

# CONSEQUENTIALISM

*Brad Hooker*

## The definition of consequentialism

A consequentialist theory evaluates things exclusively in terms of consequences. For example, beliefs could be evaluated in terms of the consequences of holding them (an approach often called pragmatism). The most common forms of consequentialism, however, focus on the evaluation of acts and sets of rules. And the rest of this chapter will stick to that focus.

The most familiar kind of consequentialism is the kind of utilitarianism maintaining that an act is morally right if and only if no alternative act has consequences containing greater *welfare* (or net benefit), impartially assessed. In utilitarians' *impartial* assessment of welfare, a benefit (i.e. addition to someone's welfare) or harm (loss to welfare) to any one individual gets the same weight as the same size benefit or harm to anyone else. Thus, benefits and harms to everyone count equally, no matter what his or her ethnic group, religion, wealth, education, political views, talent, or conscientiousness; all that matters is the size of the benefits and harms.

This kind of utilitarianism is *maximizing* in the sense that it calls for an act with unsurpassed consequences (anything less than the best isn't good enough, according to act-utilitarianism). And this theory is *direct* in the sense that the acts are assessed solely and directly by their own consequences (not, for example, by whether the acts are allowed by the rules with the best consequences). Thus the theory is known as maximizing act-utilitarianism.

While maximizing act-utilitarianism is the most familiar form of consequentialism, defining consequentialism so that this is the *only* form is a naive mistake. Many self-described consequentialists explicitly reject maximizing act-utilitarianism. Some consequentialists think *welfare* is not all that matters. Some reject *utilitarian impartiality*. Some reject a requirement to *maximize*. And some deny that acts are to be solely assessed *directly* by their consequences.

Given the disputes among consequentialists, consequentialism must not be characterized in such a way as to imply that all forms of it are welfarist, impartial in the way specified, maximizing, and direct. How should it be characterized? What all consequentialists about the morality of acts agree on is that, where

there are differences in the value of consequences, these are always, directly or indirectly, decisive in the moral evaluation of acts.

### What makes consequences better or worse?

While the most familiar form of consequentialism is maximizing act-utilitarianism, the oldest form is probably maximizing act-egoism. Maximizing act-egoism evaluates acts in terms of nothing but the consequences *for the agent*. Such egoism is, in terminology made popular by Derek Parfit (1984:143) and Thomas Nagel (1986: 152–3), *agent-relative* in the special sense that *the value of consequences depends on their relation (or connection) to the agent*. Maximizing act-egoism claims that the only consequences that matter are the ones affecting the agent's good.

Just as egoistic evaluation of consequences is agent-relative, so is purely altruistic evaluation. Imagine Jack and Jill each evaluate consequences *only* in terms of what is most beneficial (or least harmful) overall to everyone else. In other words, Jack cares about the consequences for everyone else but not for himself, and Jill cares about the consequences for everyone else but not for herself. Their evaluations of consequences will sometimes diverge, since sometimes what will be best overall for everyone except Jack will be different from what will be best overall for everyone except Jill. (Such divergence is most obvious, of course, when what is most beneficial for Jack is not what is most beneficial for Jill and no one else's welfare is affected.)

Pure egoism seems to be very rare. Nearly everyone would accept that other people's good matters morally *to at least some extent*. Pure altruism might be even rarer.

Indeed, many people accept that each agent is (at least often if not always) allowed to attach greater importance to consequences for himself or herself than to consequences for others. On such a view, while everyone's good matters, the agent's evaluation of consequences may legitimately be somewhat biased in a self-serving direction. Such a view is certainly not *purely* egoistic, but it still contains a large agent-relative component, since this view allows the value of consequences to depend on their connection to the agent. Another view with a large agent-relative component is the view that possible benefits to the agent should be given *some* weight in the agent's practical thinking but need not be given as much weight as the same size benefits to others. This view is altruistic but not purely so.

The most common forms of altruism involve special concern for those with whom one has special connections. Consider a mother who attaches more value to benefits for her own child than she does to the same size benefits for any child that is not hers. The special concern, which again is agent-relative in Nagel's and Parfit's terminology, comes in the priority the mother gives to benefits for *her* child. Many people have such agent-relative concerns for their family, their friends, their colleagues, and members of their community or country.

