

1

Value, welfare, and morality

R. G. FREY AND CHRISTOPHER W. MORRIS

In addition to an account of motivation, every normative ethical theory must contain an account of value, and the link that is often forged between what is valuable and what it would be right or what we ought to do is human welfare or well-being. While this linkage is not new, it has become a source of considerable controversy in contemporary ethics, not least because of the general reconsideration of utilitarianism that has been underway for some time now. Indeed, it often seems that there is as much disagreement about the nature of value and its relationships to welfare and morality as there is about the substantive issues on which our normative theories are supposed to pronounce.

Both in philosophy and the social and policy sciences, two general pictures of value and of its link to welfare and morality haunt contemporary discussions in the Anglo-American world. Both general pictures are inherited from David Hume and Jeremy Bentham and from the tradition of political economy founded by Adam Smith and David Ricardo. From Hume, there develops an instrumental view of rationality, of reason in the service of the passions; while reason can assist one in obtaining one's ends, what these ends are depends upon one's passions. This widely influential view of rationality is associated today with desire- or preference-based conceptions of welfare (and with conventionalist or even contractarian accounts of justice). From Bentham and nineteenth-century economics, there develops the rival though quite friendly view of utilitarianism, in which the focus is on aggregation of pleasures and pains over persons generally and in which reason assists one in the pursuit of this impersonal though welfarist end. As they are fleshed out and given substance, these two pictures come to form general accounts of value, welfare, and morality, and they form the background for and underlie the treatment of numerous issues in value and moral theory. It may be well, then, to indicate at the very outset of this volume some of what goes toward fleshing out these views.

The neo-Humean view – it differs in some respects from Hume's actual account of matters – starts with an instrumental conception of practical

reason, according to which persons are rational to the extent that their behavior is an efficient means to their ends. These ends, however, are not discovered by reason but are determined by desire or, in contemporary parlance, preference. There are quarrels among neo-Humeans as to the coherence and other conditions (e.g., information) that rationality can impose on preferences, but these intramural disputes need not detain us now.

All neo-Humeans agree in forbidding substantive conditions to be imposed on the content or nature of preferences; rationality has essentially to do with the manner in which agents take their preferences to guide their actions. To speak, then, of the “rationality of ends” is inappropriate, for this would be to ask for something more than a purely instrumental conception of rationality and to suppose that human ends are set by something other than or in addition to preference.

Importantly, it is almost certainly the case that, while neo-Humeans reject conditions on the content or nature of preferences, they will be constrained to accept some conditions on admissible preferences if they wish to hold, as many do, that instrumental rationality is to be understood as the *maximization* of the satisfaction of one’s preference, or that the measure of preference is utility. For maximization, properly understood, requires at least a weak ordering of alternatives or options, that is, a complete and fully transitive ranking of alternatives. The point is logical: One cannot *maximize* unless one can (1) compare all alternatives with respect to, for example, preference or goodness; and (2) order them in a way that satisfies transitivity. There is no greatest or best if x and y cannot be compared in the relevant manner, or if x is better than y , y better than z , and z better than x .

On the neo-Humean view, preferences guide action, with the result that action is typically explained by reference to the agent’s preferences and beliefs. This conception of rationality has an explanatory function, which is widely appealed to in the economic and social sciences, which in turn is taken to testify to its explanatory power. Indeed, it might be held that this explanatory power is nowhere more apparent than in everyday life, where human action seems commonly understood and explained by appeal to what the actor believed and desired.

Neo-Humeans agree that desires or preferences provide reasons for action, though there are disagreements, related to those mentioned above, as to whether *all* preferences do; but virtually all neo-Humeans agree that *only* preferences provide reasons. This has a rather striking result: Preference-independent interests and needs provide no reasons for action, whether the interests be one’s own or those of others. Of course, it is highly likely that one’s interests or needs will figure among

one's preferences, that is, that one will have preferences, of greater or lesser strength, with regard to the objects of one's interests or needs. But this in no way affects the point about reasons. As indicated, this claim that only preferences provide reasons for action is accepted by the vast preponderance of neo-Humeans. (It is worth noting that some neo-Humeans, such as David Gauthier, allow considerations that are not themselves preferences, albeit preference-based, to provide reasons for action.)

