bit of the higher pleasure remains more valuable than any quantity of the lower, no matter how large the finite mass of lower pleasure is assumed to become.[55] If the people competently acquainted with both kinds disagree about whether one kind is qualitatively superior to the other, then the majority's judgment must again be accepted as final:

> And there needs be the less hesitation to accept this judgement respecting the quality of pleasures, since there is no other tribunal to be referred to even on the question of quantity. . . . When, therefore, those feelings and judgement [of the competent majority] declare the pleasures derived from the higher faculties to be preferable *in kind,* apart from the question of intensity, to those of which the animal nature, disjoined from the higher faculties, is susceptible, they are entitled on this subject to the same regard.[56]

Although it may sound counterintuitive at first, Mill is insisting that utility is a complex multidimensional phenomenon whose higher dimensions are infinitely more valuable than the lower dimensions. In effect, after estimating the quantities and qualities of pleasures to be expected from different possible outcomes, any rational individual has plural preference orderings arranged in a hierarchy with the highest kind of pref-

[54] Mill, *Utilitarianism,* 211.

[55] For an extended discussion of this interpretation of Millian qualitative superiorities as infinite superiorities, see Jonathan Riley, "Millian Qualitative Superiorities and Utilitarianism, Parts 1 and 2," *Utilitas* 20 (2008). It is important to keep in mind that infinity is not a real number or quantity: there is no such thing as a completed infinity in nature. Thus, it cannot be maintained that a higher pleasure is equal in value to an actual infinite amount of lower pleasure.

[56] Mill, *Utilitarianism,* 213; emphasis in the original.

https://doi.org/10.1017/S0265052509090128 Published online by Cambridge University Press

erences at the top, followed by the next highest kind, and so forth down to the lowest. To the extent that people competently acquainted with the pleasures disagree, the majority's judgments with respect to both quantity and quality are decisive.

Mill argues that "the pleasures of the intellect, of the feelings and imagination, and of the moral sentiments" have "a much higher value as pleasures than . . . those of mere sensation."[57] More specifically, his argument seems to be that the mental feelings of pleasure associated with our intellectual capacities are infinitely more valuable as pleasures than the inchoate sensations of pleasure registered by our body, that is, our "animal nature" when assumed to be "disjoined" from our intellectual capacities of forming ideas and propositions, reasoning, imagining, remembering, and so forth. Moreover, within the mental category, the kind of pleasant feeling associated with the moral sentiment of justice is qualitatively superior to any competing kinds of pleasures, he suggests, although aesthetic pleasures that do not conflict with justice may be qualitatively supreme. It is worth emphasizing that he purports to derive this ranking of different kinds of pleasures from the "feelings and judgement" of the majority of those who are competently acquainted with the different kinds. Most if not all reasonable and impartial people will, in his view, endorse such a qualitative preference ordering, at least upon reflection.[58]

This qualitative preference ordering of the majority with competent experience of the different kinds of pleasant feelings is a very special sort of preference ordering, namely, a lexicographical (or lexical) ordering. In a lexical ordering of alternatives, a particular feature, or type of feature, of the alternatives has absolute priority over other features for determining the ordering. A dictionary serves as the model: the letter 'a' at the start of a word puts the word ahead of all words that begin with other letters, no matter how many other letters compose the words. Under hedonism, where the pleasure (including freedom from pain) expected from possible outcomes is the only feature of those outcomes that has intrinsic value, a lexical ordering arises if there are higher pleasures which have absolute priority over lower pleasures for determining the ordering. This is the case in Mill's pluralistic version of hedonism, where the infinite superiority of a higher pleasure implies that it has absolute priority over lower pleasures. The higher kind's infinite superiority over lower kinds produces discontinuities of intrinsic value—no finite amount of lower pleasure, however large, is equal in value to any amount of higher pleasure. As is well known, such discontinuities are characteristic of lexical orderings.

[57] Ibid., 211.

[58] This sketch of Mill's higher pleasures doctrine requires further clarification, but I cannot provide it here. For further discussion of the different kinds of pleasures and of why it may be reasonable to regard the higher kinds as infinitely superior in value to the lower kinds, see Jonathan Riley, "Millian Qualitative Superiorities and Rational Agency," unpublished working paper.

https://doi.org/10.1017/S0265052509090128 Published online by Cambridge University Press

The lexical ordering of kinds of pleasures must not be conflated with an ordinary ordering that can, in principle, be represented by a continuous real-valued mathematical function.[59] The lexical ordering is a very special ranking because it captures the discontinuities of value produced by the infinite superiority of higher pleasures over lower ones. In contrast, a simple ordinal ranking of pleasures is a purely quantitative measure: pleasures ranked higher in such an ordering are judged merely to be larger than pleasures ranked lower in the ordering. Given that the ordering can be represented by a continuous real-valued function, there is an implicit assumption that the alternatives being ranked differ in value only to some finite degree. Under hedonism, where pleasure is the sole intrinsic value, the implication is that the pleasures to be expected from different outcomes differ only to a finite extent. Thus, an ordinary ordinal ranking may be employed to rank different estimated amounts of pleasure of the same kind or quality. But such a ranking cannot do the work of a lexical ordering of different kinds of pleasure in Mill's sense.[60]