We have seen that one kind of agent-relativity involves a bias concerning the agent's own good, and another kind comes in special concern for others with special connections with oneself, such as one's family and friends. Yet another kind of agent-relativity focuses on the connection between the agents and their own actions, as the examples below illustrate.

Suppose Jill is in an awkward situation where the only way she can prevent three other people from telling lies is to tell a lie herself. Suppose Stephanie can prevent four other people from stealing only by stealing something herself. Suppose Rae can prevent five other people from killing innocent others only by killing an innocent person herself. Each of Jill, Stephanie, and Rae thinks that her primary duty is not to commit acts of lying, or stealing, or killing, rather than to minimize the number of acts of lying, stealing, or killing in general. Each of them thinks of herself as having agent-relative duties focused on acts of her own doing.

The opposite of agent-relativity (in Nagel's and Parfit's terminology) is *agent-neutrality*. Agent-neutral evaluation of consequences is not biased towards (or against) benefits to the agent or towards benefits to individuals with special connections to the agent. And agent-neutral evaluation of consequences is not biased towards (or against) acts with special connections with the agent (in particular, acts of the agent's doing). Therefore, agent-neutral evaluation should, in principle, be the same for everyone.

The question of what makes consequences better or worse cannot be answered without determining whether the consequences are to be evaluated agent-neutrally or agent-relatively. Most discussions of consequentialism assume that consequences are to be evaluated agent-neutrally. But influential agent-relative forms of consequentialism have been advanced.

On virtually every form of consequentialism yet advocated, at least a large part of what makes consequences better or worse is how much welfare, or net benefit, they contain. But what constitutes net benefit? There are three main views.

Hedonists hold that net benefit consists in pleasure minus pain. Pleasures and pains, or at least the ones that constitute additions to or reductions in welfare, are introspectively discernible and either attractive or aversive to the person experiencing them. Suppose Jack's life project turns out to be a failure but he never finds this out. Furthermore, suppose his pleasures are not reduced in some indirect way by this failure. Hedonists hold that failure of Jack's life project does not reduce his welfare. This is because hedonists think that one's welfare is determined solely by how one's life feels from the inside, and that this depends on whether one believes one's desires have been fulfilled, not on whether they really have been fulfilled.

Another main view of welfare holds that a person's welfare is constituted by the fulfillment of his or her desires, whether or not the person knows the desires have been fulfilled. This view is often called the desire-fulfillment (or preference-satisfaction) theory of welfare.

The main argument in favor of the desire-fulfillment theory over hedonism is that many people's self-interested concern extends beyond their own pleasures and pains, enjoyments and frustrations. For example, many people have stronger self-interested concern for knowing the truth (especially about whether their other desires are fulfilled) than for blissful ignorance. The main argument against the desire-fulfillment theory is that some desires are so whacky that their fulfillment would not itself constitute a benefit for the people who have them (though whatever associated pleasure these people derived *would* constitute a benefit for them). Imagine someone who wants a saucer of mud for its own sake (Anscombe 1957: 70), or to count all the blades of grass in the lawns along a street (Rawls 1971: 432), or to turn on as many radios as possible (Quinn 1993: 236).

A third theory of welfare holds that hedonism is right to hold that pleasure constitutes a benefit but wrong to hold that pleasure is the only thing to do so. This theory proposes that knowledge of important matters, friendship, and significant achievement (and perhaps other things) also constitute benefits. On this theory, a life contains more welfare to the extent that it contains pleasure, knowledge of important matters, friendship, significant achievement, and perhaps some other things. A life where desires were fulfilled but were not for these things would contain little welfare, according to this theory, which Parfit dubbed the objective list theory, but is sometimes called the list theory (Parfit 1984: 493–502; Crisp 1997: Ch. 3).

As well as disagreements about which theory of welfare is best, there are disagreements about whether the amount of overall welfare is all that matters, or whether the pattern of distribution matters as well. Utilitarians hold that consequences are to be assessed only in terms of overall welfare. Many other consequentialists think that an evaluation of consequences should take into account not only how much overall welfare obtains but also its distribution.