From this point of view, some neo-Humeans go on to interpret welfare as well-being, well-being as preference satisfaction, and the latter as the satisfaction of informed and/or self-regarding preferences. The link with welfare is then trivial. This general line of approach to welfare is characteristic of "welfare economics," a normative branch of economics; it no more treats of a preference-independent notion of welfare than does neoclassical economic theory generally.

The link between instrumental rationality and morality is more problematic. In fact, doubts are expressed by some writers as to whether a genuine morality is available to creatures who are merely instrumentally rational. For instance, it seems clear that some parts of morality, such as benevolence, might be held by neo-Humeans to be found in ordinary other-regarding preference. But other parts of morality may not be so easily accommodated, such as those associated with the virtue of justice, understood in the broad sense that includes truth-telling and fidelity, as well as respect for others' rights. For, as Plato and many others noticed long ago, justice frequently asks us to do that which we do not want to do, which it is not in our interest to do, and which it even sets back our interests to do. If this is so, if, that is, justice can require of us acts (and dispositions) that are not themselves utility-maximizing, then how is the neo-Humean account of rationality or value, based in preference, going to make room for justice? Put summarily, in a phrase of James Griffin's, what road do we follow from the good to the right?

One idea has been to adapt Hume's conceptions of justice and property to the task. Hume thought of the norms of justice and property as conventions that serve the general interest and that, as a result, one has reason to support, insofar as one's interests are included in those of the public. To be sure, Hume recognized that one's interests and those of others may be in opposition, as when one finds oneself, as he puts it, in the company of ruffians; in such situations, as between humans and animals, justice ceases to bind. Importantly, then, the mere fact that our interests are included in those of the general public does not entail that we individually have reason to support the norms that serve the public interest, as Hume's "sensible knave" and Hobbes's "Foole" both

recognized. This much is evident from the variety of familiar situations that are structured like the Prisoners' Dilemma. Nevertheless, the general strategy suggested by Hume is clear, and, in one form or another, it is developed, for example, by J. L. Mackie, David Gauthier, and Gilbert Harman.

The neo-Humean, then, begins with an instrumental conception of practical rationality, constructs a conception of welfare that provides reasons only insofar as it is preference-based, and develops a conception of morality, or, perhaps more accurately, justice, as a set of mutually advantageous or agreeable conventions. Additional features of this view will emerge in the discussion of the second, utilitarian position, to which we now turn.

If practical rationality requires individuals to maximize the satisfaction of their preferences, and if morality requires of us that we be impartial to the ends of one person over another, then it might seem natural to construe morality as asking that we maximize preference satisfaction overall. Certainly, it has seemed so to many moral and political philosophers, to economists and social theorists, and to numerous social reformers, starting with the British Philosophical Radicals.

As the neo-Humean and utilitarian traditions are in various respects close allies, many of the same views regarding the relations between desires, well-being, and reasons have their proponents here as well. Early utilitarians tended to understand welfare in terms of mental states, variously adumbrated as pleasure, happiness, benefit, and so on, and to think that people were moved to action by considerations of pleasure, happiness, and the like. (This view of motivation did not *require* that agents be construed as psychological egoists.) Later utilitarians came to find the mental-state view of intrinsic value and utility too confining and shifted to a desire-satisfaction account, wherein it is true that we desire more in our lives than merely pleasure or happiness. Of course, not all utilitarians have made the switch, and some remain mental-state theorists, for whom the fundamental datum of value is any of various experiences called pleasure, happiness, and so on. And other utilitarians seem to have endorsed both views, in that they apply the desire-satisfaction view in the case of humans and the mental-state view in the case of other animals.

Moreover, some utilitarians have refused either to identify preference with welfare or to take preference instead of welfare as the object of importance for morality. For these theorists, seemingly a growing class, it is well-being, understood as something independent of preference, that is what ought to be maximized. The debate here is more complicated than with neo-Humeans; for part of what is at issue is whether one's

welfare, understood as independent of one's preferences, provides one with reasons. In addition, does the welfare of others, all others, provide one with such reasons? It would seem that it must, if, indeed, we are to believe that morality demands that we maximize the well-being of all.