Strictly speaking, the competent majority's lexical ordering of the different qualities of pleasures is a lexical meta-ranking, or a ranking of rankings defined over the domain of possible outcomes. The lexical ranking is a qualitative ranking of plural kinds of quantitative rankings: for each kind, the majority competently acquainted with that kind of pleasure estimates the various quantities of it to be expected from the outcomes and ranks the outcomes accordingly in terms of more or less of that kind of pleasure.[61] The quantitative ranking of outcomes in terms of a

[59] In his excellent discussion of pluralistic or qualitative hedonism, Edwards suggests that Millian qualitative superiorities might be captured by a simple ordinal ranking of the different kinds or qualities of pleasant feelings (Edwards, *Pleasures and Pains,* 68–72, 111–19). Like Edwards, Fred Wilson, in *Psychological Analysis and the Philosophy of J. S. Mill* (Toronto: University of Toronto Press, 1990), also interprets Mill as groping for an ordinal utility scale that orders different qualities of utility from higher to lower without trying to quantify how much a higher quality differs from a lower one (ibid., 220–23, 253, 275–93).

[60] As a technical matter, the fact that a lexical preference ordering cannot be represented by a continuous real-valued mathematical function means that a utility function, in the purely formal sense of 'utility' used by modern economists, cannot be used to represent the information contained in the lexical preference. Thus, an unbridgeable gap is created in this case between the economists' purely formal sense of 'utility' and philosophical conceptions of utility as preference-satisfaction, whatever the motivations or reasons assumed to underlie the preferences. In particular, an ordinal utility function might be used to represent the information about the estimated quantities of pleasure to be expected from outcomes contained in an ordinary hedonistic preference ordering restricted to a given kind or quality of pleasure. But a utility function cannot be used to represent the information about the different qualities of pleasures to be expected from outcomes contained in the lexical preference ordering.

[61] Given that higher intellectual, moral, and aesthetic capacities must be developed to appreciate higher pleasures such as that associated with the moral sentiment of justice, the majority competently acquainted with a higher kind of pleasure is not necessarily the same majority as the majority competently acquainted with a lower kind of pleasure. *A fortiori,* the majority competent to decide on a qualitative preference ordering may be an educated minority in some social contexts.

https://doi.org/10.1017/S0265052509090128 Published online by Cambridge University Press

higher kind of pleasure takes absolute priority over any quantitative ranking in terms of a lower kind. In effect, the lower kind of quantitative ranking will count for an overall hedonistic assessment of any pair of outcomes if the higher kind of quantitative ranking expresses indifference with respect to that pair, which is necessarily the case if the one outcome does not differ from the other in terms of some source or object of the higher kind of pleasure. Qualitative evaluation thus meshes nicely with quantitative evaluation within Mill's hedonistic doctrine.

How does the lexical ranking of different kinds of pleasant feelings resolve the quandary noted earlier in connection with the reliance on purely ordinal utility information? In other words, how does the competent majority's ranking of higher and lower pleasures remove the quandary about whether the interpersonally noncomparable orderings should embody estimates of personal happiness or of the general happiness? This is one of the most remarkable features of Mill's approach, to wit, the way in which the higher kind of pleasant feeling associated with the moral sentiment of justice brings about harmony between general happiness and any person's own happiness. As I have already indicated, Mill suggests that most people who are competently acquainted with the different kinds of pleasures agree, at least upon reflection, that the pleasure of justice is infinitely more valuable than any competing kinds of enjoyments. This higher pleasure of justice, he explains, is an enjoyable feeling of "security" that can only be fully experienced under an effective and enduring social code that impartially distributes equal rights and correlative duties to all. The rights which are distributed and enforced by society are the source of the individual's security:

> When we call anything a person's right, we mean that he has a valid claim on society to *protect* him in the possession of it, either by the force of law, or by that of education and opinion. If he has what we consider a sufficient claim, on whatever account, to have something guaranteed to him by society, we say that he has a right to it.[62]

After making clear that everyone feels the immense importance of enjoying security from grievous injuries at the hands of others because "on it we depend for all our immunity from evil, and for the whole value of all and every good, beyond the passing moment," Mill goes on to say that this "extraordinarily important and impressive kind of utility . . . cannot be had, unless the machinery for providing it is kept unintermittedly in active play."[63] Among other things, a legal code of justice must be enacted, and the rights and duties distributed under it must be continually enforced

[62] Mill, *Utilitarianism,* 250; emphasis added.
[63] Ibid., 251.

https://doi.org/10.1017/S0265052509090128 Published online by Cambridge University Press

through the cooperative efforts of various government officials. Everyone must be taught the importance of the rule of law and of respecting others' rights even on those occasions when legal sanctions against violation may be inexpedient:

> Our notion, therefore, of the claim we have on our fellow creatures to join in making safe for us the very groundwork of our existence, gathers feelings round it so much more intense than those concerned in any of the more common cases of utility, that the difference of degree (as is often the case in psychology) becomes a real difference in kind. The claim assumes *that character of absoluteness, that apparent infinity, and incommensurability with all other considerations,* which constitute the distinction between the feeling of right and wrong and that of ordinary expediency and inexpediency.[64]

Given that the security afforded by equal rights and duties is infinitely more valuable than any competing enjoyments, the social code of justice has lexical priority over any competing considerations of expediency. An individual's equal rights can never be legitimately overridden without his consent to promote other people's happiness because even a bit of the higher pleasure of security—no matter who feels it—is intrinsically more valuable than any finite amount of lower pleasures—no matter how many different persons are assumed to experience them.