Some consequentialists think that the fact that a distribution of welfare is more equal than all others is always a consideration in that distribution's favor (Temkin 1993). Call this view *equality for its own sake*. A prominent objection to equality for its own sake is that there seems nothing attractive about equality where it can be achieved only by "leveling down" better-off individuals to the level of worse-off individuals (Raz 1986: 227; Parfit 1997). To preserve much of the spirit behind equality for its own sake without inviting the objection about leveling down, many consequentialists have moved to the view that, while equality of welfare does not always have something in its favor, benefits to the worse-off always matter somewhat more than the same size benefits to the better-off (Parfit 1997). This view is called *prioritarianism*.

A third prominent view about the distribution of welfare is egalitarian and prioritarian only up to a point. This view holds that the pressure to equalize welfare or choose benefits for the worse-off stops once everyone is above some threshold of welfare. If absolutely everyone has very high levels of welfare, there need be nothing objectionable about some people's having higher levels

than others have, according to this view (Skorupski 1992; Miller 1992; Crisp 2006: Ch. 6; cf. Casal 2007). Call those who hold this threshold view *sufficientarians*. Of course there is vagueness, debate, and uncertainty about where the threshold of sufficiency is.

Finally, there is the view that what really matters is not whether consequences increase equality of welfare, not whether consequences bring larger benefits to the worse-off, and not whether consequences would push people up to some sufficiency level of welfare, but instead whether people get what they *deserve* (Feldman 1997: 158–70, 203; Kagan 1999). On this view, undeserved inequalities of welfare are bad, but deserved inequalities of welfare are good.

Thus, there are disagreements among consequentialists about all of the following: (1) whether consequences are to be evaluated in an agent-neutral or agent-relative way, (2) whether welfare is constituted just by net pleasure, or by desire-fulfillment, or by some objective list of items, and (3) whether not only maximum overall welfare but also one or another pattern of distribution guide the assessment of consequences.

### **Act-consequentialism: maximizing vs. satisficing vs. scalar**

Although the most familiar kind of act-consequentialism requires the agent to maximize, much attention has recently been directed at forms of act-consequentialism that do not require maximization. Part of the explanation of the retreat from *maximizing* act-consequentialism is that the most familiar agent-neutral versions of maximizing act-consequentialism seem excessively demanding. The most familiar agent-neutral versions of maximizing act-consequentialism are maximizing act-utilitarianism and versions that conjoin concern for utility with concern either for equality, or for the plight of the worst-off, or for getting everyone up to some level of sufficiency. Each such theory calls on each relatively well-off individual to make huge sacrifices for the badly-off.

There are over a billion people in the world who are very badly-off. As long as there are efficient aid agencies with access to the badly-off, a relatively well-off person could make a small contribution that would be of huge benefit to one of the badly-off. Suppose a relatively well-off person contributes \$5 to Oxfam, and this saves someone's life. Well, this relatively well-off person could contribute another \$5 and save another life. And again, and again, and again, and many times more. Indeed, the relatively well-off person would have to reduce herself to near poverty before further sacrifices from her would be as large as the benefits for others that her further contributions would produce, or before her sacrifices would no longer be required by egalitarian, prioritarian, or sufficientarian principles.

Such a requirement to come to the aid of strangers requires a huge reduction in the agent's own good, unless the agent is already quite badly-off. This requirement

seems excessive. Working assiduously to achieve an optimal outcome in agent-neutral terms will typically require the agent actively to sacrifice more, and more often, than it is reasonable to demand. (For some alternative explanations of what the demandingness objection is supposed to be, see Hooker 2009: §2.)

A related objection to maximizing act-consequentialism is that it leaves the agent vanishingly little moral freedom (Vallentyne 2006: 23–8). According to maximizing act-consequentialism, the set of permissible acts contains only those acts whose consequences are not less good than the consequences of any alternative act. In some situations, there will be two or more acts whose consequences are not less good than the consequences of any alternative act. In such situations, maximizing act-consequentialism leaves the agent morally free to decide among these two or more acts. In most situations, however, there will be only one act with consequences not less good than the consequences of any alternative. So, in most situations, maximizing act-consequentialism will restrict the agent's "choice" to one alternative. In this way, the theory is excessively restrictive.

