

Practical Equilibrium: A Way of Deciding What to Think about Morality

Ben Eggleston

(PhilPapers version, uploaded in May 2019)

citation for published version:

Ben Eggleston, "Practical Equilibrium: A Way of Deciding What to Think about Morality," *Mind* vol. 119, no. 475 (July 2010), pp. 549–584.

DOI: 10.1093/mind/fzq036.

abstract:

Practical equilibrium, like reflective equilibrium, is a way of deciding what to think about morality. It shares with reflective equilibrium the general thesis that there is some way in which a moral theory must, in order to be acceptable, answer to one's moral intuitions, but it differs from reflective equilibrium in its specification of exactly how a moral theory must answer to one's intuitions. Whereas reflective equilibrium focuses on a theory's consistency with those intuitions, practical equilibrium also gives weight to a theory's approval of one's having those intuitions.

1. Introduction

A philosopher who wants to refute, embarrass, or otherwise discredit a moral theory is as likely as not to do the following: sketch a hypothetical situation in which an agent is faced with a decision to make, identify the option that the theory in question picks out as the right thing to do, and then point out—with apparent horror (but also, all too often, badly concealed glee)—that most people would quite confidently regard that option as very much the *wrong* thing to do. What gives attacks of this sort whatever bite they are felt to have is the assumption that a moral theory is flawed if its implications for particular cases conflict with the moral judgements that most people are inclined to make. This assumption, long influential in moral philosophy, has become especially entrenched since Rawls's *A Theory of*

Justice gave it, in the framework of reflective equilibrium, pride of place among ways of deciding what to think about morality.

Although reflective equilibrium is a fundamentally sound way of deciding what to think about morality in general and moral theories in particular, my aim in this paper is to describe and defend a variant of reflective equilibrium that I call ‘practical equilibrium’. As I mention below (in Sect. 7), others have effectively imagined and employed the idea of practical equilibrium, but they have not explicated and justified it as such. As its name suggests, practical equilibrium is a close cousin of reflective equilibrium. It shares with reflective equilibrium the general thesis that there is some way in which a moral theory must, in order to be acceptable, answer to our ordinary moral intuitions; in other words, it shares with reflective equilibrium the thesis that a moral theory must be capable of being in some sort of harmony, or *equilibrium*, with intuition. But it differs from reflective equilibrium in offering a broader, more flexible account of what enables a moral theory to be in harmony with intuition. In this paper, I briefly sketch reflective equilibrium and characterize practical equilibrium in contrast to it, and I argue that practical equilibrium improves on reflective equilibrium as a way of deciding what to think about morality in general and choosing among moral theories in particular.

2. What reflective equilibrium is

Reflective equilibrium is a way of deciding what to think about morality, with a focus on choosing among moral theories, that is based on the idea of consistency between theory and intuition—consistency, that is, between the moral theories being evaluated and a broad class of intuitions.¹ This is not to say, of course, that

¹ Defenders of reflective equilibrium may worry that by casually referring to moral ‘intuitions’, I am saddling reflective equilibrium with meta-ethical commitments that ought to be kept separate. But I disavow any such implication: on the contrary, I acknowledge that reflective equilibrium needs to presuppose little more than the existence of sincerely held moral judgements—these are the moral intuitions in question—and does not need to presuppose or imply any thesis about the epistemic status of those judgements or the faculty by which they are made. For discussion of the meta-ethical neutrality of reflective equilibrium by some of its defenders, see Shaw 1980 (pp.

reflective equilibrium dictates that we take all of our intuitions at face value, or that we place naïve and uncritical faith in every intuition that anyone wants to bring to the table. On the contrary, reflective equilibrium requires that we discard those intuitions that we doubt the reliability of. Only those judgements that survive such scrutiny are kept and attain the status of considered judgements (Rawls 1951, pp. 5–7; 1971, pp. 46–53; and Shaw 1980, p. 129); these are, in Rawls’s phrase, the ‘provisional fixed points’ to which a moral theory may appropriately be held to answer (Rawls 1971, p. 20).

In addition, reflective equilibrium allows us to fine-tune our intuitions in order to make them accord with a moral theory that, for whatever reason, we find compelling. One reason for this is that, as William Shaw writes,

We know that our considered judgments are subject to occasional irregularities, inconsistencies, and distortion. ... [So,] when presented with a set of principles which gives an appealing account of our sense of justice [or morality] we may wish to revise some of our judgments to conform with it. (1980, p. 129)

But this is not the only reason for the revisability of intuitions. Another is that it can be reasonable for one to revise them simply due to the appeal of a conflicting moral theory or a conflicting general moral principle. As Rawls writes (in the original but not the revised edition of *A Theory of Justice*),

Moral philosophy is Socratic: we may want to change our present considered judgments once their regulative principles are brought to light. And we may want to do this even though these principles are a perfect fit. A knowledge of these principles may suggest further reflections that lead us to revise our judgments. (1971, p. 49)

This means that in coming up with a theory that answers to our intuitions, reflective equilibrium tells us to ‘work from both ends’ (Rawls 1971, p. 20), considering ‘all possible descriptions [i.e. moral theories] to which one might plausibly conform

129–30), Daniels 1996a (p. 4), and Hooker 2000 (p. 15). I maintain this meta-ethical neutrality with regard to moral intuitions throughout this paper.

one's judgments together with all relevant philosophical theories for them' (Rawls 1971, p. 49). So what a theory needs to match are not our initial considered judgements, but rather those that we have after we reflect on various theories, arguments for and against them, and how our intuitions would have to be revised in order to be consistent with one or another of them.² These judgements, 'duly pruned and adjusted' (Rawls 1971, p. 20), are the ones that a moral theory must, according to reflective equilibrium, match.

T. M. Scanlon writes that reflective equilibrium is open to two interpretations. On the descriptive interpretation, reflective equilibrium is a way of 'characterizing the conception of justice [or morality] held by a certain person or group' (2003, p. 142)—it is a way of arriving at an accurate portrait of a held moral view (2003, p. 143). On the deliberative interpretation, reflective equilibrium is a way of 'figuring out what to believe about justice [or morality]' (2003, p. 142). Scanlon endorses the deliberative interpretation as primary (2003, p. 147 and p. 149), and that is the understanding of reflective equilibrium that is operative in this paper. For this paper is concerned with reflective equilibrium as a way of deciding what to think about morality, including deciding what moral theory to affirm.

Scanlon claims that reflective equilibrium 'is the best way of making up one's mind about moral matters and about many other subjects' and is, moreover, 'the only defensible method' (2003, p. 149). I agree that reflective equilibrium is fundamentally sound, and in this paper I urge the revision of just one aspect of it: its commitment to the thesis that a moral theory must, in order to be acceptable, ultimately be *consistent* with intuition. This thesis, which I call the consistency thesis, is not unique to reflective equilibrium; as a result, my critique of reflective equilibrium will implicitly be a critique of any other method of deciding what to think about morality generally, and about moral theories in particular, that embraces the consistency thesis but develops it differently.³ But although there are

² The revisability of intuitions is discussed at greater length by Rawls (1975, p. 289), Daniels (1979, pp. 26–8; cf. 1980a, p. 60, and 1980b, p. 71), Tersman (1993, p. 49), and Shaw (1999, p. 98).

³ The prevalence of the consistency thesis in moral theory evaluation is frequently noted in the literature. Frey writes, in a discussion of tests for moral theories, that 'One such test—

many possible consistency methods, as one may call them, I will focus on reflective equilibrium because of what I perceive to be its fundamental soundness and overall sophistication. Indeed, practical equilibrium is a variant of reflective equilibrium, and depends largely on the latter's merit for its own plausibility as a potentially superior alternative.

3. What practical equilibrium is

As a variant of reflective equilibrium, practical equilibrium is based on reflective equilibrium's fundamental insight that deciding what to think about morality is a matter of achieving an equilibrium (of some kind) among all of one's moral beliefs at all levels of generality. What practical equilibrium denies is the consistency thesis, which we have just seen to be an essential component of reflective equilibrium. Now, one way of denying the consistency thesis is to argue for an even stronger claim from which a denial of the consistency thesis follows: the claim that a moral theory need not answer to intuition *in any way at all* (whether in the way that the consistency thesis specifies or otherwise). On this view, when intuition and theory are in conflict, we should just say 'so much the worse for

historically, doubtless *the* test—is whether the results of the application of the theory are in accordance with our “ordinary moral convictions” (1977, p. 96). According to Shaw, ‘Pick up any recent journal or Moral Problems anthology, and it seems as if everyone is going about ethics in a similar way’ (1980, p. 127). Gibbard reports that ‘The method of counterexample ... is widely thought to provide decisive reasons for rejecting those ethical theories against which it can be brought to bear’ (1982, p. 77), and Sinnott-Armstrong claims that this is ‘The most common way to choose among moral theories’ (1992, p. 399). Tersman writes that reflective equilibrium ‘stands out today as the single most widely discussed ... candidate for a methodology in ethics. Sometimes, one gets the impression that it is the only candidate taken seriously in debates over moral epistemology’ (1993, p. 28). According to Brandt, ‘Reliance on such intuitions is extremely widespread among philosophers at the present time, often in the belief that there is no other base from which philosophical thinking about values and morality can start’ (1996, p. 5). Finally, Hooker claims that ‘most contemporary moral philosophers—no matter what their views on the metaphysics, epistemology, and language of morals—apply the same reflective-equilibrium methodology in normative ethics’ (2000, p. 14), and he cites others making this same point (2000, p. 14, n. 20).

intuition'.⁴ But such a view is too sweeping. As Ross claims in *The Right and the Good*,

to ask us to give up at the bidding of a theory our actual apprehension of what is right and what is wrong seems like asking people to repudiate their actual experience of beauty, at the bidding of a theory which says, 'only that which satisfies such and such conditions can be beautiful'. ... [T]he request is nothing less than absurd. (1930, p. 40)

Now, it is apparent from Ross's work that he also holds that nothing short of consistency between theory and intuition could avert the 'absurd[ity]' he mentions.⁵ But although I will soon argue that he and others are wrong to require consistency in this context, I do share the view that a moral theory must, in order to be acceptable, answer (in a way to be specified shortly) to our considered judgements about right and wrong. A moral theory that condemned our intuitions as thoroughly groundless and misguided would have no plausible claim to justification. So I join proponents of reflective equilibrium in maintaining that a moral theory must, in order to be acceptable, be capable of being in harmony with intuition; the only issue I wish to dispute is how this harmony should be characterized.