The qualitative superiority of the pleasure of justice over competing kinds of pleasures implies that any rational hedonist will give absolute priority to considerations of justice over competing considerations. Moreover, his maximization of his own happiness is logically compatible with his fellows' maximization of their personal happiness and thus with maximization of the general happiness regarded as nothing but the simultaneous maximization of everyone's personal happiness. His own security is maximized if and only if everyone else's security is also maximized because *equal rights are distributed to all.* Codes that distribute equal rights are the sole source of the higher pleasure of security associated with the moral sentiment of justice. No competing kinds of pleasures should ever override this enjoyable feeling of security according to most people competently acquainted with the different kinds. Thus, a rational individual's own happiness necessarily coheres with other rational individuals' personal happiness to the extent of their equal rights. Neither personal utility nor general utility can be promoted by violating rights.

To even form a moral sentiment of justice, an individual must be able to identify the social rules of justice with which to comply. Until he knows the particular code which ought to be accepted, he cannot know the particular equal rights and correlative duties which ought to be recog-

[64] Ibid.; emphasis added.

https://doi.org/10.1017/S0265052509090128 Published online by Cambridge University Press

nized by everyone within his community as belonging to him and anyone else in like circumstances. But to establish rules and rights that are publicly endorsed by his society in its laws and conventions, the individual must participate with his fellows in a political process, given that an ideal observer is not available to determine the best moral and legal code. An open process of free discussion and debate is essential for fallible beings to assess proposals and converge on an optimal code, that is, a code that impartially distributes those particular equal rights and duties which, at least so far as fallible people can tell, maximize the amount of security enjoyed by anyone and everyone who possesses them. In short, the individual's sentiment of justice presupposes a utilitarian political procedure to select an optimal security-maximizing code upon which any just individual must rely to guide his conduct.

As discussed earlier, Mill apparently takes for granted a political procedure such as majority rule which works with purely ordinal utility information. In this context, such information is contained in interpersonally noncomparable individual orderings defined over distinct proposed codes and rights. Any rational individual would consistently rank the codes in terms of the estimated amount of personal security that he expects from the relevant rules and rights. In doing so, he also implicitly ranks them in terms of estimates of general security because the codes distribute equal rights to all. Moreover, given that the individuals are competently acquainted with the kind of pleasures involved, majority judgments can be plausibly seen as maximum-likelihood estimates of the genuine amounts of general security associated with the different codes. In this process, there is no assumption that any individual can come up with a precise measure of his relative preference intensities. Nor is there any assumption that meaningful interpersonal comparisons can be made. If virtually all adults have developed the high intellectual and moral capacities required to estimate the amounts of security reasonably to be expected from the distinct codes and the equal rights distributed by them, then each individual should be given an equal vote.[65] In that special case, the ordinal utilitarian political process used to select an optimal code of justice is equivalent to some form of egalitarian democracy. More generally, however, in social contexts where some individuals are significantly more competent than others, Mill recommends a form of popular government in which the more educated are given plural votes to enable

[65] It should be recalled that a scoring function may be used to remove inconsistencies in majority rule. (See note 49 above.) Strictly speaking, majority rule per se does not give equal weights to different persons' preferences. Rather, it impartially counts the different noncomparable preference orderings without attempting to weigh them relative to one another, and it reflects the shared preferences of the greater number. A scoring function can be superimposed to mimic the way in which ballots are distributed to voters, so that one person's preferences may be counted for exactly as much as another's in the stronger sense that each gets equal weight in the political decision. But the ballot typically covers only the voter's top-ranked alternative instead of his entire ordering.

https://doi.org/10.1017/S0265052509090128 Published online by Cambridge University Press

https://doi.org/10.1017/S0265052509090128 Published online by Cambridge University Press

them to elect a sufficient number of representatives to influence legislative debates and impede incompetent and unjust legislation.[66]

It emerges that, for Mill as interpreted, rational agency is linked logically to an unusual plural ordinal utilitarianism by virtue of the higher pleasure of the moral sentiment of justice. A rational individual necessarily endorses justice understood in terms of equal rights because there is no higher competing kind of enjoyment than his feeling of security associated with this moral sentiment. And he necessarily endorses utilitarianism because the sentiment of justice itself inherently involves a utilitarian political procedure that relies solely on purely ordinal utility information—information contained in noncomparable individual preferences defined over distinct proposed social codes of equal rights—to impartially determine an optimal code that distributes the particular rights and correlative obligations which most if not all competent (yet fallible) people (or their representatives) decide are best for protecting every individual's vital interests.

Clearly, a rational individual thus understood is not a narrowly selfish egoist but rather an enlightened moral agent who appreciates that he has higher kinds of utilities or "permanent interests as a progressive being," in particular, a higher interest in the kind of security that can only be experienced fully under a code of equal rights.[67] He sees himself as a social being who strongly prefers to be in unity with his fellows to the extent of their reciprocal rights. Moreover, he understands that the kind

[66] Mill was a firm defender of popular government with a universal franchise, although he was also far more inclined than either his father James Mill or Bentham to institute constitutional checks and balances against unrestrained majoritarianism. His notorious device of plural votes based on education should be regarded as a check against abuses by ignorant or tyrannical majorities. But the device as he describes it would not permit the educated minority to elect so many representatives that they could pass legislation opposed by the representatives of the popular majority, keeping in mind that he also advocates a system of proportional representation. In the special case of an ideal society of highly educated citizens, there would be no need for plural votes. For further discussion of his argument that a form of constitutional representative democracy is the best form of government for any civil society, see Jonathan Riley, "Mill's Neo-Athenian Model of Liberal Democracy," in Nadia Urbinati and Alex Zakaras, eds., *J. S. Mill's Political Thought: A Bicentennial Reassessment* (Cambridge and New York: Cambridge University Press, 2007), 221–49.