Satisficing act-consequentialism does not require the agent to produce the best possible outcome. It instead requires the agent to choose acts with "good enough" consequences (Slote 1985, 1992). The main attractions of satisficing act-consequentialism are that it can avoid the charges of being excessively demanding and excessively restrictive. But in order to have these attractions, satisficing act-consequentialism must offer a criterion of "good enough" consequences that typically falls some distance short of "best available consequences."

One of the main lines of objection to satisficing act-consequentialism focuses on the obvious difficulty of finding a stable and non-arbitrary specification of "good enough." Do consequences need to be 50 percent as good as the best available? Or 75 percent as good? Or sometimes 50 percent and sometimes 75 percent, and sometimes 90 percent as good?

The second main line of objection to satisficing act-consequentialism is that this theory allows agents to do *less* than the best for others *even when doing the best for others would involve no greater sacrifice for the agent* (Mulgan 2001: 129–42). Why would morality allow agents to benefit others less when doing more for others would cost the agent not one bit more?

Some consequentialists have developed theories that were supposed to share satisficing act-consequentialism's attractions (satisficing act-consequentialism is neither too demanding nor too restrictive) without permitting agents to choose an act that would benefit others less though some other act would benefit others more and not involve any greater sacrifice from the agents. For example, Douglas Portmore has developed a consequentialist theory that mixes altruistic concern for others' good with agent-relative favoring of benefits for oneself (Portmore 2001, 2003, 2005, 2008). On such a theory, the agent is permitted to choose a lesser benefit for oneself over a greater benefit for someone else, as long as the difference is not too great.

Another kind of theory developed in reaction to the difficulties that maximizing and satisficing share is scalar act-consequentialism. This theory assesses acts in terms of the relative goodness or badness of their consequences but jettisons the categories of “morally required,” “morally optional,” “morally permissible,” and “morally wrong” (Slote 1985: Ch. 5; Railton 1988; Norcross 2006). Because scalar consequentialism jettisons those categories, it absolves itself of the responsibility to identify the boundary between “morally permissible” and “morally wrong” and (within the category of “morally permissible”) the boundary between “morally required” and “morally optional.”

Is it sensible for a theory of morality to jettison these categories? The concepts of moral guilt and blame, and the categories of morally guilty and morally blameworthy, are of absolutely central importance. And they are parasitic on the concept of moral wrongness. Normally, an agent is morally guilty and his act morally blameworthy only if what he did was morally wrong. So, by jettisoning the concept of moral wrongness, scalar act-consequentialism threatens the concepts of moral guilt and blameworthiness. There would have to be an overwhelmingly powerful argument in favor of this theory in order for us to accept the conceptual purges it dictates.

### Actual vs. expected value of consequences

So far, our discussion has presumed to compare the actual consequences of one act with the actual consequences of each alternative act. But this sort of “full-information,” God’s-eye point of view might seem quite irrelevant to the situation of normal agents, who are usually uncertain what the consequences of an available action would be. One way of dealing with such uncertainty is to think in terms of *expected values* of outcomes.

To calculate the expected value, multiply the value of each possible outcome times the probability of that outcome, and then add together these products. Table 37.1 gives a very simple example, where there are only two available

Table 37.1 Calculating expected values.

	Value of possible outcome	Probability of this possible outcome	Expected value of possible outcome	Expected value for option
Alternative A	20	.5	$20 \times .5 = 10$	$10 + 1 = 11$
	2	.5	$2 \times .5 = 1$	
Alternative B	4	.7	$4 \times .7 = 2.8$	$2.8 + 4.8 = 7.6$
	16	.3	$16 \times .3 = 4.8$	

alternatives to choose between, and each of them has only two possible outcomes.

This example is artificial in many important ways. In the example, there are only two possible alternatives. Each alternative has only two possible outcomes. The value of each possible outcome is quantifiable, and the probability of that possible outcome is known. Real life is rarely so simple. Still, understanding often starts with the simplest cases.

In a world of uncertainty, of course, how agents deal with probabilities and risks matters. Agents are blameworthy for taking unnecessary risks of very bad outcomes for others even if the bad outcomes did not in fact come about. In other words, blameworthiness is tied to expected values of choices.