We saw above that according to reflective equilibrium (or any other way of deciding what to think about morality that is based on the consistency thesis), a moral theory must be consistent with an agent's most firmly held intuitions (which, for reflective equilibrium, are the ones that remain after the initial intuitions have been 'pruned and adjusted'). More precisely, a moral theory must (possibly in

⁴ This view is described in similar terms by Sprigge (1965, p. 270), Frey (1977, pp. 95–7), Daniels (1996a, pp. 3–4), and Shaw (1999, pp. 26–7 and pp. 86–8). It is advocated by Hare (1971a, p. 122; 1978, p. 27; 1981, pp. 11–12 and p. 40), Nielsen (1972, p. 229), Smart (1973, p. 68), Brandt (1979, p. 20), Gauthier (1986, p. 269), and Sencerz (1986, p. 77).

⁵ He writes that 'The main moral convictions of the plain man seem to me to be, not opinions which it is for philosophy to prove or disprove, but knowledge from the start' (1930, p. 20, n. 1). Later, he adds that 'I would maintain ... that what we are apt to describe as "what we think" about moral questions contains a considerable amount that we do not think but know, and that this forms the standard by reference to which the truth of any moral theory has to be tested, instead of itself having to be tested by reference to any theory' (1930, p. 40).

conjunction with certain empirical propositions about the usual consequences of actions of certain kinds, etc.) logically entail, or imply, those intuitions. (This is, of course, a stronger relationship than mere consistency—but it requires consistency.⁶) But according to practical equilibrium—and here is the essential idea of this way of deciding what to think about morality—entailing those intuitions is not the only way for a moral theory to achieve the requisite harmony; another way for a moral theory to achieve the requisite harmony is for it to approve of them (or recommend them, or endorse them) as intuitions that it is morally good for that person to have. Entailment is good, but so is approval.⁷

Because practical equilibrium follows reflective equilibrium in recognizing entailment as a way for a moral theory to be in harmony with intuition, it is not the case that practical equilibrium requires approval and gives no weight to entailment—rather, its slogan could be phrased as ‘either entailment or approval’.⁸

⁶ Since entailment is a stronger relationship than consistency, a moral theory that is consistent with all of one’s intuitions need not entail all of them—it could entail only some of them. According to reflective equilibrium, the more thoroughly a moral theory entails one’s intuitions, the better; but reflective equilibrium rightly allows that a moral theory can be acceptable even if it is merely consistent with many of one’s intuitions rather than entailing them. Lurking here are important questions about the most plausible possible formulation of reflective equilibrium, but I will set these questions aside since my claims about reflective equilibrium do not depend on the answers to them.

⁷ The consistency/entailment points discussed in the previous footnote are paralleled by the following points that arise in the context of practical equilibrium. A moral theory might not approve of all of the intuitions that it does not disapprove of—it might be neutral with respect to the moral desirability of some of them. Practical equilibrium regards approval of one’s intuitions as counting more strongly in favour of a theory than mere neutrality towards them, but it is most plausibly construed as allowing for some degree of mere neutrality, too. As was the case with reflective equilibrium, these points raise important questions regarding the most plausible possible formulation of practical equilibrium, but I will set them aside in order to focus on how practical equilibrium differs from reflective equilibrium.

⁸ Just as a theory can approve of an intuition without entailing it, a theory can entail an intuition without approving of it. It is also possible for a theory to both entail and approve of an intuition, or to neither entail nor approve of an intuition.

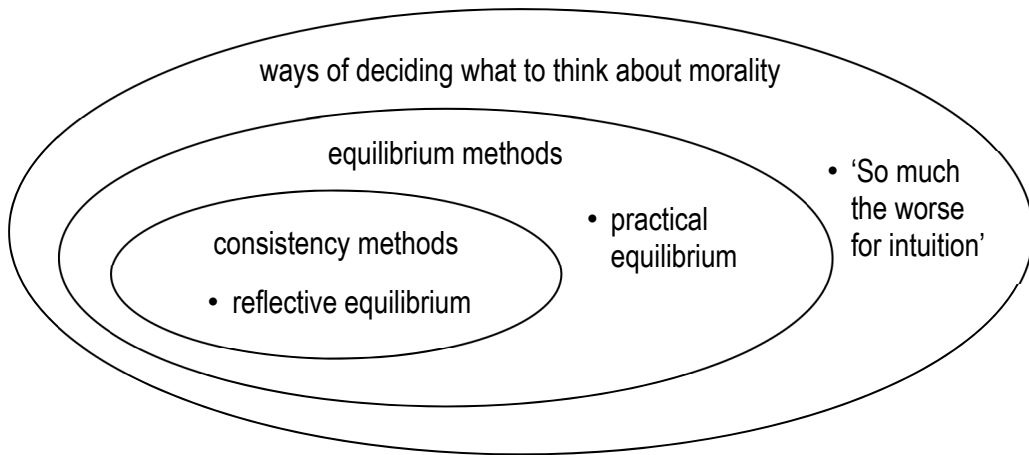
Thus, if an agent evaluating a moral theory would end up finding it to be in harmony with intuition under reflective equilibrium, she would end up finding it to be in harmony with intuition under practical equilibrium as well. Moreover, practical equilibrium does not require the agent to choose either entailment or approval as the overall mode by which a given moral theory might be in harmony with her intuitions—for each of her intuitions, either entailment or approval could be the mode by which the theory in question could be in harmony with that particular intuition. Thus, compared to reflective equilibrium, practical equilibrium is more permissive, in the sense that an agent evaluating moral theories using practical equilibrium would end up finding more of them to be in harmony with intuition than if she were using reflective equilibrium. (Then, as I discuss later in this section, each method would direct the agent to use other considerations—beyond harmony with intuition—in order to choose a single moral theory to affirm.)⁹

Both reflective equilibrium and practical equilibrium are perspectives of first-person deliberation: they are ways for a single person to decide what to think about morality (though, of course, multiple persons might engage in this activity collaboratively as they seek to make up their respective minds). For reflective equilibrium, the intuitions to be entailed by a theory being evaluated by an agent are the intuitions of the agent herself, not necessarily the intuitions that are most prevalent in her society: the entailment the agent looks for is the theory's entailment of her own intuitions, not the theory's entailment of the intuitions of any group. For practical equilibrium too, the intuitions to be either entailed or approved of by the theory are those of the agent herself. In looking for harmony between

⁹ Reflective equilibrium and practical equilibrium, as well as being understood as imposing requirements on moral theories (as just described), can also be understood as imposing requirements on agents. Specifically, they can be understood as establishing criteria for when an agent's affirmation of a given moral theory is reasonable. Reflective equilibrium maintains that if an agent affirms a given moral theory, her intuitions must (for the most part) be entailed by the theory. In contrast, practical equilibrium maintains that if an agent affirms a given moral theory, her intuitions must (again, for the most part) be either entailed or approved of by the theory. This is another aspect of practical equilibrium's more permissive conception of what it takes for a moral theory to be in harmony with intuition.

theory and intuition, the agent looks for both the theory’s entailment of her intuitions, in the manner just described, and the theory’s approval of her intuitions. And in looking for the theory’s approval of her intuitions, she is concerned not with the theory’s approval of her intuitions as general social norms, but with the theory’s approval of her intuitions as ones that it is morally good for *her* to have. Of course, the intuitions that a theory approves of for some particular agent might also be intuitions that the theory would approve of as general social norms—indeed this might usually be the case. And it will be convenient, below, to speak of ‘our’ intuitions, or the intuitions of some other group. But, throughout, when an agent is using practical equilibrium to decide what to think about morality, the intuitions that matter are her own.

This, then, is the general idea behind practical equilibrium, and I will try to show what difference it might make in the activity of deciding what to think about morality in a moment. But first I want to present the following diagram, which shows some things about the relationships between practical equilibrium and other ways of deciding what to think about morality.



The largest oval includes all ways of deciding what to think about morality, and the second-largest oval includes equilibrium methods, which are distinguished by the thesis that a moral theory must be in some sort of equilibrium with intuition, or must answer to intuition. An example of a non-equilibrium method is any that says that when theory and intuition clash, then ‘so much the worse for intuition’. The smallest oval includes consistency methods—methods requiring consistency between theory and intuition. Reflective equilibrium is a consistency method,

of course. In contrast, practical equilibrium, while following reflective equilibrium in being an equilibrium method, is not a consistency method, since it allows that the requisite equilibrium can be achieved not only through consistency between theory and intuition, but also through a theory's approval of one's intuitions.