[67] Mill tells us that he regards "utility as the ultimate appeal on all ethical questions; but it must be utility in the largest sense, grounded on the permanent interests of man as a progressive being" (*On Liberty,* 224). With respect to the interest in security associated with justice, security is a kind of pleasure or satisfaction (including freedom from pain) which can only be had in the context of a code of equal rights. The amount of security experienced by any individual varies across different codes, depending on whether he is regarded as a member of the relevant community to whom equal rights are distributed and, if so, on the contents of the rights distributed to all members. If he is excluded from the group, then he experiences no security. For any given population, general security can be maximized only if equal rights are distributed to all individuals, none being excluded. Mill realizes that the idea of justice as equal rights has been clouded historically by mistaken notions of expediency that call for excluding this or that subset of people from the rights distribution. But this is simply another way of saying that history is a record of injustice, of unfair societies in which general security has not been maximized. See, e.g., Mill, *The Subjection of Women,* in *CW,* XXI.

https://doi.org/10.1017/S0265052509090128 Published online by Cambridge University Press

of satisfaction which he experiences as security is infinitely more valuable to him than any satisfactions that conflict with it, and thus he refuses to sacrifice even a bit of it for any finite amount of lower enjoyments that involve violations of others' equal rights.[68] Security is a pleasure or satisfaction that can and should be enjoyed *permanently,* he believes, in harmony with many other kinds of enjoyments that are compatible with (and, in some cases, made possible by) the code of equal rights and duties. That is ultimately why he agrees with his fellows that the machinery for providing security—the legislature that enacts a legal code, the executive that enforces the code, the judiciary that hears and decides cases and controversies as they arise, the system of education that impresses the importance of equal justice on all, and so forth—must be "kept unintermittedly in active play."[69]

At the same time, the machinery must not be conflated with the pleasure or satisfaction of security itself. Strictly speaking, it is a separate issue to what extent legal institutions and social practices that rely on coercion are expedient for providing security. Indeed, we can imagine a social utopia where rational individuals jointly frame a moral code and conscientiously perform their duties in accord with the rules despite the absence of any threat of legal or social punishment for noncompliance. In such a utopia, the guilt—the kind of suffering associated with doing something morally wrong—which any individual expects to experience from breaking the rules is, by assumption, sufficiently powerful to enforce his voluntary compliance with the code: legal penalties and public stigma are not needed. Outside utopia, however, the issue of how best to employ various forms of coercion to discourage noncompliance by unreasonable people must be considered. Even so, as Mill emphasizes, this does not imply that it is generally expedient to invariably employ legal sanctions for the purpose. In some situations, expedience may justify the use of public stigma without legal penalties, for example, or even reliance on the internal sanctions of conscience alone.[70] *A fortiori,* there is no suggestion that government should expend virtually all of its resources on enforcing the laws of justice by identifying, prosecuting, and punishing lawbreakers. Rather, the political authorities must recognize the limits within which legal coercion can expediently operate, and they must also expend sufficient resources to sustain basic rights of subsistence, disability insurance, health care, education, and so forth, not merely rights to physical safety and private property.

[68] For Mill, the lexical priority of a code of justice over competing considerations is justified by the infinite superiority of this higher kind of utility over competing kinds. In contrast, Kant and John Rawls in their different ways trace the absolute priority of the code to a self-justifying principle of practical rationality or procedure of public reasoning that transcends and constrains any person's idea of his utility or good. See, e.g., John Rawls, *Political Liberalism,* expanded edition (New York: Columbia University Press, 2002).

[69] Mill, *Utilitarianism,* 251.

[70] Mill, *On Liberty,* 224–25.

VI. A Credible Maximizing Utilitarianism?

Plural ordinal utilitarianism arguably avoids the calamities that plague the more familiar versions of maximizing utilitarianism, crude or otherwise. General utility is conceived in such a way that it cannot conflict with personal utility. Absolute priority is given to a code of equal rights and correlative duties in virtue of the qualitative superiority of the pleasure of security associated with the moral sentiment of justice over any competing kinds of pleasures. Since codes of equal rights are the sole source of this higher kind of pleasure, any rational individual will participate with his fellows in a political process that impartially counts their respective rankings of distinct proposed codes for the purpose of selecting an optimal security-maximizing code. Each person's ranking is based on his estimates of the quantities of personal security he expects to experience under the different codes. But his estimates are at the same time impartial estimates of general security because he is constrained to rank codes of *equal* rights. The political outcome will be a code that distributes those particular equal rights which most of the participants believe will best protect not only their own but everyone else's vital personal interests.[71] Such a code, effectively enforced, simultaneously maximizes every person's own security as well as the general security thus conceived. Given human limitations, this political process is reasonably expected to take place in a gradual and piecemeal fashion, with different parts of a code being chosen at different times and rules being repeatedly reconsidered to correct mistakes and allow for new information.

It is worth emphasizing that any rational participant in the political process is required only to assess and rank different codes of equal rights in light of his own security. He is not required to deliberate about any independent notion of the general welfare. His prudent evaluation of the rules and rights that he expects will best protect his own vital concerns is sufficient input because it is also implicitly an impartial evaluation of the means of maximizing general security, even though he may not be consciously thinking of the general welfare. The overwhelming importance of their own security to most rational people may reasonably be trusted to lead them to converge on an optimal code of equal rights which genuinely does maximize general security.