Now wrongness and blameworthiness have close conceptual ties. Given this, wrongness seems to be conceptually tied to expected value rather than to actual value of choices. Admittedly, in many contexts the choice that would actually have had the best consequences is the one we *wish* we had made. Nevertheless, *moral* assessment of choice is more about expected value than about actual outcomes. Maximizing act-consequentialism can take on-board these ideas about moral assessment – by holding that an act is morally wrong if there is some alternative act whose consequences have greater expected value.

### Decision procedures

It would be a mistake, however, to think that maximizing act-consequentialism is the view that on every occasion an agent *should decide which act to do* by ascertaining which act has the greatest expected value. Trying to decide what to do on that basis is often not what has the highest expected value, for the following reasons:

- (a) People often lack information about the probable effects of their choices and, without such information, could not calculate expected value.
- (b) Where they lack this information, they also often lack the time needed to get the information.
- (c) Even if they had the information, calculating expected values is typically unpleasant and time-consuming and thus a cost in itself.
- (d) Human limitations and biases are such as to make people inaccurate calculators of the expected overall consequences, especially where self-interest interferes.
- (e) There would be a breakdown of trust in society if people knew that others, with all their human limitations and biases, were always making their moral decisions by trying to calculate expected values. For if people knew others were deciding in this way, they could not confidently predict that others would routinely behave in certain ways (e.g. not attack, not steal, not break their promises, not lie, etc.).



If maximizing act-consequentialism's recommended procedure for making day-to-day moral decisions is not to try to calculate expected values, what is it? Act-consequentialists say that the right procedure to follow when making decisions is to follow rules against attacking others, stealing, breaking promises, lying, and so on, unless following these rules is more or less certain to result in far worse consequences than breaking them would. Deciding what to do on the basis of these rules has greater expected value than deciding what to do by trying always to calculate expected values.

But now act-consequentialism is in a seemingly paradoxical position. On the one hand, the theory holds that an act is morally permissible if and only if there is no alternative act whose consequences have greater expected value. In its account of moral permissibility, act-consequentialism makes no reference to rules about killing, promise breaking, attending to welfare of others to whom one has special connections, etc. On the other hand, act-consequentialism tells agents to make their day-to-day moral decisions by following such rules. An agent who follows such rules might well feel confused when told that his act was nevertheless impermissible because some other act had a bit more expected value (cf. Parfit 1984: 31–40; Streumer 2003; Lang 2004).

### Rule-consequentialism

Rule-consequentialism tells agents to make moral decisions by following certain rules, and the theory ties moral permissibility to these rules. It holds that an act is morally permissible if it is permitted by rules selected for their consequences. So rule-consequentialism does not get itself into the bind of specifying a criterion of moral permissibility that can conflict with its injunction about how to make moral decisions.

Is rule-consequentialism to be formulated in terms of rules with the *best actual consequences* or in terms of rules whose *consequences have the greatest expected value*? For reasons much like those mentioned earlier, rule-consequentialism is best formulated in terms of expected value (Hooker 2000: 72–5).

If rules are to be selected by their expected value, is this expected *agent-neutral* value (such as welfare for everyone, equality, or getting everyone above some threshold of sufficiency) or *agent-relative* value (such as the agent's own welfare and the welfare of those specially connected to the agent)? Well, rule-egoism is a form of rule-consequentialism. But obviously rule-egoism eschews the very attractive idea that rules are to be assessed in terms of the benefits and harms *for everyone*, not merely for the agent or some subset of everyone. Thus, the only kinds of rule-consequentialism discussed in the rest of this chapter are ones that assess rules in terms of agent-neutral value. (However, as will be explained below, the content of rule-consequentialist duties will be mostly agent-relative.)

Sometimes rule-consequentialism has been formulated as holding that an act is morally permissible if general *conformity* with the rule has the greatest expected value. Arguably, however, formulating rule-consequentialism purely in terms of the expected value of conformity with rules pushes rule-consequentialism into “extensional equivalence” with act-consequentialism, which means that though the theories have different criteria of permissibility, they end up selecting exactly the same acts as permissible. Act-consequentialism holds that it is wrong to attack others, or steal, or break a promise, etc., *only* when such acts have less expected value than not doing them has. The objection is that rule-consequentialism would have to agree, because of the benefits of perfect conformity with rules forbidding such acts only when these acts do not maximize expected value.