So that is where practical equilibrium stands among ways of deciding what to think about morality. What practical equilibrium amounts to in the activity of moral theory evaluation can be seen particularly clearly in the context of utilitarianism (specifically, act utilitarianism, on which I will focus), since this theory is very commonly accused of failing the test of reflective equilibrium.¹⁰ By holding that an act is right if its consequences are at least as good, in terms of well-being, as those of any other act that could have been performed instead, utilitarianism implies that an act may be right even if it is an instance of lying, stealing, injustice, or any of many other bad things. Consider, for example, the implication of utilitarianism that outraged Anscombe so much: the implication that, in certain circumstances, it may be right to procure 'the judicial punishment of the innocent' (p. 19).¹¹ Intuitively, of

¹⁰ According to Hare, 'The commonest trick of the opponents of utilitarianism is to take examples of such thinking, usually addressed to fantastic cases, and confront them with what the ordinary man would think. It makes the utilitarian look like a moral monster' (1976, p. 222). In addition, Frey calls 'vast' the 'number of people who have objected that an untrammelled act-utilitarianism can and does produce morally shocking results' (1977, p. 99), and Gibbard writes that 'The method of counterexample has been used chiefly against act utilitarianism' (1982, p. 76). For example, Rescher writes of the 'substantial violence' that utilitarianism in at least one guise 'seems prepared to do' to 'elemental considerations of justice and common-sense morality' (1966, p. 48). And even Hooker, himself a consequentialist, rejects act consequentialism (of which act utilitarianism is one version) on reflective-equilibrium grounds (1990, p. 67; 1991, p. 269; 1994, pp. 313 and 314; 1995, p. 29; 1996, p. 538; and 2000, pp. 147-58).

¹¹ This has been an issue for utilitarians at least since the time of Godwin, who writes that 'I may put an innocent man to death for the common good, either because he is infected with a pestilential disease, or because some oracle has declared it essential to the public safety' (1793, p. 368). Although he denies that such activity would count as punishment (1793, p. 368), he writes of the infliction of suffering that 'An innocent man is the proper subject of it, if it tend to good' (1793, p. 370). He adds that 'A guilty man is the proper subject of it under no other point of view' (1793, p. 370).

course, we resist this claim, holding that punishing the innocent is never justified.¹² According to reflective equilibrium, this conflict counts against utilitarianism, because it is a case—or, rather (and more seriously), a fairly general class of cases—in which the theory fails to match intuition. When we ‘ask’ utilitarianism whether punishing the innocent is ever permissible, the theory gives the wrong answer.

But according to practical equilibrium, what it is important to ‘ask’ utilitarianism is not only whether our intuition against punishing the innocent is *correct*, but also whether it is good that we have this intuition. If the theory says that it *is* good that we have this intuition, then although the theory fails (for reasons mentioned above) to answer to intuition in the way specified by reflective equilibrium (or, thus, in the first of the two ways specified by practical equilibrium), it does answer to intuition in the second of the two ways specified by practical equilibrium.

So what does utilitarianism say about our intuition against punishing the innocent? Does it approve of it as one that it is good that we have? As a preliminary to answering this specific question, I want to address, briefly, the more general question of how it is even possible for a moral theory to say that it is good that we think, intuitively, that some kind of conduct is never justified, while itself implying that, sometimes, such conduct is justified. The key to seeing how this is possible is realizing that from the point of view of a moral theory, what someone *believes*, in so far as this is under her control, is fair game for moral appraisal, just as much as how someone *behaves*, in so far as *this* is under her control, is fair game for moral appraisal. From the point of view of utilitarianism, beliefs are to be judged in the same terms as acts: in terms of conduciveness to well-being. So what it is good that we believe is not necessarily what is true, but what it would be most useful for us to believe. No doubt the usefulness of a belief will sometimes depend on its truthfulness—as Mill said, ‘The truth of an opinion is part of its utility’ (1859, p. 233)—but to the extent that a belief’s truth and utility coincide, it will be only in

¹² But Ross, surprisingly, allows that ‘The interests of the society may sometimes be so deeply involved as to make it right to punish an innocent man “that the whole nation perish not”’ (1930, p. 61; cf. p. 64).

virtue of this latter property that it is good (according to utilitarianism) that it is held. Now it should be noted that utilitarianism does not—any more than any other theory that assesses beliefs instrumentally instead of in terms of their truth—presuppose that our beliefs are entirely under our control. It is perfectly compatible with Hume’s claim that ‘we can naturally no more change our own sentiments, than the motions of the heavens’ (1740, p. 517). All utilitarianism says, on this topic, is that *to the extent* that our beliefs are under our control, we ought to shape them in useful directions, not necessarily in accordance with the truth.

Having addressed the general question of how it is even possible for a moral theory to say that it is good that we think, intuitively, something that the theory itself implies is false, I want to return to the specific question of whether utilitarianism approves of our intuition that punishing the innocent is never justified, even while implying that this intuition is false. We can begin to answer this question by thinking about what our society would be like if we lacked a wholesale ban, both intuitively and legally, on punishing the innocent. If people believed, on the contrary, that punishing the innocent is sometimes justified, then our society would be a much less pleasant place to live: people would be ‘in constant fear of becoming sacrificial victims on the altar of utility’ (Scarre 1996, p. 105), and they would surely (and rightly) be afraid for their loved ones, too. And their fear would scarcely be lessened if they were to reflect on how rare the cases in which utilitarianism would actually require this are, because they would also have to worry about all the cases in which people in positions of power would *think*, however erroneously, that greater well-being would result from punishing the innocent. So the disadvantages in terms of well-being, if it were commonly believed that punishing the innocent is even occasionally justified, would be enormous.

Fortunately, a non-utilitarian view—the more restrictive one according to which punishing the innocent is never justified—is rather widely held, and so many people live largely free of the worries just described. And this is surely a greater benefit than that which is forgone when those few truly utility-maximizing occasions of punishing the innocent are passed over (particularly since there are, it seems likely, *very* few of them). Many utilitarians have argued for this conclusion. Timothy Sprigge, for example, claims that

such things as the punishment of an innocent man should present themselves to our thoughts as intrinsically bad. ... A society in which punishment of the innocent did not repel would be in all sorts of ways a wretched one. (1988, p. 223)

And R. M. Hare observes that

Our whole system of justice is founded on the premise that nobody is to be punished ... for offences that they have not ... committed. It would take an inconceivable shift in opinion to abandon this principle, and the consequences of its abandonment would be dire. (1997, p. 145)

So it seems very likely that it is more useful for it to be commonly believed that punishing the innocent is never justified than, as utilitarianism strictly implies, that punishing the innocent is sometimes justified.

So far I have been focusing on the issue of punishing the innocent, and how utilitarianism approves of the intuition that such conduct is never justified. Similar claims can be made in regard to the intuitions that utilitarianism approves of for many other issues, such as being truthful¹³ and respecting individual rights.¹⁴ Indeed, Sidgwick goes so far as to claim that 'Common-Sense morality' in general may be preferable to utilitarianism as a morality for people to intuitively embrace

¹³ Mill, for example, writes that 'the cultivation in ourselves of a sensitive feeling on the subject of veracity is one of the most useful, and the enfeeblement of that feeling one of the most hurtful, things to which our conduct can be directed' (1861, p. 223). See also Shaw 1999 (p. 89, p. 111, and p. 258).

¹⁴ Sprigge writes: 'there is every reason to institute some legal rights which put an absolute bar on sacrificing individuals in certain ways. And I believe we should go further and say also that the most desirable constructed moral world is one in which certain rights are thought of as absolute, or at least as well nigh absolute' (1988, p. 220; cf. 1989, p. 15 and p. 18). Scarre observes that 'A society that treated its citizens as expendable in the interest of the greater good would not be a setting for happy lives; it would be a breeding ground for insecure neurotics' (p. 168).

and live by.¹⁵ The upshot is that while we may have many firmly held intuitions that utilitarianism refuses to *entail*—intuitions not only about never punishing the innocent, but about a broad range of issues—we may have relatively few intuitions that utilitarianism refuses to *approve of*. In this way, while utilitarianism may fail to answer to intuition in the way required by reflective equilibrium, it may succeed in answering to intuition in the other of the two ways disjunctively required by practical equilibrium.

Practical equilibrium maintains that a theory can be in harmony with our intuitions in either of two ways: not only by entailing them, but also by approving of them. But not every theory that approves of our intuitions is thereby justified, according to practical equilibrium. Here, again, practical equilibrium borrows from reflective equilibrium. An agent who regards reflective equilibrium as the best way of deciding what to think about morality would be concerned not only with proposed theories' entailment of her intuitions, but also with other aspects of them, such as what arguments can be given for them, what moral ideals they stand for, what conceptions of human nature and human flourishing they espouse, what conceptions of personal and political relationships they espouse, the extent to which they possess traditional theoretical virtues such as simplicity and power, and so on. To distinguish practical equilibrium from reflective equilibrium, I have focused on the latter's specification of the harmony that needs to obtain between theory and intuition. But reflective equilibrium ultimately involves bringing all relevant considerations to bear on the process of evaluating moral theories and deciding what to think about morality in general.

The same is true of practical equilibrium: the agent should bring all relevant considerations to bear, and a certain degree of harmony between theory and intuition is only a necessary, not a sufficient, condition for theory justification. The only difference is that instead of interpreting the necessary harmony in terms of

¹⁵ He writes that although 'Common-Sense morality is really only adapted for ordinary men in ordinary circumstances ... it may still be expedient that these ordinary persons should regard it as absolutely and universally prescribed, since any other view of it may dangerously weaken its hold over their minds' (1907, p. 466).

entailment only, the agent would interpret it in terms of both entailment and approval. But such harmony need not be the agent's main reason for affirming a moral theory (assuming she ends up affirming one). To return to the example of utilitarianism, she may find that theory compelling primarily because of her commitment to certain conceptions of individualism, well-being, impartiality, and maximization. Still, she would recognize the necessity of some sort of harmony between utilitarianism and her intuitions. In looking for that harmony, she may find that utilitarianism is in harmony with many of her intuitions not in virtue of entailing them (for it might not), but in virtue of approving of them. And she may, finally, regard that kind of harmony between theory and intuition as perfectly satisfactory. The point of practical equilibrium, quite simply, is that this agent's way of deciding what to think about morality is entirely defensible, and (as argued below) better than reflective equilibrium's narrower perspective in which only entailment, and not approval, is credited with establishing harmony between theory and intuition.