Moreover, there is no need for interpersonal comparisons or cardinal measurement for any individual either to make these prudent evaluations of his own security or to form a moral sentiment of justice which involves his compliance with the optimal code of equal rights established through the political process. To propose or assess any rule that distributes equal rights and correlative duties of a particular content to the various people

[71] I am ignoring the relatively rare occasions on which the scoring procedure may fail to select a Condorcet winner when it exists.

https://doi.org/10.1017/S0265052509090128 Published online by Cambridge University Press

involved in a given situation or type of situation, any competent individual must imagine himself at each personal position with the particular rights and duties distributed to that position, must estimate how much personal security the rights at a given position may be expected to give him in that position, and then must aggregate over the positions to form an estimate of his own security overall were he ever to actually find himself in such a situation. This is necessary given that he cannot know beforehand which of the various positions he will occupy and thus might occupy any of them with equal probability. The individual may be assumed to form a personal preference ordering defined over alternative possible rules and rights proposed to govern the relevant situation, with the top-ranked alternative being the rule that distributes the particular set of rights and duties across the positions which he estimates will maximize his overall personal security, whether conceived as his average security at each position or his total security at all the positions combined. There is never a need for the individual to sympathize with other people by adopting their personal preferences or the estimates of security which they expect to experience at any position. Nor does he require a cardinal measure of the gains or losses of his own security as he moves across the personal positions in the given situation. Even without such precise utility information, he can roughly estimate how much overall personal security to expect under any proposed code of rights and duties in comparison to other proposals. For any pair of proposals, he can compare in terms of more or less the amounts of security he expects on a position-by-position basis, and then come up with a rough judgment about which proposal offers him more overall security than the other. The lack of precision may lead to significant ranges of indifference in his personal preference ordering because he cannot see that his overall security alters across the distinct proposed rights and duties within any such range. But the indifference is not problematic. In any case, the personal orderings defined over any given possible alternatives serve as input into the political process, and majority voting results in an optimal code that distributes equal rights and duties to all in various situations. This optimal code must be established before any individual can complete the formation of his moral sentiment of justice. Until he can recognize the particular code of equal rights which he ought to recognize, he does not know how he should act to maximize general security in any situation.[72]

[72] Unlike the sophisticated utilitarianisms discussed earlier, Millian utilitarianism supposes that any prudent and impartial individual must be willing to universalize his prescriptions of the rights which he thinks are best for promoting his own security. Thus, he is willing to assert a general rule that distributes equal rights and correlative duties which all rational agents might adopt with a view to promoting their own security. In addition, he must be willing to comply with the rule which emerges as the outcome of the impartial political process, even if the rule favored by the majority is not his personal favorite. In contrast, both Hare and Harsanyi require, in effect, that the individual must be willing to universalize his prescriptions so as to maximize an independent notion of general welfare

https://doi.org/10.1017/S0265052509090128 Published online by Cambridge University Press

By adopting the perspective of a prudent individual who seeks to maximize the amount of security he feels for his own vital interests under the constraint that his fellows must be afforded like treatment, we can infer the nature of the various equal rights and correlative duties which are likely to be distributed by an optimal code. These particular rights may be expected to include claims to be protected from unprovoked physical attacks by other people, for instance, as well as rights of criminal due process, rights to keep enough of the fruits of one's own labor and saving to be able to provide a living for oneself and one's family, rights of contract, rights to basic support from others when in childhood, disabled, injured, or involuntarily unemployed, and so forth.[73] But no competent judge will vote for equal rights to be free of a bad headache at the cost of another's life. His own security would obviously be compromised by such a right because at any time *he* could be the one called upon by the correlative duty to die merely to prevent others from suffering a headache. Nobody's vital interests are at stake merely because they have a headache. The pain of a headache is a relatively trivial kind of harm that is qualitatively inferior to any grievous injury of the sort that ought to be prevented by right.[74]

The focus on personal security also suggests interesting Millian utilitarian lines of response to hard cases such as those listed earlier (in Section III). Consider the case of the Carthaginians required by the Romans to surrender Hannibal or face extermination. Mill makes clear what he thinks the "morality of utility" implies in such a situation:

which is fully understood only by a fictitious impersonal observer in possession of the requisite cardinal and interpersonally comparable utility information of homogeneous quality.

[73] Instead of giving extra weight to the utilities of the worst-off people when calculating overall good, as in Brad Hooker's rule consequentialism, Mill's utilitarianism as interpreted protects the worst-off by impartially distributing an equal claim to subsistence or to some minimum level of income guaranteed by the general taxpayer. Security from abject poverty is certainly something that most rational and impartial agents can recognize as a part of justice, at least in societies where a threshold level of collective wealth has been achieved. In addition, the Millian doctrine prescribes moral duties of charity, which are imperfect in the sense that they do not correlate to anybody's rights, and thus, unlike perfect duties of justice, admit some personal discretion as to the occasions of their performance. Supererogatory acts of giving to others are also viewed as praiseworthy, although they are not morally required. This tripartite classification of acts of redistribution into just, charitable, and supererogatory is seen as a dynamic social construction that evolves as the members of society develop their intellectual, moral, and practical capacities. As society advances, competent majorities may choose to reclassify acts formerly seen as acts of charity into acts of justice, for instance, so that a right to some level of poor relief comes to be accepted as morally required. Or acts of sacrifice formerly viewed as supererogatory may come to be seen as morally obligatory on at least some occasions. For further discussion, see Jonathan Riley, "Mill's Extraordinary Utilitarian Theory of Morality," *Politics, Philosophy, and Economics* 8 (2009). Hooker's distinct approach to these matters is discussed in his *Ideal Code, Real World*, 147–74.