Talk of well-being understood as independent of one's preferences raises a tantalizing possibility: Can there be articulated an "objective" account of well-being, one that construes well-being to consist in certain factual conditions of human flourishing *and* that measures degree of success in achieving these factual conditions in some equally factual way? Some utilitarians have been intrigued by this possibility and by questions of whether individuals are the best judge of their well-being, of whether their word is final on the matter, of whether they can be deceived, and so on. Something like a measure of well-being as consisting in, among other things, caloric intake might illustrate the point: It becomes possible to plot well-being on a scale that can be applied to persons indiscriminately, and possession of such a measure would in turn seem to make it possible for one to refuse to take as definitive in the matter the individual's own statement of how well his life is going. Doctors in hospitals, for instance, frequently take themselves to have such measures, and economists and policy theorists interested in judgments of "social welfare" have commonly attempted to construct such scales.

Several points now arise of the utmost importance to the debate between neo-Humeans and utilitarians. A few words on each of these will have to suffice to show their crucial relevance.

First, as traditionally understood, utilitarianism is a maximizing theory that would have us maximize total preference or welfare satisfaction, that is, the greatest total amount of satisfaction distributed across all persons or, rather, all sentient beings. Let us call this the greatest total good. Some versions of the theory are "internalist" and claim that we each have reason to seek the greatest total good. Other versions are "externalist" and claim only that the greatest total good is the goal of morality; whether each of us has reason to seek it is a contingent matter. All these versions, whether internalist or externalist, however, seem committed to the view that the greatest good is of value, in some way, to all persons, that is, to each and every person. Is this the case? While it is easy to realize that my welfare, or that of my family and my friends, matters to me, and relatively easy to see how the welfare of (most) members of my community matters to me, it is not obvious that the general or collective welfare, understood as the greatest total good, matters to me. If it does not matter to me, then it seems unlikely that

the greatest total good can provide the motivational force to action that utilitarians, whether internalist or externalist, have thought that it provided.

Second, to claim that the greatest total good is of value, in some way, to all persons is to claim that the value of the greatest good is agent-neutral. By contrast, neo-Humeans claim that the good of others, much less the greatest total good, is of value to someone only if it is the object of one of her preferences. For them, the value of the satisfaction of someone's preferences or welfare is agent-relative. How exactly the agent-neutral/agent-relative distinction is to be drawn is a matter of increasing controversy. But the general idea is clear enough: Something has agent-relative value to the extent that it has value from particular perspectives, those of the particular valuers to whom it is valuable, whereas something has agent-neutral value to the extent that it is valuable, necessarily, from the perspective of all, or rather, of each and every valuer. (And while an agent-relative value could be a value from the perspective of all, it need not be.) Utilitarians generally affirm, and neo-Humeans generally deny, that the value of the greatest total good is agent-neutral.

Clearly, the dispute between utilitarians and neo-Humeans is of special theoretical importance. For if one could establish that there is no agent-neutral value, most versions of utilitarianism would be in trouble; whereas if one could show that there is no agent-relative value (as G. E. Moore perhaps thought) or merely that there is some agent-neutral value, most neo-Humean positions would be in trouble. The establishment of either thesis, then, would seem to carry import for the adequacy of versions of the position in question.

It is important to stress a difference between utilitarian and neo-Humean positions here. Utilitarians, as well as many philosophers who include some consequentialist principles in their accounts of morality, would appear to think that there is at least one agent-neutral value, but many seem also to hold that there are agent-relative values as well. Endorsement of the existence of agent-neutral value need not bar them from recognizing the existence of agent-relative value. Neo-Humeans, however, typically deny the existence of agent-neutral value and recognize only agent-relative value; all value whatsoever is agent-relative. This is a bold thesis that denies a possibility that many utilitarians and consequentialists allow. And it is not only they who allow it; recent work by Thomas Nagel affirms the existence of both agent-neutral and agent-relative value, without endorsing the claim that the greatest total good is an agent-neutral value.