The example I just gave involves an agent who finds utilitarianism compelling because of her commitment to certain ideals of individualism, well-being, impartiality, and maximization, and who wants to know whether she can reasonably regard utilitarianism as being in harmony with her intuitions. It might be objected that if an agent is more committed to principles and intuitions that are inconsistent with utilitarianism than she is to those ideals that make utilitarianism compelling to some people, then that agent is perfectly entitled to reject utilitarianism and subscribe to a moral theory that entails (or is at least consistent with) her intuitions. But such a decision is perfectly compatible with, and at home within, the perspective of practical equilibrium. For the point of practical equilibrium is not to say that an agent must subscribe to any moral theory that manages to approve of her intuitions, regardless of whether she finds it independently compelling. Moreover, the point of practical equilibrium is not to say that consistency between a theory and one's intuitions is irrelevant to one's rationally deciding what moral theory to affirm. Rather, the point of practical equilibrium is to say that if an agent does find some moral theory independently compelling, then she need not regard some degree of inconsistency between that

theory and her intuitions as preventing that theory from being in harmony with her intuitions. The requisite harmony can hold in virtue of the theory's approving of her intuitions, even if not in virtue of thoroughgoing consistency between the theory and her intuitions.

4. Not just for consequentialists

It might appear from the foregoing that only utilitarians and other consequentialists could have reason to be interested in practical equilibrium, and in the merits of it as a way of deciding what to think about morality. For it might appear that only a consequentialist theory could fare better under practical equilibrium than under reflective equilibrium. But all a theory needs to do in order to fare better under practical equilibrium than under reflective equilibrium is to approve of agents' having intuitions that are inconsistent with the intuitions that it logically entails. And this can be done not only by consequentialist theories, but by deontological ones as well.

Admittedly, it cannot be done by the most prominent deontological theory, that of Kant—or so, at least, I am resigned to conceding, due to the following line of reasoning. Kant says that the categorical imperative prohibits people from treating humanity, wherever we find it (whether in ourselves or others), exclusively as a means (1785, p. 4:429). If one's humanity is to be understood in terms of one's rationality (as maintained in, for example, Hill 1980), it seems clear that the categorical imperative would prohibit any agent from inculcating or maintaining in herself false beliefs—about the requirements of morality or anything else. Indeed, doing so would seem to be as bad as lying to someone else (which of course is one of Kant's principal examples of wrongdoing): for, in each case, one is corrupting someone's rationality. On this basis it seems compulsory to conclude that Kant's moral theory cannot approve of agents' having intuitions that are inconsistent with the intuitions that it logically entails. It would be advantageous to me to be able to refute this argument about Kant's moral theory, since I would like to be able to show that even Kantians might be interested in the merits of practical equilibrium as a

way of deciding what to think about morality. But I find the foregoing argument plausible, and will not contest it here.¹⁶

Having seen what feature of Kant's moral theory prevents it from faring better under practical equilibrium than under reflective equilibrium—the refusal to view agents' rationality as a means—we can try to imagine other, non-Kantian, deontological theories that eschew this position and, thus, that might well fare differently under these two ways of deciding what to think about morality. We can make progress imagining such a theory if, in addition to excluding the Kantian fixation on agents' rationality, we also envision moral principles that are especially complicated, such as principles of the form 'An act is impermissible if it is of type A, unless it is also of type B, in which case it is permissible, unless it is also of type C, in which case it is impermissible after all, unless it is also of type D, in which case it is permissible after all'. Envisioning complicated principles helps because the more complicated a theory's principles are, the more likely they are to be too complicated to be consistent with an agent's intuitions, which may well be relatively simple. If we make some further assumptions, as in the extended example given below, we end up with a deontological theory that would be more likely to be affirmed by an agent deciding what to think about morality using practical equilibrium than by an agent using reflective equilibrium.

To develop an example of this possibility, let us work with the view of Frances Kamm, a leading deontological theorist. In her aptly named *Intricate Ethics*, the key principles are complicated indeed—'very intricate', Kamm writes (2007, p. 4). For example, Kamm's definitive statement of her Principle of Permissible Harm is

¹⁶ A superficially tempting avenue of refutation is to cite those passages in which Kant encourages us to treat animals well so that we do not 'stifle the instinct of humaneness within us' (1793, p. 27:710), encourages us to avoid discontent (1785, p. 4:399; and 1797, p. 6:388), and encourages us to cultivate our sympathetic feelings (1797, p. 6:457). These encouragements might appear to reflect a willingness, on Kant's part, to view our deliberative capacities more instrumentally than the argument given in the text acknowledges. But for Kant these measures are a means of enlightening and informing rational deliberation, and are quite different from the kind of corruption of one's rational capacities that is involved in having false beliefs. I am grateful to Christine Korsgaard for directing me to these passages.

more than 300 words long (2007, p. 186, n. 78, and p. 188, n. 29). Now suppose that a deontological theorist whom we will call Sam regards Kamm's moral principles as correct, but also claims the following:

Kamm's principles, such as the Principle of Permissible Harm, are correct. But they are also much more complicated than our considered judgements. The resulting inconsistencies cause Kamm's theory to look unacceptable to most agents deciding what to think about morality using reflective equilibrium. But there is still hope. My theory incorporates the principles of Kamm's theory, but also states that it is a duty of agents to take measures that will help them avoid acting wrongly. These measures include, but are not limited to, cultivating in themselves whatever moral intuitions will, when employed by them in their everyday decision-making, tend to minimize the aggregate moral seriousness of their own wrongdoing. Call this the principle of training oneself for rightness, or TOR. Because of the complexity of principles such as the Principle of Permissible Harm, TOR requires agents to cultivate in themselves much simpler intuitions, such as adherence to the doctrine of double effect (the basic idea of which the Principle of Permissible Harm develops with extreme rigour—and all the complications that rigour requires). As it turns out, the simpler intuitions required by TOR are virtually identical to the considered judgements I mentioned above—the ones that are inconsistent with Kamm's principles. Of course, they are inconsistent with my principles, too, because mine are Kamm's plus TOR. So, my theory, like Kamm's, looks unacceptable to agents deciding what to think about morality using reflective equilibrium. But because my theory approves of the simple intuitions people have, agents who take the perspective of practical equilibrium view my theory much more favourably.

Sam, then, is an example of a deontological theorist who could have reason to be interested in practical equilibrium.

I deliberately put these claims into the mouth of a hypothetical deontological theorist not only because I do not affirm deontological principles, but also because I am not prepared to affirm all of Sam's other claims. But obviously I need to establish that Sam's claims do, at least, represent a reasonable deontological position. To that end, let me turn to some objections that might be raised against Sam's view.

First, it might be thought that Sam's view is not truly deontological. For it might be thought that for a truly deontological view, it is not simply outcomes that matter; rather, what matters is the agent's diligence in trying to do what the view says is right. Sam's view does not quite fit this mould. Rather than directing agents, on every occasion of action, to try to do exactly what it holds to be right, it directs agents to cultivate intuitions that will usually lead them to act rightly, but that might sometimes lead them to act wrongly. But does this prevent Sam's view from being deontological? There are at least two aspects of Sam's view that arguably warrant calling it a deontological one. First, Sam's view (following Kamm's view) holds that agents can be permitted and even required to perform acts that result in overall consequences that are not as good as the overall consequences that would result from some other available act. Second, Sam's view (specifically, its TOR) takes an *agent-centred* approach to the minimization of the aggregate moral seriousness of wrongdoing: it does not say that agents should cultivate in themselves whatever intuitions will minimize the aggregate moral seriousness of *all* agents' wrongdoing; rather, it says that agents should cultivate in themselves whatever intuitions will minimize the aggregate moral seriousness of their *own* wrongdoing. Regardless of whether these features of Sam's view ultimately warrant calling it a deontological one, I would submit that they make Sam's view dissimilar enough to utilitarianism and other standard forms of consequentialism to show that practical equilibrium has a wider application than might have previously been apparent.

Second, Kamm's own method of theorizing about morality relies heavily on considering 'as many case-based judgments of yours as prove necessary' (2007, p. 5). It might be thought, then, that the principles she arrives at are obviously going to be consistent with intuition. But it must be noted that the intuitions Kamm probes are not necessarily widely held ones; they are her own, and she acknowledges that her approach involves 'rely[ing] on intuitions even at great levels of complexity'

(2007, p. 5). In light of this, and in light of the complexity of the principles themselves, it would be entirely reasonable for a theorist such as our imagined Sam to maintain that although the principles are quite rightly *derived* from Kamm's careful consideration of *her* intuitions, it must be conceded that most people's intuitions are not so finely tuned. Indeed, Kamm herself acknowledges Thomas Nagel's remark that 'my deontological intuitions, at least, begin to fail above a certain level of complexity' (Nagel 1986, p. 180; quoted in Kamm 2007, p. 5). And one reviewer of Kamm's book writes that 'I suspect that many readers will find that their supply of strong intuitions cannot live up to the demands of Kamm's argument' (Woollard 2008, p. 232). So it should not be assumed that, because of the case-based way in which Kamm derives her principles, her theory will match most agents' intuitions. On the contrary, it is very reasonable for Sam to worry that most agents' intuitions might be oblivious to many of the fine distinctions drawn by Kamm's principles.

Third, it might be wondered how Sam could reasonably believe that agents' allegiance to Kamm's principles could result in more aggregate wrongdoing, by the lights of those very principles, than would result from agents' continued adherence to their simpler intuitions. But if even moral philosophers find Kamm to be operating at a dauntingly high level of complexity, it is surely reasonable for Sam to think that if people in general tried to comply with Kamm's principles, they would get so bogged down in the intricacies of them that their conduct would go awry more often than if they just stuck to their simpler intuitions. To be more precise, the aggregate moral seriousness of the wrongdoing that would result from misapplication of those principles might well exceed the aggregate moral seriousness of the wrongdoing that would result from the exclusion of certain intricacies from agents' intuitions. Of course, the best scenario of all, from Sam's point of view, might be for all agents to make themselves capable of flawlessly applying Kamm's principles. But just as utilitarians reasonably abjure such unrealistic scenarios in trying to ascertain what sorts of intuitions it would be best for agents to inculcate in themselves, so Sam is quite reasonable in concerning himself with the relative levels of wrongdoing that may realistically be expected to

result from agents' adherence to Kamm's principles versus their adherence to their simpler intuitions.

