

PETER RAILTON

REPLY TO BEN EGGLESTON

1. INTRODUCTION

Ben Eggleston's comment nicely weaves together several strands of the approach to normative ethics that informs the essays in Part II of *Facts, Values, and Norms*. I quite agree with what he says concerning how and where the strands might reinforce one another, so I will concentrate in my reply on the places he finds strains and gaps – especially his questions concerning the nature of valoric consequentialism and how it might differ from more familiar forms of consequentialism.

As mentioned in the *Précis* above, the chief aim of valoric consequentialism is to translate as directly and faithfully as possible into the domain of practical ethics what I think of as the guiding consequentialist idea: the touchstone in moral assessment is the effect of what we do and how we are on the good of all concerned. I take “good” in the broad sense that extends beyond welfare, and includes all the forms of intrinsic value realized in and through our lives. The basic evaluative category of valoric consequentialism is *fortunateness*. Roughly, the fortunateness of x – where x ranges over acts, motives, experiences, practices, act-tendencies, etc. – is a function of how much x contributes to the overall good of those affected; x is more fortunate than y , other things equal, to the extent that x contributes more than y to this good.

2. ACT UTILITARIANISM, RULE UTILITARIANISM, MOTIVE UTILITARIANISM, AND COOPERATIVE UTILITARIANISM

Straight off we see differences between valoric consequentialism, on the one hand, and act or rule utilitarianism, on the

other. The latter are theories of *moral rightness*, an essentially deontic, binary notion, as applied to individual actions.¹ Moral fortunateness is an evaluative notion, a matter of degree, and directly applicable to a wide range of moral phenomena.

Does this difference in the entry-point of consequentialist assessment into practical ethics matter? Consider: By an act-utilitarian standard, all suboptimal acts are morally wrong alike. By a rule-utilitarian standard, all acts out of compliance with an acceptance-optimal code are wrong alike. But from a valoric consequentialist perspective, a near-optimal act may be almost indistinguishable morally from an optimal one, and crucially different morally from an act far from optimality. Such differences in magnitude become especially salient when we consider an agent's acts over time. Looking through the aperture of right and wrong, we see in the evaluative foreground the relative frequency or proportion of right acts. Yet imagine two well-disposed agents who face a similar series of choices, one of whom devotes equal energy to detail in almost every decision while the other is more selective in her attention. Over a span of time, the first, let us say, acts optimally with high consistency, but loses sight of the forest for the trees and acts quite non-optimally in a few very consequential decisions. The second acts sub-optimally much more frequently, but by focusing attention and effort on the larger choices, acts in these cases much closer to optimally. From the standpoint of consequentialism's guiding idea, mere frequency of right action seems to have little significance compared with overall advancement of the good.

Similar differences arise when we consider the evaluation of practical dispositions and motives. Are the best dispositions from a consequentialist standpoint those that lead most reliably to right action? The example above suggests otherwise—selective attention might be a more fortunate practical attitude in general, despite a higher frequency of right actions. Moreover, a good deal of the value contributed by attitudes and motives to the value of an agent's life and the lives of those around her is not a matter of action. Thus, in

his criticism of orthodox utilitarianism, Bernard Williams challenges its “act-adequacy premiss” (1973, 120). And in his defense of “motive utilitarianism,” Robert Adams urges that ethicists should reject “the assumption that ‘What should I (try to) do?’ is *the* ethical question, and that we are engaged in substantive *ethical* thinking only insofar as we are considering *action-guiding* principles” (Adams, 1976, 474). He offers *motive utilitarianism* not as an alternative theory of right action, but as an evaluative theory in its own right that applies the guiding consequentialist idea *directly* to motives. In effect, he adopts a valoric consequentialist approach to the ethics of motive, asking which motives are more fortunate than others – either in general, or in a specific setting – taking into account all the ways in which motivation can enrich or impoverish lives. Note that his question is not equivalent to asking which motives it would be optimal to *encourage* or which motives would arise from internalizing an *ideal code* – it is a matter of which motives it is most fortunate to *have*, either in general or in specific circumstances.

A satisfactory consequentialist ethics of action should expand its evaluative focus yet further beyond the rightness of individual acts, to incorporate the potential value of joint action involving multiple agents. As Donald Regan (1980, chs. 2–4) shows, all of the agents in a group can *individually* satisfy an act- or rule-utilitarian standard of right action, and yet fail to coordinate with one another, even when coordination would be much more fortunate overall.² Like Adams, Regan locates the flaw of act and rule utilitarianism in their “exclusive act-orientation” (see Regan, 1980, 114). To have the *adaptability* required to cooperate beneficially in an open-ended class of contexts, agents must not simply act rightly, given their situation, but *be* the way required to create situations: actively disposed to read situations in terms of the opportunities they afford for beneficial coordination, to imagine how the situation looks from the standpoint of others and thereby envisage which points of coordination or strategies are salient to them, and to trust one another enough to initiate unsecured cooperative behavior. Like successful communica-

tion in open-ended conversations, successful coordination in an open-ended class of contexts requires more than a collection of agents, each of whom applies a similar standard or follows the same rules, and an ethics of action should reflect this.