Although conformity with rules is hugely important, the process of internalizing rules and their ongoing acceptance can have consequences in addition to compliance with them. For one thing, people’s knowing that Jack accepts certain rules might lead them to do certain acts or have various feelings, even though Jack never has an opportunity to comply with these rules. For another thing, suppose that, while compliance with rule *A* would have *slightly* greater expected value than compliance with rule *B*, the time, effort, and other costs involved in getting rule *A* internalized would be *much* greater than those involved in getting rule *B* internalized. These additional consequences should be counted in a rule-consequentialist assessment of possible rules. So most philosophers now accept that rule-consequentialism is better formulated in terms of acceptance or internalization than in terms of mere compliance.

Thus formulated, rule-consequentialism holds that an act is morally permissible if it is allowed by the rules whose general acceptance (including the costs of getting them accepted) has the greatest expected value. Usually, “general acceptance” is interpreted as “full acceptance by a large percentage of people.” Permissibility is determined by rules selected by the expected value of their acceptance by a collection of people, not merely acceptance by the individual. In this sense, rule-consequentialism is typically put forward as a “collective” rather than an “individual” form of consequentialism.

Rule-consequentialism needs to be formulated in terms of acceptance by a large percentage of people, not in terms of acceptance by every single person (though universal acceptance can be an ideal). The reason not to formulate it in terms of acceptance by every single person is that many moral problems simply would not exist if every single person fully accepted rules against attacking others, stealing, breaking promises, etc. For example, there would not need to be rules about permissible defense against attackers in a world where there were no attackers. Of course, rule-consequentialism would *prefer* for every single person to accept the best rules. But it had better gear its moral rules for a less ideal world.

Rule-consequentialism seems to accord well with widespread views about permissibility. Constraints on attacking others, stealing, breaking promises, lying, and so on can be justified by the fact that acceptance of such constraints by a large percentage of people is crucial for security and thus has

high expected value. Duties to be especially concerned about the welfare of those with whom one has special connections can also be justified by their high expected value, since human nature is such that a world without such special concerns is likely to be a miserable place (Hooker 2000: 136–41).

According to standard forms of rule-consequentialism, these constraints and duties of special concern have agent-neutral justification but agent-relative content. The constraint on attacking others, for example, is the duty not to attack others oneself, not the duty to minimize instances of attacks by agents generally. More obviously, the duty to be especially concerned about the welfare of those with whom one has special connections will require different agents to be concerned about different others.

Such foundationally agent-neutral rule-consequentialism also endorses a more general duty to come to the aid of others, because of the benefits of this duty's acceptance. Now, will the general duty about aid that rule-consequentialism endorses be excessively demanding? General compliance with a more demanding duty to aid has higher expected value than general compliance with a less demanding duty to aid. However, the time and energy and emotional costs in getting a more demanding duty to aid internalized by a large percentage of people will at some point outweigh the added benefits of compliance with the more demanding duty. For this reason, foundationally agent-neutral rule-consequentialism justifies a less demanding duty to aid than agent-neutral act-consequentialism does.

## Conclusion

There are forms of agent-relative act-consequentialism and of agent-neutral rule-consequentialism that accord much better with intuitive ideas about constraints, about duties of special concern, and about limits on the duty to aid than agent-neutral act-consequentialism can. But agent-relative act-consequentialism has no place for the attractive idea that moral assessment is foundationally impartial in an agent-neutral way. Of consequentialist theories, only foundationally agent-neutral rule-consequentialism manages to achieve the conjunction of (a) building this kind of impartiality into the foundational level of assessment and (b) justifying the constraints, duties of special concern, and a limit on the more general duty to aid that seem intuitively compelling.

*See also* Utilitarianism to Bentham (Chapter 13); John Stuart Mill (Chapter 16); Welfare (Chapter 54); Population ethics (Chapter 61).

## References

---

- Anscombe, E. (1957) *Intention*, Oxford: Blackwell.  
 Casal, P. (2007) "Why Sufficiency Is Not Enough," *Ethics* 117: 296–326.  
 Crisp, R. (1997) *Mill on Utilitarianism*, London: Routledge.