Another interesting possibility bears mention as well: Neo-

Aristotelian virtue theorists such as Philippa Foot, while antagonistic to neo-Humeanism, deny the existence of any agent-neutral value. Thus, while neo-Humean accounts entail that all value is agent-relative, it should not be thought that rejection of preference-based accounts commits one to accepting (some) agent-neutral value, as shown by neo-Aristotelian ethics. We should guard against thinking, as is not uncommon, that consequentialist theories and various neo-Kantian alternatives exhaust the theoretical options we have.

Third, it is often said today that utilitarianism does not take seriously the “separateness of persons,” a charge usually contested by utilitarians. The idea, in part, is that to attribute to morality a global-maximizing structure is *ipso facto* to minimize the importance of or to accord derivative significance to the separate points of view of individuals. Put differently, to sum desire satisfactions or utilities across persons, as global-maximizing theories do, is to place less significance on the fact that these utilities are those of particular and distinct individuals. Of course, the formal requirement of utilitarianism, that each is to count for one and no one for more than one, ensures that no individual is overlooked in determining the greatest total good; but the good of any particular individual is summed with the goods of others to form the greatest total good, and the distinctness of these individuals ceases to be of fundamental moral concern.

A number of contemporary philosophers regard this failure to accord primary weight to the distinctiveness of persons and personal perspectives to be a central flaw of utilitarianism. They believe that no moral theory can be adequate if it fails to take seriously the different perspectives of separate agents. One can, in part, think of this issue as having to do with the identity of the appropriate “objects” of moral theory; utilitarians understand well-being or the good as the object of primary moral concern, whereas others, such as neo-Kantians and many contemporary natural rights theorists, take individuals to be the objects of moral theory. The central issue here, then, to use Sidgwick’s terms, is between moral socialism and moral individualism.

It is important to guard against taking talk of agent-neutral value to refer to value in some agent-independent sense. Thus, when utilitarians claim the greatest total good to be an intrinsic, that is, a noninstrumental value, their critics have sometimes been led to speculate about whether there are other values, agent-independent values, that would persist even in a world of no valuers. But this is not what utilitarians take intrinsic value to be. Whatever is held by them to be valuable in and of itself in no way severs all connections with agents or sentient beings. The greatest total good is not valuable in virtue of its fostering or fa-

cilitating the realization of some other value; that is why it is said to be of “intrinsic,” or better, of “ultimate,” value. But that is not to say that it is an agent-independent value. Agent-neutral value refers to the fact that something is, necessarily, valuable to all agents, and that is not the same as claiming that something is valuable independently of any reference to agents at all. We should distinguish between two senses of the term ‘intrinsic’: Some value might be intrinsic in the sense of being “ultimate” or noninstrumental, and some value might be intrinsic in the further sense of being inherent in the world, a valuer-independent property of objects. These different senses are often conflated.

Some environmental ethicists regard the value of a functioning biotic community as intrinsic in both senses and even think that this value would persist or remain even if agents or valuers gradually ceased to be part of the picture. Neither neo-Humean nor utilitarian positions – nor, indeed, the positions of Foot or Nagel – hold much comfort for these environmentalists. This said, however, there may be a perfectly straightforward way of accommodating a portion of the Green movement without severing the link between intrinsic value and agents. We can speak of the “non-use” or “existence” value of something, which refers to people’s preferences that something not be used (or consumed, spoiled, polluted) but continue to exist (as it is). Just as people value the existence of the spotted owl, so they value the existence of a particularly beautiful part of the countryside or, more generally, of unspoiled forests or unpolluted rivers. In this way, the “intrinsic” value of inanimate nature may be identified with non-use or existence value. But this value makes direct reference to people’s preferences. It need not be value independent of all valuers; the part of the Green movement that seeks value of this sort in nature may find little comfort from contemporary value theory.

Neo-Humeanism and utilitarianism, then, are traditions that to a great extent dominated discussions of moral value and so ethical theory during the last two centuries, at least in the Anglo-American world. Even today, when numerous theorists depart from these traditions and when it is probably fair to say that no tradition is dominant, neo-Humeanism and utilitarianism nevertheless form the background against which those new discussions take place. Indeed, in the case of utilitarianism, it remains in all its versions the object of attack, even as its opponents attempt to spell out new alternatives, whether with respect to rightness and the adequacy of consequentialisms of different stripes or with respect to goodness and the adequacy of utilitarian value theory.