Finally, it might be pointed out that many deontologists will want to say that agents' simpler intuitions should be regarded, by Sam, as having only *pro tanto* status—as reflecting moral considerations that are valid, but that can be overridden by further considerations. This observation is entirely reasonable, and may well be an element of the most sensible understanding of a view such as the one I have attributed to Sam. For Sam's view is not that agents' simpler intuitions are truly correct moral principles; remember that his theory consists of Kamm's principles (in all their intricacy), plus TOR (the principle of training oneself for rightness). So in Sam's view, what agents' simpler intuitions have going for them is not that they tell the whole story about morality, but that they may well tell as much of the story about morality as it is morally desirable for agents to have in mind. This is what makes his theory more likely to be affirmed by an agent deciding what to think about morality using practical equilibrium than by an agent using reflective equilibrium.

This does not mean, of course, that Sam would be entitled to say that any agent whose intuitions were approved of by his theory necessarily ought, then, to affirm his theory. For as I said at the end of my discussion of utilitarianism in the previous section, the point of practical equilibrium is not that an agent must subscribe to any moral theory that manages to approve of her intuitions, regardless of whether she finds it independently compelling. Rather, the point of practical equilibrium is that if an agent does find some moral theory independently compelling, then she need not regard some degree of inconsistency between that theory and her intuitions as preventing that theory from being in harmony with her intuitions, since the requisite harmony can hold in virtue of the theory's approving of her intuitions. So although practical equilibrium will not license deontologists such as Sam to declare anyone to be required to affirm their theories (any more than it will license utilitarians to declare anyone to be required to affirm their theories), it will enable deontologists such as Sam (along with utilitarians) to rebut objections claiming that their theories fail to be in harmony with certain intuitions because of inconsistencies between their theories and those intuitions.

The foregoing provides at least a prima facie justification of the claim that practical equilibrium is not necessarily of interest only to proponents of consequentialist theories. Still, this section has been highly abstract, relying on a hypothetically advocated and thinly sketched moral theory. Thus, sceptics about the non-consequentialist relevance of practical equilibrium may remain somewhat dissatisfied. And it must be admitted that, in the present circumstances of contemporary moral philosophy, practical equilibrium will, indeed, be of interest primarily to proponents of consequentialist theories. Let me turn, then, from explaining the content and potential impact of my proposal to defending it as an improvement on reflective equilibrium.

5. Defending practical equilibrium, part 1: the psychology of the agent

Probably the greatest obstacle to the acceptance of practical equilibrium as a way of deciding what to think about morality is a cluster of concerns about the psychology of the agent who has used practical equilibrium to decide what to think about morality, and who has ended up affirming a moral theory that she would not have ended up affirming if she had used reflective equilibrium. I will address these concerns in this section, before offering (in the subsequent two sections) some affirmative considerations suggesting that practical equilibrium is an improvement on reflective equilibrium.

One natural concern is that practical equilibrium, with its emphasis on the practical value of relatively simple intuitions for most people, is to be thought of as a way of legitimating a moral theory which is to be employed by privileged rulers—say, Platonic guardians—in their supervision and manipulation of proles who cannot be trusted to employ the more sophisticated principles that the guardians themselves follow. For example, it might be thought that practical equilibrium is to be thought of as a way of legitimating what Bernard Williams calls ‘Government House utilitarianism’—‘an outlook favouring social arrangements under which a utilitarian élite controls a society in which the majority may not itself share those beliefs’ (Sen and Williams 1982a, p. 16; see also Williams 1973, pp. 138–40).

Obviously it would reflect badly on practical equilibrium if this were its impact on moral theorizing.

But as described above, practical equilibrium follows reflective equilibrium in being a perspective of first-person deliberation: a way for each person to decide, for herself, what to think about morality, including what moral theory (if any) to affirm. It is not a perspective of social planning, in which people are assigned to different roles and statuses based on their different attributes. Of course, there is nothing in practical equilibrium to prevent an agent from deciding that part of what she thinks about morality is that her society should be organized on the model of something like government-house utilitarianism. But the same is true of reflective equilibrium—in each case, it is up to the agent to have her own reasons for rejecting such a view. For an agent taking the perspective of practical equilibrium, these reasons may well include scepticism about whether the arguments for government-house utilitarianism are really good, as well as the observation that government-house utilitarianism would seem to require her to have a lot of intuitions that are very uncongenial to her. For example, she might be repulsed by the idea of being either a perpetrator or a victim of a scheme of large-scale manipulation. For such an agent, deciding what to think about morality in the manner of practical equilibrium rather than reflective equilibrium would bring her no closer to affirming anything like government-house utilitarianism.

This concern pertaining to schemes such as government-house utilitarianism is just one of the concerns about agents' psychology that practical equilibrium may excite. A further concern has to do with inconsistency in agents' moral beliefs. We have seen that according to practical equilibrium, a moral theory can be in harmony with an agent's intuitions even if it is inconsistent with them. As a result, an agent who follows the guidance of practical equilibrium in order to decide what to think about morality may end up with significant inconsistencies in her web of moral beliefs. On this basis, it might be tempting to regard practical equilibrium as an irrational way of deciding what to think about morality.¹⁷

¹⁷ As one would expect, the importance of consistency in moral thinking is confidently asserted almost everywhere that it is not simply taken for granted. Shaw writes that 'The testing

In response to this concern, one point to be made at the outset is that practical equilibrium is not the *source* of the inconsistency in the agent's web of beliefs. The inconsistency was there all along, though it may have been latent or unnoticed by the agent until she deliberately undertook the activity of deciding what to think about morality and discovered that she held certain ideals, or was amenable to certain arguments, that led her to affirm a moral theory whose implications were somewhat at odds with her intuitions. Practical equilibrium is not the reason she found herself being pulled in different directions by her own intuitions and other beliefs.

Still, we might ask more of a way of deciding what to think about morality: we might expect it to eradicate inconsistency, rather than merely refraining from causing or exacerbating it. One reason we might expect this is that it may seem plainly hypocritical for one to affirm a moral theory whose implications, to a considerable extent, one rejects. But the feigning of attitudes that one does not

of principles against our intuitions about particular cases ... fits well into the solid philosophical practices of constructing counterexamples and devising *reductios*. Consistency is a basic constraint on reason; we cannot accept a principle and refuse its consequences' (1980, p. 128). DePaul writes that 'surely when a person has inconsistent beliefs these cannot be rational, and surely the way to correct this problem is to reject [one of] the belief[s]' (1987, p. 473). Tersman writes that allowing inconsistency is 'unreasonable—a desideratum of a theory of justification is surely that it must not yield that each of a set of inconsistent beliefs is justified for the same person at the same time' (1993, p. 84; cf. p. 92). Finally, Blackburn writes that 'If our beliefs are inconsistent ... then something is wrong. Similarly if our attitudes are inconsistent ... then something is wrong' (1998, p. 309); but he adds that 'sometimes ... it is virtuous to be in two minds about things' (1998, p. 310).

Even many critics of reflective equilibrium accept the consistency thesis. Singer writes, 'If the reader simply cannot accept a moral judgment that follows from a moral theory, he must modify the theory, or else drop it altogether. To this extent the reflective equilibrium method is sound' (1974, p. 516). And Sencerz writes that 'it is uncontroversial that an *ideal* moral judgment is to be ... a member of a coherent system' (1986, p. 85; emphasis in original).

One author, however, who appears not to take the consistency thesis for granted is Amartya Sen, who observes that 'There could be good *instrumental* reasons for a moral theory to require inconsistent moral beliefs' (1982, p. 34). It should be noted, though, that he also expresses sympathy for the approach of reflective equilibrium (1982, p. 14).

actually hold is no part of the kind of thinking about morality that practical equilibrium makes room for. If an agent affirms a moral theory because she finds the arguments for it compelling and because it approves of her intuitions, then that does not mean that she will conceal either her affirmation of the theory or her conflicting intuitions. On the contrary, she might (and presumably should) openly affirm both the theory and her intuitions. In doing so, the agent need not (and probably would not) be saying that she adopted those intuitions because they were approved of by the theory; remember that the agent (in this story) had those intuitions anyway. She would simply be truthfully reporting that she has all of those beliefs, inconsistent though they may be. The value of this kind of inconsistent affirmation might have to do with explaining why one thinks about a particular issue in a certain way, or with persuading another person to adopt some or all of the beliefs one is affirming.

Of course, it would be possible for such an agent to believe that the best way to get other people to have the intuitions that the theory she affirms says *they* ought to have is for her to lie about what theory she affirms. So, as with government-house utilitarianism, there is nothing in practical equilibrium itself to preclude these undesirable results. But, again as with government-house utilitarianism, the same is true of reflective equilibrium. In each case, the agent herself must have intuitions (such as intuitions against systematically deceiving and manipulating other people) or other beliefs (such as beliefs in the ineffectiveness of attempts at large-scale deception and manipulation) that prevent her deliberation about morality from leading to hypocrisy.