- (2006) *Reasons and the Good*, Oxford: Clarendon Press.
- Feldman, F. (1997) *Utilitarianism, Hedonism, and Desert: Essays in Moral Philosophy*, New York: Cambridge University Press.
- Hooker, B. (2000) *Ideal Code, Real World: A Rule-Consequentialist Theory of Morality*, Oxford: Clarendon Press.
- (2009) “The Demandingness Objection,” in Tim Chappell (ed.) *Moral Demandingness*, London: Palgrave, pp. 148–62.
- Kagan, S. (1999) “Equality and Desert,” in Louis Pojman and Owen McLeod (eds) *What Do We Deserve?*, New York: Oxford University Press, pp. 298–314.
- Lang, G. (2004) “A Dilemma for Objective Act-Consequentialism,” *Politics, Philosophy & Economics* 3: 221–39.
- Miller, D. (1992) “Distributive Justice: What the People Think,” *Ethics* 102: 555–93.
- Mulgan, T. (2001) *The Demands of Consequentialism*, Oxford: Clarendon Press.
- Nagel, T. (1986) *The View from Nowhere*, New York: Oxford University Press.
- Norcross, A. (2006) “Scalar Act-utilitarianism,” in Henry West (ed.) *Blackwell Guide to Mill’s Utilitarianism*, Boston, MA: Blackwell, pp. 217–32.
- Parfit, D. (1984) *Reasons and Persons*, Oxford: Clarendon Press.
- (1997) “Equality and Priority,” *Ratio* 10: 202–21.
- Portmore, D. (2001) “Can an Act-Consequentialist Theory Be Agent-Relative?,” *American Philosophical Quarterly* 38: 363–77.
- (2003) “Position-Relative Consequentialism, Agent-Centered Options, and Supererogation,” *Ethics* 113: 303–32.
- (2005) “Combining Teleological Ethics with Evaluator Relativism: A Promising Result,” *Pacific Philosophical Quarterly* 86: 95–113.
- (2008) “Dual-Ranking Act-Consequentialism,” *Philosophical Studies* 138: 409–27.
- Quinn, W. (1993) *Morality and Action*, New York: Cambridge University Press.
- Railton, P. (1988) “How Thinking about Character and Utilitarianism Might Lead to Rethinking the Character of Utilitarianism,” *Midwest Studies in Philosophy* 13: 398–416.
- Rawls, J. (1971) *A Theory of Justice*, Cambridge, MA: Harvard University Press.
- Raz, J. (1986) *The Morality of Freedom*, Oxford: Clarendon Press.
- Skorupski, J. (1992) “Value and Distribution,” in M. Hollis and W. Vossenkuhl (eds) *Moralische Entscheidung und rationale Wahl*, Scientia Nova, Munich: Oldenbourg, pp. 191–207.
- Slote, M. (1985) *Common-sense Morality and Consequentialism*, London: Routledge & Kegan Paul.
- (1992) *From Morality to Virtue*, New York: Oxford University Press.
- Streumer, B. (2003) “Can Consequentialism Cover Everything?,” *Utilitas* 15: 237–47.
- Temkin, L. (1993) *Inequality*, New York: Oxford University Press.
- Vallentyne, P. (2006) “Against Maximizing Act Consequentialism,” in James Dreier (ed.) *Contemporary Debates in Moral Theory*, Boston, MA: Blackwell, pp. 21–37.

### Further reading

---

- Kagan, Shelly (1998) *Normative Ethics*, Boulder, CO: Westview Press.
- McNaughton, David and Rawling, Piers (1991) “Agent-Relativity and the Doing–Happening Distinction,” *Philosophical Studies* 63: 167–85.
- Mulgan, Tim (2007) *Understanding Utilitarianism*, Stocksfield, UK: Acumen.
- Pettit, Philip (1994) “Consequentialism and Moral Psychology,” *International Journal of Philosophical Studies* 2: 1–17.
- Smart, J. J. C. and Williams, Bernard (1973) *Utilitarianism: For and Against*, Cambridge: Cambridge University Press.