[74] Recall that the qualitative difference implies a discontinuity of value in any list of pains or dissatisfactions such that any finite amount of injustice—the kind of suffering that is associated with violations of equal rights—is regarded by most rational people as incomparably worse than any competing finite amount of lower kinds of dissatisfactions or annoyances, however large.

https://doi.org/10.1017/S0265052509090128 Published online by Cambridge University Press

> [T]he morality of utility requires that the people should fight to the last rather than comply with the demand: not only because of the special tie between the community and each of its members, and between the community and a benefactor who in the case supposed is demanded as a victim precisely because of the greatness of his services—but also for a more general reason—namely the reason which makes it right that a people inferior in strength should fight to the death against the attempt of a foreign despot to reduce it to slavery. For such iniquitous attempts, even by powers strong enough to succeed in them, are very much discouraged by the prospects of meeting with a desperate though unsuccessful resistance. The weak may not be able finally to withstand the strong if these persist in their tyranny, but they can make the tyranny cost the tyrant something, and that is much better than letting him indulge it gratis.[75]

The Carthaginians must fight together to defend their general security, that is, their equal rights not to be killed or enslaved by the Roman aggressors. This is not only because the moral sentiment of justice involves a feeling of social unity which binds any rational person to his fellows in a joint commitment to a code of equal rights. Any rational individual can appreciate the great importance for his own security of cooperating in the fight. By arbitrarily demanding the surrender of Hannibal, the Romans have shown that they are capable of arbitrarily demanding the surrender of any Carthaginian citizen who attracts their notice. There is no reason to expect them to refrain from repeating the demand, given their view of Rome's interests. Indeed, there is no reason to expect them to refrain from destroying or enslaving Carthage even if Hannibal and others are surrendered, especially because this goal will be easier to achieve after the surrender. Thus, a prudent agent understands that fighting gives him at least as much chance as capitulating does to avoid death or slavery. Only a fool would believe otherwise.[76]

Analogous things can be said about the case of an individual required by an unjust kidnapper to kill an innocent victim or face the execution of multiple others (including himself) at the hands of the kidnapper and his followers. A utilitarian morality along Mill's lines requires the individual to fight even to the death, perhaps by taking the gun, knife, or other weapon intended for use against the innocent and using it instead against the kidnapper. Similarly, a sheriff required by an unruly mob to hastily hang an innocent stranger or face widespread violence should threaten to shoot the mob leaders and, if necessary, actually do so, even if he dies in the aftermath. Unless his threat or act of violence brings the mob to its

[75] Mill, Letter to Thornton, in *CW*, XV, 854.

[76] As it turned out, the relentless Romans eventually did destroy Carthage at the end of the Third Punic War in 146 B.C.

https://doi.org/10.1017/S0265052509090128 Published online by Cambridge University Press

senses, these people are hardly likely to rest content with a single victim. After the rampage is over, sober heads, after duly investigating the incident, would in any case properly execute or incarcerate any sheriff who was found to have recklessly hung an innocent person. Under the circumstances, the iniquity of the hanging, committed in haste, would easily be discovered.

As for the runaway trolley car case, it does not seem prudent in terms of personal security to vote for an equal right to invariably be rescued at the cost of another innocent person's life whenever his death would save five or ten or even five hundred people. This is not so much because we can never be certain about how successful the person's rescue attempt will be, although the uncertainty is a relevant consideration. Even if success were certain, it can be inferred that personal security would not be noticeably increased by a rule that distributes equal rights of this nature. Given that there is little likelihood of being in such a runaway trolley car situation, the probability of being one of the people who would be saved from accidental death at the cost of another's life is not noticeably different from the probability of being in a position to rescue the others by deflecting the grave harm onto oneself. Keep in mind that the measure of the amount of security to be expected from different rights and duties (including the option of no rights or duties) is purely ordinal: there is great imprecision in the estimates (even the competent majority's estimates) of the gains or losses of security to be expected from the different options. Given the tiny probabilities involved in the situation (one chance in a billion of being the doomed rescuer, say, and five hundred chances in a billion of being one of those saved by his death), the tiny difference of probabilities is swamped by the imprecision of the ordinal measure. It follows that, in light of our impoverished utility information, the proposed right in question cannot be justified as a source of increased personal security, even assuming that we can ignore the possibility of mistakes resulting in more deaths because of botched rescues. An optimal code of justice thus should not distribute such rights. At the same time, any individual's willing sacrifice of his own life to save others in these situations can be viewed as a noble and beautiful supererogatory act deserving of high praise.[77]

This is not to say that there can never be a justification in terms of security for equal rights whose correlative duties may require the individual to sacrifice his own life or other vital interests for the benefit of

[77] Police, firefighters, medical personnel, and other professionals may well be much more likely than the ordinary person to find themselves in tragic situations akin to the runaway trolley case. A separate argument might be made for rules that distribute professional obligations to risk one's life to rescue a greater number of other people from accidental death in these situations. Such obligations would become part of the job description. In any case, it would not be surprising if professionals developed some such convention to govern interactions among themselves.

https://doi.org/10.1017/S0265052509090128 Published online by Cambridge University Press

https://doi.org/10.1017/S0265052509090128 Published online by Cambridge University Press

others.[78] Depending on the likelihood of being in the situation and the numbers of people who can be saved, even the trolley car example can be construed to justify such a right. But consider instead a disaster scenario involving a fanatical terrorist who is threatening to destroy an entire city or even the whole country with advanced weapons. Suppose that there is a single innocent person in a position to stop the attack by setting off a bomb that will kill both himself and the terrorist. Unless the terrorist is stopped, the vast majority of the population will be annihilated. More specifically, any individual is virtually certain to die in the attack, whereas the probability of being the individual in a position to stop the terrorist is one in a billion.