A further concern about agents' psychology, beyond concerns pertaining to government-house utilitarianism and hypocrisy, has to do with another objection that goes back to Williams: an objection he raised against Hare's account of the 'critical' and 'intuitive' levels of moral thinking. In Hare's view, a person might affirm a particular moral theory at the critical level, while making judgements about particular cases at the intuitive level that are inconsistent with the theory she affirms. Of course, this is just the sort of situation that practical equilibrium would allow an agent to find herself in, if she already had some degree of inconsistency in her web of beliefs. (As I mention below, Hare is prominent among the theorists whose work anticipates the idea of practical equilibrium itself.) According to

Williams, if the agent affirms a consequentialist theory but her intuitions have significant non-consequentialist content, then her intuitions will provide her with a view of particular situations that, according to Williams, cannot be 'combine[d]' with the view afforded by the theory she affirms (1988, p. 189). He goes on to offer the following explanation of his objection:

In saying that you 'cannot combine' these two things, I do not mean that as a matter of psychological fact it is impossibly difficult. People indeed have thoughts that they describe in these terms—Hare himself has said that he does. The point is that the thoughts are not stable under reflection. (1988, p. 190)

Williams's objection seems to be about the rational workability of affirming a theory that is substantially inconsistent with one's intuitions. How, in short, can a moral agent operate on these two levels at the same time?

This question might pertain to both thought and action. In regard to thought, the question might express a suspicion that the agent's moral thinking is vulnerable to being overrun by an unruly jumble of moral beliefs. After all, if there is no consistency constraint, and if there are no standards allowing only some beliefs in and keeping others out, then it might seem that all bets are off. But a celebration of aimless inconsistency is no part of the moral psychology being considered here. On the contrary, far from giving the agent no standards for what to believe (and thus implicitly licensing all sorts of beliefs), the moral theory the agent affirms could well imply fairly tight constraints on the agent's beliefs and could well subject them to rigorous standards of assessment derived from the theory's doctrine of what sorts of conduct agents should dispose themselves to perform. Admittedly, as noted earlier, people cannot just choose what beliefs to have: in so far as an agent is ever required by her moral theory to have certain beliefs, the sense to be made of this notion is that the agent is required to perform those actions, of those available to her, that tend to shape her beliefs in the most desirable ways. But regardless of how severe this limitation happens to be in practice, worries about doxastic chaos are unmotivated. In place of a consistency constraint, other criteria apply.

Of course, as just mentioned, these other criteria are not fully determining, since people have only limited control over the beliefs they find compelling. Reflection and external influences inevitably affect the evolution of a person's beliefs, and although one can exercise some control over the way these processes unfold (e.g. by conscientiously reflecting on certain lines of reasoning, or by deliberately exposing oneself to certain external influences), internal and external influences unavoidably come unbidden. Moreover, it seems that when people are aware of inconsistencies in their thinking, they are naturally inclined to reflect on those inconsistencies and to iron them out when they can. So Williams's suggestion that inconsistencies between theory and intuition are not 'stable under reflection' may be particularly well-phrased.

But none of this reveals any serious problem with the psychology of an agent who, after following the guidance of practical equilibrium, does happen to affirm a moral theory that is inconsistent with her intuitions. Of course, as just discussed, her beliefs may well evolve so as to diminish or eradicate the inconsistency in her thinking. In that case, practical equilibrium would continue to counsel her, as always, to give weight to both entailment and approval. On the other hand, her beliefs, with their inconsistencies, may remain steady. All practical equilibrium says is that if that happens, then the agent can still rightly regard the theory she affirms as answering to her intuitions, and thus can rightly regard herself as having decided what to think about morality in a way that is at least as defensible as, and is arguably more defensible than, any way (such as reflective equilibrium) that insists on consistency between theory and intuition.

In regard to action, the question that Williams's objection raises might express doubt as to whether the agent herself will be able to be decisive in particular situations. But despite the complexity of the sort of psychology we are now considering, this matter need not be particularly mysterious. When an agent has a moral decision to make, it can be assumed that her decision will result from both the moral intuitions she has and the general moral principles she affirms, including those that constitute the moral theory she affirms. Her intuitions (regardless of their relation to her general principles) can be more or less firm, as can her affirmation of her general principles. In some situations, the agent's

intuitions might be so strongly motivating that she is uninterested in consulting her general principles. In some other situations (such as ones in which her intuitions are absent, weak, or unclear in their import), she might rely entirely on her general principles. In still others, she might consult both her intuitions and her general principles. Of course, given the complexity of the agent's psychology, it might be impossible for an external observer to predict the agent's choices with a high degree of accuracy. But this does not mean that the agent cannot conduct herself as a fully functioning moral agent.

The account of moral decision-making given in the previous paragraph, which concludes my response to concerns about thought and action raised by Williams's objection, is essentially explanatory: it explains how an agent with inconsistent moral beliefs can still be decisive in particular situations. But there is also, beyond Williams's objection, a related normative question: If an agent is in a situation in which her relevant moral intuitions conflict with the dictates of the moral theory she affirms, which beliefs ought to have authority in her decision-making? In other words, which beliefs would it be right for her to act on? Perhaps surprisingly, practical equilibrium does not purport to answer this question. Instead, practical equilibrium regards this as a substantive moral question to which different answers might be given by different acceptable ways of thinking about morality. For example, one agent who has followed the guidance of practical equilibrium in order to decide what to think about morality might end up with the belief (among her many moral beliefs at various levels of abstraction) that the right way to resolve conflicts between theory and intuition is to do what theory requires. In contrast, another agent who has followed the guidance of practical equilibrium might end up believing that intuition should prevail. A third might end up believing that theory should be followed in some circumstances but intuition in others. In principle, practical equilibrium is open to all of these possible ways of answering the question of the relative standing of theory and intuition in cases of conflict. So although practical equilibrium puts constraints on an agent's thinking about morality—requiring that the theory she affirms answer to her intuitions (in the broad sense discussed earlier)—its constraints do not include the thought that some particular answer to this normative question is the only acceptable one. This

question is an important one, but it is one for every agent to consider for herself, as part of her process of deciding what to think about morality.

6. Defending practical equilibrium, part 2: finding a moral theory we can live with

In the last section, I responded to a cluster of concerns that are probably the greatest obstacle to the reception of practical equilibrium as a way of deciding what to think about morality. In the next two sections, I want to offer two arguments affirmatively supporting the claim that practical equilibrium can justifiably be regarded as improving on reflective equilibrium. These arguments stem from considering exactly why we tend to be interested in what reflective equilibrium focuses on. As a way of deciding what to think about morality, including moral theories, reflective equilibrium focuses on the implications of those theories for particular situations and issues, such as the punishment of the innocent. The thought that I want to emphasize here is that our interest in these implications has a distinctly practical dimension that is absent or highly attenuated in contexts in which we are probing the implications and assessing the merits of, for example, scientific theories. This practical dimension, I submit, gives rise to two specific sets of considerations in support of practical equilibrium, one involving an interpersonal perspective and the other involving an intrapersonal perspective. I will discuss the first set of considerations in this section.

I begin with the thought that one of the reasons that we tend to be interested in the implications of a given theory when evaluating it is that we want to know whether it is, we might say, a theory we can live with, in the following sense: if the theory were widely accepted in a given society, then we could imagine living in that society and finding life there to be to our liking. Now I take it that the primary way in which the acceptance of a given theory in a society affects what it is like to live there is by affecting the behaviour of the people who live there. So if we are interested in what it would be like to live in a society of people who accept some theory, then we are interested in what sort of behaviour we could expect from those people, and what sort of interactions we could expect to have with them.

Normally, in order to anticipate the behaviour of people who accept some theory, we assume that they will act as the theory prescribes and so we look to the implications of the theory for specific cases. And certainly there are some moral theories whose implications for specific cases are reliable indicators of how people who accept those theories would behave. But whenever a theory approves of our having intuitions that are at odds with its implications—whenever, that is, the intuitions that the theory approves of differ from the ones that the theory entails—then those approved-of intuitions, and not only the theory’s implications, need to be taken into account in order to anticipate the behaviour of people who accept that theory. Of course, the intuitions that the theory approves of might not be widely held even among people who affirm the theory, if those intuitions are too alien to their ways of thinking—being approved of by the affirmed theory might not be enough to make them widely held. But if the intuitions that the theory approves of are easily occurring and easily sustained (as is, for example, the intuition that punishing the innocent is never justified), then it would be reasonable for an agent considering such a theory to think that the intuitions that the theory approves of, as well as those that it entails, would be among the intuitions held by people who affirm the theory. That is, it would be reasonable for an agent considering such a theory to think that the intuitions of people who affirm that theory would be some combination or mixture of the intuitions that the theory entails and the intuitions that the theory approves of. Intuitions related to the theory in both of these ways, not just intuitions related to the theory in the first way, would govern such persons’ conduct and would, for all practical purposes, define their moral characters.

So to the extent that our concern with a theory’s bearing on particular cases is based on our concern with whether it is a theory we can live with, we want to focus not only on the intuitions that the theory entails, but also on those that it approves of—the acceptability of a moral theory depends on both of these things. This is the position of practical equilibrium.

7. Defending practical equilibrium, part 3: 'I have my heart in the right place'

The last section was predicated on the thought that the distinctly practical dimension of the evaluation of a moral theory gives rise to a set of considerations involving an interpersonal perspective: that of whether the theory is one that we can live with in society with other people. This section, in contrast, involves an intrapersonal perspective, having to do with whether an agent evaluating a theory finds that its conception of what sort of person she ought to be harmonizes with her own.

To take up this perspective, recall the consistency thesis. This thesis can be understood as embodying the idea that there are certain moral facts or propositions that we *know*, in a sense, and that a theory that conflicts with these thereby implies that our thinking contains various factual mistakes about morality. But we should consider whether this is the only way in which a theory can clash with intuition. I submit that another way in which a theory can clash with intuition is by implying that certain of our intuitions are morally bad or, in an extreme case, that we are immoral or vicious people. What we find objectionable, I claim, is not only a theory's being inconsistent with certain of our intuitions, but also a theory's positing an ideal of what kind of people we ought to be—including what kind of intuitions we ought to have—that is at odds with what kind of people we are and intuitively think we ought to be. This line of thought provides another way of arriving at my earlier claim that there are two ways for a theory to be in harmony with intuition: not only entailment, but also approval.