3. "MIXED THEORIES" AND SOPHISTICATED CONSEQUENTIALISM

The most straightforward way of expanding the scope of consequentialist practical ethics beyond orthodox theories of the rightness of individual acts is to develop a "mixed" normative theory, embracing several forms of consequentialism at once, neither of which is given evaluative primacy. Adams suggests what he calls "act-and-motive utilitarianism," according to which individuals should both avoid acts "gravely detrimental to utility" and avoid being "in a bad motivational state" (Adams, 1976, p. 477). For example, an act-and-motive utilitarian agent would likely be disposed to steer clear of a life-choking concern always to act for the best or of a coordination-defeating concern to require assurance from others before cooperating. At the same time, she would possess a sufficiently active conscience and critical awareness to realize that sometimes it is of cardinal importance to identify and do the most fortunate thing, or to require evidence of cooperativeness – even though this goes strongly against her motivational grain.

Sophisticated consequentialism (Railton, 1984) seeks to make room for just such "mixed" ways of thinking about how to live. A sophisticated act consequentialist, for example, can approve of agents who possess a mixture of commitments, loyalties, motives, and spontaneous cooperative dispositions, since these are a necessary part of the most valuable and contributory lives, even though they will also sometimes lead agents not to attend to utilitarian considerations or to act wrongly by act-utilitarian standards. Just which commitments, loyalties, motives, and dispositions to have, with degree of strength and deliberative autonomy, and with what place for explicitly consequentialist reasoning or at what cost

to rightness in conduct, will of course depend upon what is most fortunate given an agent's actual context and capacities.

4. VALORIC CONSEQUENTIALISM

However, Eggleston is correct in his suspicion that I now believe the limitations of act utilitarianism *as a consequentialist theory of morally right action* call for more than sophistication. Moreover, Eggleston and I share the view that a more credible consequentialist theory of right action is not to be found in the form of rule utilitarianism, however ingeniously improved. Hence my attempt to introduce valoric consequentialism as an alternative. This brings me to Eggleston's two chief concerns, both of which arise when a valoric consequentialist mixed theory incorporates "a non-act-utilitarian standard of rightness."

(1) *If act- and rule-utilitarian accounts of morally right action are inadequate from a valoric perspective, what would a more adequate consequentialist theory of right action look like?* Eggleston insists that any consequentialist worth her salt must regard the act that would bring about the greatest value³ as "morally noteworthy," and the valoric consequentialist concurs, deeming it the *most fortunate* act. In moral and prudential deliberation alike, the most fortunate act has a strong claim to be the one we have most reason to perform. Optimization is not the only game in town, but it certainly enjoys widespread recognition as a compelling standard of rational choice in the prudential case, and act utilitarianism gains some of its credibility by deploying this standard from a moral point of view.

However, even though the optimal choice is in one obvious decision-theoretic sense *right*, for a choice to be *morally right* is a different matter. The concept of "moral rightness" comes to the consequentialist already well-embedded in our moral thought and practice. If "morally right" meant "optimal from a moral point of view," "morally wrong" would naturally be "less than optimal from a moral point of view." But moral wrongness goes with notions of blameworthiness,

condemnation, resentment, and guilt, and we do not typically dispense these for mere suboptimality. Moreover, our reactive attitudes characteristically attach not only to individual acts, but to act-types characterized in ways substantially independent of their consequences – e.g., breaking a promise, betraying a friend or spouse, etc. We would experience betraying a friend, even for his own good, as shameful – something understandably resented and calling for apology. By contrast, failure to benefit a friend *maximally* when one is already doing him a good turn does not feel shameful, and resentment on his part or a demand for apology would strike us as churlish.

In light of its embedded nature, moral rightness is most naturally treated by the valoric consequentialist indirectly, through some variant of *contextualism*. Relative to a moral context – which includes prevailing moral practices, the state of knowledge, and feasible (“nearby”) options – there are certain expectations any reasonable person might have of herself or others. Failure to live up to these expectations naturally brings down our reactive attitudes. Some expectations seem reasonable in almost any context – to take promises seriously, to be honest, to refrain from actively harming innocents, etc. Other expectations vary in strength or relevance by context – to assist others in need, to cooperate actively, to do one’s share. Right and wrong with regard to individual responsibility for the well-being of others will differ for a nomadic society vitally dependent on resource sharing, a prosperous, egalitarian, or functional welfare state, a highly-inegalitarian post-industrial society with decaying social services.⁴