Interestingly enough, too, in view of the recent surge of interest in

substantive ethics, metaethical considerations and considerations that are part of the epistemology of morals have resurrected themselves, in some forms similar to what they were in the postwar Anglo-American world and in other forms different. The matter of the “objectivity” of moral judgments is of considerable interest to many philosophers today. We might ask whether moral judgments are capable of being true (or false). Some “expressivists” and “projectivists” deny that they can have truth-values. Others disagree. Among the latter there is an additional debate as to whether the truth of moral judgments is independent of their conditions of verification.

Traditionally, it used to be thought that metaethics is, or ought to be, neutral between competing normative ethical theories. But an important point of contention between expressivists and their opponents is whether the former’s analyses can in fact account for the content of our moral judgments. Can emotivist or expressivist analyses, for instance, recognize the agent-neutral or even agent-independent content of some moral judgments? The resolution of these disputes may have considerable importance for the questions we have raised about the relations between value, welfare, and morality.

The essays that follow take up a number of the above issues and numerous others as well. They not only give, collectively, an overview of the sorts of controversies sketched above, but they each seek as well to address and resolve some of those controversies. They aim, thus, to advance discussion of acutely significant issues having to do with the connections among value, welfare, and morality.

In defense of a “projective” account of the content of moral judgments, Simon Blackburn, in “The Land of Lost Content,” enumerates the ways projectivism provides a better story than its rivals. Its naturalism makes metaphysical sense, he argues, as well as explaining how it is that such judgments can be known and how skepticism can arise concerning them. Additionally, sense must be made of the motivational power of ethical judgments, and this, Blackburn thinks, is a central strength of projectivism.

In “Putting Rationality in its Place,” Warren Quinn raises an objection against a type of subjectivism that is influential in recent moral philosophy, what he calls “potential noncognitivism.” It is the view, identified with J. L. Mackie, that morality, conceived as involving genuine beliefs, is defective and must be reconceived along expressionist lines. What Quinn wishes to attack is the part of this subjectivism that understands an individual’s moral judgments as providing reasons; indeed, subjectivists think that such judgments provide reasons only because of the

noncognitive attitudes that they express. Quinn questions this claim, and he sketches an alternative view that reasons are provided by the good that acts realize or the bad that they avoid.

“The moderate Humean view,” as John Broome characterizes it in “Can a Humean Be Moderate?”, is the position that while reason does not constrain the ends one may hold, it does require that a person’s preferences be consistent in certain ways. These “moderate” requirements are the subject matter of decision theory, as many if not most practitioners view the field. A more radical Humean view is more extreme: It is the view that no preference can be irrational. Such a position does not even rule out someone taking inefficient means to achieve his ends. Broome argues that the moderate view is untenable, leaving only the extreme view to defend the field.

L. W. Sumner, in “Welfare, Preference, and Rationality,” wishes to undermine both the received view of welfare as consisting in the satisfaction of an individual’s preferences and the received view of rationality as consisting in the individual’s maximization of her preferences. Taken together, the two views entail the view that rationality requires the maximization of one’s own welfare, a position that he takes as absurd.

Drawing analogies between the role of preference in contemporary moral philosophy and that of perception in classical empiricism, Arthur Ripstein in “Preference” also attacks the uses to which the concept of preference has been put by recent philosophers. Extending antiempiricist arguments to preference-based moral theories, Ripstein argues that either they provide no account of practical reason or the good, or they presuppose an independent account of reason or good.

Without denying that one’s preferences often provide one with reasons, David Copp argues in “Reason and Needs” that meeting one’s basic needs is central to achieving the good from one’s own standpoint and that needs and their satisfaction are independent of preference. Copp defends three theses about needs: that a person’s needs are not determined by his or her preferences, that someone’s needs provide him or her with reasons, and that rationality does not require one to do something that would prevent one from meeting one’s basic needs.

A number of contemporary philosophers, including Copp, have appealed to second-order desire or preferences – desires or preferences about (first-order) desires or preferences – to explain or to construct basic notions. David Lewis, for instance, has recently argued for identifying valuing with a kind of second-order desire. Gilbert Harman is critical of all such attempts and argues against several recent versions in “Desired Desires.” He thinks that valuing something involves desiring