Here is another way of making this point. Whatever a person's moral intuitions may be, we can impute to her two distinct claims (though each might be expected to be held by her implicitly rather than explicitly). One is simply the claim that her intuitions themselves are true—mostly if not entirely. Let us call this her truth claim. The other is the claim that it is morally good for a person such as herself, a person in her situation (whatever situation that may be), to have the intuitions she has—again, mostly if not entirely. Let us call this her goodness

claim.¹⁸ Now suppose this person is considering a particular theory to see whether it seems, to her, to be a good way to think about morality. Should she consider only whether this theory agrees with her truth claim, or should she also consider whether this theory agrees with her goodness claim? I submit that both of these things should matter to her in her assessment of the theory that she is considering. If the theory agrees with both of her claims, then she should credit that theory with being in harmony with her intuitions in two distinct, independently important, ways. Its agreement with her truth claim is not the only thing that matters; its agreement with her goodness claim matters, too.

To alter this example slightly, suppose that the theory she is considering agrees largely with her goodness claim, but not so thoroughly with her truth claim. (She might, for example, be considering a form of utilitarianism of the kind discussed in Sect. 3, or a form of deontology of the kind discussed in Sect. 4.) Such a theory, I claim, would deserve some credit for being in harmony with her intuitions. Admittedly, if another theory were to agree largely with her truth claim, but not so much with her goodness claim, then that theory would also deserve some credit for being in harmony with her intuitions. My point is not that agreement with an agent's truth claim fails to count as contributing to harmony between theory and intuition, but only that agreement with an agent's goodness claim also counts as contributing to harmony between theory and intuition.

This, of course, is the position of practical equilibrium, in contrast to reflective equilibrium. Reflective equilibrium, with its exclusive focus on a theory's entailment of an agent's intuitions, gives weight only to a theory's agreement with the agent's truth claim. In contrast, practical equilibrium, by giving weight to both a theory's entailment of an agent's intuitions and its approval of her intuitions, gives weight to both a theory's agreement with the agent's truth claim and its agreement with her goodness claim. Practical equilibrium's more expansive conception of what it means for a theory to answer to intuition enables it to give weight to both of the important claims to which a person's moral intuitions implicitly commit her.

¹⁸ I am grateful to Shelly Kagan for suggesting the distinction between these two claims.

I have been arguing that when an agent is deciding what to think about morality generally, and about a moral theory in particular, it is right for her to give weight not only to whether the theory implies the accuracy of her moral judgements, but also to whether it implies approval of her as a moral agent. This way of thinking builds on the work of many utilitarian writers, including Sidgwick, Sprigge, and Hare. But preceding all of these is Godwin, who famously considers the problem of whom to save if both his valet and the Archbishop Fénelon were in danger, and he could save either of them but not both. He writes that if he were confronted with such a situation, he ought to save the archbishop; and he adds that he ought to save the archbishop even if the other person were his father, brother, benefactor, or other loved one. But even the hard-headed Godwin, while insisting that he ought to save the archbishop, admits that if he were to save his father instead, then even though he would be acting wrongly,

every man will respect in me the sentiment of filial affection, will acknowledge that the feeling by which I am governed is a feeling pregnant with a thousand good and commendable actions, and will confess, according to a trite, but expressive, phrase, that at least I have *my heart in the right place*, that I have within me those precious and inestimable materials out of which all virtuous and honourable deeds are made. (1801, p. 187)

Now this remark comes in a passage in which Godwin is trying to show that his view of morality, which is essentially utilitarianism, is more congenial to an ordinary person's intuitions than it might be thought to be. So what Godwin means here, presumably, is that even though utilitarianism is inconsistent with some of an ordinary person's intuitions (specifically, the intuitions that might lead one to save one's father in such a situation), it is in harmony with those intuitions in an equally deep and important way, by approving of them. Practical equilibrium, but not reflective equilibrium, acknowledges this as a way in which a moral theory can be in harmony with intuition.

This virtue of practical equilibrium is reinforced by some reflection on a passage from George Eliot's *The Mill on the Floss*. In this novel, Maggie Tulliver's

father is in an ongoing conflict with Philip Wakem's father (who is, in effect, the Tullivers' landlord, Mr. Tulliver having lost ownership of the mill in a lawsuit). Her brother Tom warns her not to associate with Philip, on the grounds that doing so would be humiliating to their father, but Maggie's feelings are too strong and she befriends, secretly meets with, and falls in love with Philip. When Tom confronts her, he points out that even she knows, at some level, that what she has been doing is wrong. He says, 'If your conduct, and Philip Wakem's conduct, has been right, why are you ashamed at its being known?' (Eliot 1860, p. 449). Here is Maggie's reply:

I don't want to defend myself: I know I've been wrong—often, continually. But yet, sometimes when I have done wrong, it has been because I have feelings that you would be the better for if you had them. (Eliot 1860, p. 449)

Unfortunately it is not entirely clear what feelings Maggie is referring to. She might mean friendliness towards the deformed (Philip is ostracized because of the curvature of his spine), or openness to falling in love, or something else. But in any case, it is clear that she means that it is better that she has certain feelings, even if they lead her to act wrongly sometimes, than that she have the inhibitions that would guard her from those errors. So even though Maggie accepts the correctness of Tom's moral intuitions, she also appreciates the possibility—a possibility ignored or unnoticed by Tom—that the best kind of person for her (or her brother, or anyone) to be is one whose character and conduct are driven by intuitions that are different from (and, indeed, inconsistent with) those admittedly correct ones.¹⁹

¹⁹ This discussion has affinities with certain passages in R. M. Adams's 'Motive Utilitarianism'. See, in particular, his examples of Martha and Mary (1976, p. 475) and his suggestion that a utilitarian agent should have 'a vigorous desire to *live well*, in terms of the overall utility of his life, but not necessarily to act *rightly* on every occasion' (1976, p. 477). But Adams seems unlikely to embrace practical equilibrium as a way of deciding what to think about morality and moral theories. He suggests that instead of giving act utilitarianism credit for approving of certain desirable motives, we should modify it in the direction of rule utilitarianism so that its implications more closely match the (not strictly act-utilitarian) intuitions that would be associated with those desirable motives (1976, pp. 478–9).

If we sympathize with Maggie's perspective, we have reason to think that the wider perspective of practical equilibrium is preferable to the narrower one of reflective equilibrium. For when Maggie morally evaluates her own character, she apparently regards her having certain intuitions as contributing to her having a good character, even though they are not strictly correct. In other words, she is apparently fairly explicitly committed to what I earlier called a goodness claim—the claim that it is morally good for her to have the intuitions that she has. Moreover, she is apparently not as heavily committed to what would be her truth claim—the claim that her intuitions are true—because she acknowledges that her intuitions direct her to act wrongly from time to time. Thus, if Maggie were to consider a particular moral theory to see whether it seems, to her, to be a good way to think about morality, presumably she would regard it as important for the theory to agree with her goodness claim, and not only for it to agree with her truth claim. As we saw earlier, reflective equilibrium's exclusive focus on entailment means that it gives weight only to a theory's agreement with an agent's truth claim, whereas practical equilibrium's broader scope allows it to give weight to a theory's agreement with an agent's goodness claim, too.

Practical equilibrium, by attending to the intuitions that a theory approves of as well as the ones that it entails, evaluates a moral theory more favourably if the theory tells an agent that the kind of person she ought to be is, in fact, precisely the kind of person she is and intuitively thinks she ought to be. In this way, practical equilibrium gives credit to those moral theories that incorporate Maggie Tulliver's insight that the strictly correct intuitions may not be the constituents of the best kind of character for a person to have. Thus, if an agent is guided by practical equilibrium, she is more likely to give a moral theory credit for recognizing that she has her heart in the right place.²⁰

²⁰ A complication that arises here concerns the possibility of an agent's complying with theory-approved intuitions in a situation in which the theory also implies that those intuitions are incorrect. It might be thought that this places the moral theory in an awkward position, one whose awkwardness is brought into sharp relief by the question of what sort of account the theory should offer of such things as guilt, blame, and punishment. Should its account be act-based—maintaining that guilt, blame, and punishment are merited because of the wrongness of the act itself—or should

8. Practical equilibrium: partnering with reflective equilibrium?

Like reflective equilibrium, practical equilibrium is a way of deciding what to think about morality that involves taking everything that one thinks into account, sifting one's beliefs and weighing their relative force, seeing how they all relate to each other, and bringing them all into a broadly coherent system. All that practical equilibrium says, in departing from the model of reflective equilibrium, is that a system of beliefs can be broadly coherent in a manner not recognized by reflective equilibrium—specifically, that a moral theory can be in harmony with intuition through approval, not only through entailment. In other words, practical equilibrium offers a two-pronged conception of how a theory can be in harmony with a person's moral intuitions (i.e. by entailing them or by approving of them), in contrast to reflective equilibrium's one-pronged conception (on which approval does not count, only entailment does).

This broadened conception of harmony between theory and intuition could, of course, be absorbed into a revised conception of reflective equilibrium, obviating the need to conceptualize practical equilibrium as a variant of reflective equilibrium. Alternatively, one might wish to retain the distinct conception of practical equilibrium but acknowledge its roots in reflective equilibrium by thinking of

its account be character-based—maintaining that guilt, blame, and punishment are unmerited because of the moral desirability of the intuitions on which the agent acted?