Remarkably, security is enhanced by a rule that obliges anyone in the position of the single innocent individual to stop the terrorist by intentionally committing suicide under these conditions. Given the relative probabilities, any prudent individual appreciates that such a rule gives him a much better chance of surviving these situations, even though he will be obliged to sacrifice his life if he is unlucky enough to be the one who can stop the terrorist. General security is thus maximized by giving everyone an equal right to the innocent victim's sacrifice, which he has a correlative duty to perform. Unless he has developed a powerful sentiment of justice, however, the relevant individual may well ignore the rule because nobody—not even himself—is likely to be left alive to inflict punishment on him for violating his duty!

Even if these utilitarian lines of response to the hard cases are contestable, critics cannot pretend that equal rights are given short shrift in the responses. In particular, deontological liberals who merely assert that given rights dictate this or that response, without further justification, do not provide a convincing alternative. Plural ordinal utilitarianism offers a justification of the rights that should be distributed by pointing to their consequences in terms of personal security, keeping in mind that security is seen as a kind of utility which is infinitely more valuable than competing kinds. We may argue about the particular rights that would best promote security, but at least there is a way to resolve the argument.

It emerges that, under this Millian version of maximizing utilitarianism, it is legitimate to coerce an individual, if necessary, to make him comply with an optimal security-maximizing code of justice, but, as long as he complies, his equal rights give him capacious freedom to organize his life and conduct as he pleases.[79] This freedom allows the individual to pursue the various other kinds of pleasures available to human beings while at the same time enjoying the protection for his vital interests which

[78] Brad Hooker also makes this point (*Ideal Code, Real World,* 169).

[79] This point about freedom does not require acceptance of Mill's doctrine of a moral right to absolute liberty of so-called purely self-regarding conduct. For a defense of the self-regarding liberty doctrine, see Riley, *Mill on Liberty;* and Riley, *Mill's Radical Liberalism* (London: Routledge, 2009).

is provided by his rights and the duties they imply for others. By implication, any individual's ranking of the possible outcomes in terms of the amounts of these various other kinds of pleasures is counted impartially along with his fellows' rankings of this sort, without any attempt by society or its political representatives to interfere coercively with what these individuals do on this basis, and thus without any need for a political process such as majority voting to be repeated, given that the individuals are acting within their rights and therefore not impinging on anyone else's security. In effect, plural ordinal utilitarianism counts any person's *enjoyment of his freedom* to pursue his own projects in accordance with his rights for exactly as much as another's similar enjoyment. This pleasure of doing as one wishes within one's rights is apparently the aspect of personal happiness that Mill refers to as individuality or self-development in *On Liberty.* Thus, the normative principle of impartiality which is at the heart of maximizing utilitarianism is not overly stringent in the context of this Millian version of the doctrine. Scheffler's sweeping claim to the contrary is refuted.[80]

No doubt more needs to be said to clarify plural ordinal utilitarianism.[81] The doctrine can, for instance, consistently accommodate any individual's free and undeceived choice to give up his equal rights and security in order to engage in virtuous acts of self-sacrifice to help others beyond the call of moral duty. Provided it is done voluntarily, his sacrifice of his own vital interests, including even his life, may reasonably be thought to allow others to enjoy so much more security from their rights that general security increases despite his insecurity and suffering. As Mill puts it:

> The utilitarian morality does recognize in human beings the power of sacrificing their own greatest good for the good of others. It only refuses to admit that the sacrifice is itself a good. A sacrifice which does not increase, or tend to increase, the sum total of happiness, it considers as wasted.[82]

In this regard, the higher pleasure of security is not necessarily the *highest* kind of pleasant feeling. Mill himself suggests that nonmoral aesthetic and spiritual pleasures, including those associated with praiseworthy supererogatory acts, may be qualitatively superior to the pleasures of

[80] As a result, Millian utilitarianism as conceived can consistently endorse both "an Ideal of Purity and an Ideal of Humanity," in Scheffler's terms (*Human Morality,* 101–32). For Scheffler's claim about stringency, see Section III above.

[81] For further discussion of my interpretation of Mill's utilitarian theory of justice, see Jonathan Riley, "Justice as Higher Pleasure," to be published in *John Stuart Mill: Thought and Influence,* ed. Paul J. Kelly and Georgios Varouxakis; and Riley, "Mill's Extraordinary Utilitarian Theory of Morality." For a Millian response to Sen's Paretian liberal impossibility result, see Riley, "Liberal Rights in a Pareto-Optimal Code," *Utilitas* 18 (March 2006): 61–79.