Is there not a grand theory of moral rightness behind these judgments – in effect, an ideal moral code akin to that of rule utilitarianism? The rule utilitarian’s ideal code can be only as comprehensive and specific with regard to contexts as ordinary human capacities to learn and apply it permit, opening up the unfortunate possibility of “rule worship.”⁵ Contextualism about right and wrong, by contrast, is not a moral code but a theoretical framework for asking what is more or

less fortunate and feasible for agents across the wide range of variation in human thought, practice, and circumstance.⁶

(2) *Suppose we had in hand an adequate consequentialist theory of right action. If it is incorporated as part of a valoric “mixed” theory that also includes (say) a theory of apt motivation, does not practical deliberation require that one evaluative category have primacy in decision?* To someone holding a mixed theory of this kind, one and the same act can be both (1) the almost-inevitable upshot of the motives it is most fortunate for the agent to possess, and (2) morally wrong. Rather than collapse all dimensions of moral assessment into one, valoric consequentialism seeks to portray these diverse types of reasons and their respective degrees of severity accurately. *How much* more fortunate is the act than alternatives? *How grave* are the harms involved? How contrary would the act be to the agent’s motives and commitments? And so on.

But how are agents aware of these diverse magnitudes and dimensions to reach a univocal decision to act? If we say that fortunateness always has primacy, this risks marginalizing the notion of moral wrongness and atrophying associated reactive attitudes. On the other hand, if we say that rightness has primacy, regardless of how unfortunate the act might be, we court the risk of “rule worship.”

However, a dimension of evaluation need not be given lexical priority in order to be treated seriously. We try to balance competing demands of family and career without undue sacrifice of either – the health of each is important for the health of the other. Scientists make career-shaping decisions about which hypotheses to accept or pursue by making trade-offs among informativeness, fruitfulness, explanatoriness, and degree of confirmation. And even the strictest deontologists allow that justice must be tempered by mercy. The balances we strike in such cases seem to us far from arbitrary, whereas it can feel quite arbitrary to insist that one dimension always has primacy – as if we were closing our eyes to what is in plain sight.

Valoric consequentialism seeks to provide a general framework for moral assessment that recognizes the distinctive

place and role of various elements in our moral lives and self-understanding. An ideal rational choice theory would bring all dimensions of moral assessment into a single hierarchy or metric, which then would incorporate all the prudential, aesthetic, epistemic, and personal dimensions of our lives as well into one grand optimization problem. But absent this unlikely prospect, we face a choice: either insist on according primacy to one dimension of morality and of our lives, or accept a “mixed” theory that frankly recognizes the need for balanced judgment.

Faced with a choice between ϕ -ing or not ϕ -ing, we must decide what to do. Even so, the constituents of action – beliefs, desires, attention, emotions, effort, deliberation, cooperativeness, and commitment – come in degrees, as do our reactive attitudes and notions of responsibility. It is one of the lessons of moral dilemmas, I argued, that the manner in which one ϕ s can matter as much as whether one ϕ s. It is one thing if Sophie chooses with difficulty and regret, another thing if she does so without pang or sadness.

NOTES

¹ I am simplifying for the sake of contrast, and speaking in the objective mode for ease of exposition.

² For example, even the most sophisticated versions of rule utilitarianism (e.g., Hooker, 2000) cannot rule out the possibility that several ideal codes might be equal-best in light of all plausible desiderata (including the desideratum of closeness to extant moral views), so that two agents can follow acceptance-optimal codes to the letter without coordinating.

³ Or, to put things in the subjective mode, the act with the highest expected value.

⁴ John Rawls introduces a similar dependence of the principles of justice upon the state of development and mode of life of a society, e.g., in the transition from the general to the special conception of justice (Rawls, 1971, 42f).

⁵ The ideal code of rule utilitarianism does have a natural place within a valoric mixed theory, as part of a direct consequentialist theory of an ideally just society. For this purpose, stipulating widespread compliance (within the general constraints of human psychology) is appropriate.

⁶ Contextualism about moral right and wrong is not a form of moral relativism, since the underlying judgments of fortunateness themselves – and the theory of intrinsic good upon which they depend – are not relativized. For further discussion of relativism, see my reply to Justin D’Arms.

REFERENCES

- Adams, R.M. (1976): ‘Motive Utilitarianism’, *Journal of Philosophy* 73, 467–481.
- Hooker, B. (2000): *Ideal Code, Real World*, Oxford: Oxford University Press.
- Railton, P. (1984): ‘Moral Realism’, *Philosophical Review* 95, 163–207.
- Rawls, J. (1971): *A Theory of Justice*, Cambridge: Harvard University Press.
- Regan, D. (1980): *Utilitarianism and Cooperation*, Oxford: Oxford University Press.
- Williams, B. (1973): A Critique of Utilitarianism, in J.J.C. Smart and B. Williams, *Utilitarianism: For and Against*, Cambridge: Cambridge University Press.

Department of Philosophy
University of Michigan
435 South State Street
Ann Arbor, MI 48109-1003
USA
E-mail: prailton@umich.edu