[82] Mill, *Utilitarianism,* 218.

https://doi.org/10.1017/S0265052509090128 Published online by Cambridge University Press

justice. But these emotional enjoyments of the highest kind do not conflict with justice. Thus, an individual might nobly choose to waive his rights—even die—so as to save the lives of others, for instance, or freely devote his life to helping those in need, even though he has no moral duty to do so. Such sublime activities may bring only a tiny amount of the highest kind of satisfaction to the agent, or they may even bring him only freedom from suffering the worst kind of pain, namely, a kind of spiritual suffering associated with the idea of living in security when one could have prevented grave injury or even death to others. But, for the noble individual who is capable of experiencing them, Mill implies, feelings of this type may be infinitely more valuable than any others.

This illustrates that the Millian doctrine can accommodate some disagreement between the majority and a minority with respect to the qualitative ranking of different kinds of pleasures. A few noble individuals may freely choose to sacrifice their personal security and even lose their lives to deflect grave harms that would otherwise befall others. In effect, in situations where someone's vital interests will inevitably be sacrificed, these noble few prefer to forgo their own rights rather than see others stripped of their equal rights by factors beyond their control. These few are willing to give up what most may consider the most valuable component of personal happiness, namely, the security afforded by their equal rights, in order to experience a kind of personal happiness which the few consider qualitatively superior even to their own security. Since this highest kind of enjoyment is, for those capable of experiencing it, infinitely more valuable than any lower pleasure including security, their choice to experience even a bit of it instead of a large amount of security and comfort over a lifetime can be a rational choice for them, even if such a choice is not rational for the rest of us who are not similarly capable. There is no conflict for them between their idea of personal happiness and the happiness of everyone else: they are willing to waive their own rights, if need be, to maintain the equal rights of others and thereby promote others' security. *A fortiori,* their freedom to engage in the relevant acts of self-sacrifice promotes their own happiness while simultaneously promoting the happiness of others.

VII. Conclusion

A plural ordinal utilitarianism along Millian lines deserves serious consideration as a credible alternative to the more familiar versions of maximizing utilitarianism. By rejecting the received wisdom that cardinal and interpersonally comparable utilities of a homogeneous quality are essential to the structure of utilitarianism, the Millian variant seems able to immunize itself against objections that are fatal to the familiar versions. In particular, it is able to recognize the overwhelming importance of a code of justice and equal rights in comparison to competing considerations by

https://doi.org/10.1017/S0265052509090128 Published online by Cambridge University Press

treating the kind of pleasure or satisfaction associated with the moral sentiment of justice as qualitatively superior to, and thus infinitely more valuable than, any competing kinds of satisfactions. This kind of satisfaction, namely, security, can only be provided by codes of equal rights, and thus it is maximized for any one individual only if it is simultaneously maximized for all.

Consistently with security-maximization, the complex Millian utilitarian counting principle also, in effect, impartially counts any person's enjoyment of his freedom to do as he wishes in accordance with his rights and duties distributed under the code. Indeed, any individual can even freely give up his own rights and security so as to promote what he considers the greater happiness of others, although he has no moral obligation to do so. Most rational and impartial people may not accept that such a beautiful act of self-sacrifice brings a kind of personal satisfaction that is qualitatively superior to security itself, even if they can appreciate the beauty and may stand in awe of the noble few who are capable of such self-renunciation on behalf of others. It is thus pointless to expect even a competent majority to establish moral duties of such self-sacrifice. More importantly, given that they remain unconvinced that any aesthetic or religious pleasure of self-sacrifice is far more valuable than security, it cannot be demonstrated to the majority that their self-sacrifice (including even death) would *increase* the general happiness. After all, the pleasant security of justice can only be had in the context of a code of equal rights. All competent judges, including the noble few, agree that a person's equal rights must not be overridden without his consent because that would tend to destroy the sole source of everyone's feelings of security. The noble few may wish to waive their own rights and security, but they do not pretend to decide the issue for others. The majority wish to retain their rights and security, and, in their estimation, there is no qualitatively superior kind of satisfaction which requires them to give up the feeling of security. It cannot be assumed that an ideal impersonal observer with meaningful cardinal and interpersonally comparable utility information might, in principle, provide a definitive resolution of this disagreement between the majority and the minority. To think otherwise is to return to the delusions of the more familiar versions of maximizing utilitarianism.

The upshot is that plural ordinal utilitarianism gives rise to a conception of morality as *nonpervasive, overriding,* and *nonstringent* (or at least not overly stringent), in Scheffler's terms. Morality is nonpervasive or limited in scope: there are spheres of conduct, including a purely self-regarding realm as well as a sphere of supererogatory acts of self-sacrifice on behalf of others, which are beyond moral assessment although they remain open to prudential and aesthetic standards. Morality is also overriding: a rational agent must never allow competing considerations to dissuade him from fulfilling his moral duties to others, even though he may be free to waive his own rights on aesthetic grounds which a majority of competent

https://doi.org/10.1017/S0265052509090128 Published online by Cambridge University Press

judges reject as imprudent. Finally, morality is not overly stringent: the moral obligations of justice are not too demanding for human beings, even though an individual may be obliged to fight to the death or sacrifice his life for others in extraordinary situations. Indeed, a noble few will find their moral duties to be *too moderate:* these individuals will choose to perform supererogatory acts in accordance with their own aesthetic (if not prudential) standards. But I must leave for another occasion further discussion of these matters and their implications for our understanding of Mill's utilitarian Art of Life.

Philosophy and Political Economy, Tulane University

https://doi.org/10.1017/S0265052509090128 Published online by Cambridge University Press