I do not have space here to address this objection thoroughly, but the short answer is similar to what I said in the final paragraph of Sect. 5 about practical equilibrium's answer to the question of the right way for an agent to settle conflicts between theory and intuition. Different moral theories will presumably offer different sorts of accounts of the present matter—some act-based and some character-based—and a theory of either kind could, in principle, be evaluated favourably by practical equilibrium (or reflective equilibrium). For the question of the proper assignment of such things as guilt, blame, and punishment is internal to a moral theory, and one on which an advocate of practical equilibrium (or reflective equilibrium) can remain neutral. Still, this question does of course matter, and the way in which a theory answers it should figure in any agent's assessment of that theory, regardless of whether she chooses reflective equilibrium or practical equilibrium as her way of deciding what to think about morality.

practical equilibrium as partnering with reflective equilibrium or complementing it, rather than superseding it. Thus, it is simply for expository convenience that, throughout this paper, I characterize practical equilibrium as a freestanding (albeit derivative) alternative to reflective equilibrium. My aim is not to quarrel with the overall perspective of reflective equilibrium, but only to propose a particular revision of it. Whether the result is understood as a rival to reflective equilibrium or merely as a particular conception of it—or something in between—my aim here is only to defend the merits of this way of deciding what to think about morality.²¹

²¹ I am grateful to many people for their help with this paper. Dale Miller helped most of all, by responding to my early formulations of the main ideas and reading and commenting on multiple drafts. For probing questions and comments at presentations of early versions of this paper, I would like to thank the audience at the 2000 conference of the International Society for Utilitarian Studies at Wake Forest University, and audiences at the University of Pittsburgh, the University of Southern California, Columbia University, and the University of Kansas. Additionally, thoughtful comments were provided by David Gauthier, Donald Bruckner, Maura Tumulty, Bill Shaw, Shelly Kagan, Mark van Roojen, Pat Kitcher, and Carol Rovane. Finally, I am grateful for the detailed and very constructive criticisms provided by an editor and three referees for *Mind*.

References

- Adams, R. M. 1976: 'Motive Utilitarianism'. *Journal of Philosophy*, 73, pp. 467–81.
- Anscombe, G. E. M. 1958: 'Modern Moral Philosophy'. *Philosophy*, 33, pp. 1–19.
- Blackburn, Simon 1998: *Ruling Passions: A Theory of Practical Reasoning*. Oxford: Oxford University Press.
- Brandt, Richard B. 1979: *A Theory of the Good and the Right*. Oxford: Oxford University Press.
- 1996: *Facts, Values, and Morality*. Cambridge: Cambridge University Press.
- Daniels, Norman 1979: 'Wide Reflective Equilibrium and Theory Acceptance in Ethics', in his 1996b, pp. 21–46. Originally published in *Journal of Philosophy*, 76.
- 1980a: 'Reflective Equilibrium and Archimedean Points', in his 1996b, pp. 47–65. Originally published in *Canadian Journal of Philosophy*, 10.
- 1980b: 'On Some Methods of Ethics and Linguistics', in his 1996b, pp. 66–80. Originally published in *Philosophical Studies*, 37.
- 1996a: 'Introduction: Reflective Equilibrium in Theory and Practice', in his 1996b, pp. 1–17.
- 1996b: *Justice and Justification: Reflective Equilibrium in Theory and Practice*. Cambridge: Cambridge University Press.
- DePaul, Michael R. 1987: 'Two Conceptions of Coherence Methods in Ethics'. *Mind*, 96, pp. 463–81.
- Eliot, George 1860: *The Mill on the Floss*, ed. A. S. Byatt. London: Penguin Books, rpnt 1979.
- Freeman, Samuel (ed.) 2003: *The Cambridge Companion to Rawls*. Cambridge: Cambridge University Press.

- French, Peter A., Theodore E. Uehling Jr., and Howard K. Wettstein (eds) 1978: *Midwest Studies in Philosophy, III: Studies in Ethical Theory*. Morris, MN: University of Minnesota Press.
- Frey, R. G. 1977: 'Act-Utilitarianism: Sidgwick or Bentham and Smart?' *Mind*, 86, pp. 95–100.
- Gauthier, David 1986: *Morals by Agreement*. Oxford: Oxford University Press.
- Gibbard, Allan 1982: 'Inchoately Utilitarian Common Sense: The Bearing of a Thesis of Sidgwick's on Moral Theory'. In Miller and Williams 1982, pp. 71–85.
- Godwin, William 1793: *An Enquiry Concerning Political Justice*. In Godwin 1993, Vol. 3.
- 1801: 'Thoughts Occasioned by the Perusal of Dr. Parr's Spital Sermon'. In Godwin 1993, Vol. 2, pp. 163–208.
- 1993: *Political and Philosophical Writings of William Godwin, 7 Vols*, ed. Mark Philp. London: William Pickering.
- Hare, R. M. 1971a: 'The Argument from Received Opinion', in his 1971b, pp. 117–35.
- 1971b: *Essays on Philosophical Method*. Berkeley, CA: University of California Press.
- 1976: 'Ethical Theory and Utilitarianism', in his 1989, pp. 212–30. Originally published in Lewis 1976.
- 1978: 'Prediction and Moral Appraisal'. In French, Uehling, and Wettstein 1978, pp. 17–27.
- 1981: *Moral Thinking: Its Levels, Method, and Point*. Oxford: Oxford University Press.
- 1989: *Essays in Ethical Theory*. Oxford: Oxford University Press.
- 1997: *Sorting Out Ethics*. Oxford: Oxford University Press.
- Hill, Thomas E. Jr. 1980: 'Humanity as an End in Itself'. *Ethics*, 91, pp. 84–99.
- Hooker, Brad 1990: 'Rule-Consequentialism'. *Mind*, 99, pp. 67–77.

- 1991: 'Rule-Consequentialism and Demandingness: A Reply to Carson'. *Mind*, 100, pp. 269–76.
- 1994: 'Compromising with Convention'. *American Philosophical Quarterly*, 31, pp. 311–17.
- 1995: 'Rule-Consequentialism, Incoherence, Fairness'. *Proceedings of the Aristotelian Society*, new series, 95, pp. 19–35.
- 1996: 'Ross-style Pluralism versus Rule-consequentialism'. *Mind*, 105, pp. 531–52.
- 2000: *Ideal Code, Real World: A Rule-consequentialist Theory of Morality*. Oxford: Oxford University Press.
- Hume, David 1740: *A Treatise of Human Nature*, second edition, ed. L. A. Selby-Bigge and P. H. Nidditch. Oxford: Oxford University Press, 1978.
- Kamm, Frances 2007: *Intricate Ethics: Rights, Responsibilities, and Permissible Harm*. Oxford: Oxford University Press.
- Kant, Immanuel 1785: *Groundwork of the Metaphysics of Morals*. In Kant 1996, pp. 37–108. (Page numbers in the text refer to the Prussian Academy page numbers found in the margins.)
- 1793: *Kant on the Metaphysics of Morals: Vigilantius's Lecture Notes*. In Kant 1997, pp. 249–452. (Page numbers in the text refer to the Prussian Academy page numbers found in the margins.)
- 1797: *The Metaphysics of Morals*. In Kant 1996, pp. 353–603. (Page numbers in the text refer to the Prussian Academy page numbers found in the margins.)
- 1996: *Practical Philosophy*, tr. and ed. Mary J. Gregor. Cambridge: Cambridge University Press.
- 1997: *Lectures on Ethics*, tr. Peter Heath, ed. Peter Heath and J. B. Schneewind. Cambridge: Cambridge University Press.
- Lewis, H. D. (ed.) 1976: *Contemporary British Philosophy: Personal Statements*, fourth series. London: George Allen & Unwin Ltd.

- Mill, John Stuart 1859: *On Liberty*. In Mill 1977, pp. 213–310.
- 1861: *Utilitarianism*. In Mill 1969, pp. 203–59.
- 1969: *Collected Works of John Stuart Mill, Vol. 10: Essays on Ethics, Religion, and Society*, ed. J. M. Robson. Toronto: University of Toronto Press.
- 1977: *Collected Works of John Stuart Mill, Vol. 18: Essays on Politics and Society*, ed. J. M. Robson. Toronto: University of Toronto Press.
- Miller, Harlan B. and William H. Williams (eds) 1982: *The Limits of Utilitarianism*. Minneapolis, MN: University of Minnesota Press.
- Nagel, Thomas 1986: *The View from Nowhere*. Oxford: Oxford University Press.
- Nielsen, Kai 1972: 'Against Moral Conservatism'. *Ethics*, 82, pp. 219–31.
- Rawls, John 1951: 'Outline of a Decision Procedure for Ethics', in his 1999, pp. 1–19. Originally published in *Philosophical Review*, 60.
- 1971: *A Theory of Justice*. Cambridge, MA: Harvard University Press.
- 1975: 'The Independence of Moral Theory', in his 1999, pp. 286–302. Originally published in *Proceedings and Addresses of the American Philosophical Association*, 48.
- 1999: *Collected Papers*, ed. Samuel Freeman. Cambridge, MA: Harvard University Press.
- Rescher, Nicholas 1966: *A Constructive Critique of the Utilitarian Theory of Distribution*. Indianapolis, IN: Bobbs-Merrill Company, Inc.
- Ross, W. D. 1930: *The Right and the Good*. Oxford: Oxford University Press.
- Scanlon, T. M. 2003: 'Rawls on Justification'. In Freeman 2003, pp. 139–67.
- Scarre, Geoffrey 1996: *Utilitarianism*. London: Routledge.
- Seanor, Douglas and N. Fotion (eds) 1988: *Hare and Critics: Essays on Moral Thinking*. Oxford: Oxford University Press.
- Sen, Amartya 1982: 'Rights and Agency'. *Philosophy and Public Affairs*, 11, pp. 3–39.

- Sen, Amartya and Bernard Williams 1982a: 'Introduction'. In Sen and Williams 1982b, pp. 1–21.
- (eds) 1982b: *Utilitarianism and Beyond*. Cambridge: Cambridge University Press.
- Sencerz, Stefan 1986: 'Moral Intuitions and Justification in Ethics'. *Philosophical Studies*, 50, pp. 77–95.
- Shaw, William H. 1980: 'Intuition and Moral Philosophy'. *American Philosophical Quarterly*, 17, pp. 127–34.
- 1999: *Contemporary Ethics: Taking Account of Utilitarianism*. Oxford: Blackwell Publishers.
- Sidgwick, Henry 1907: *The Methods of Ethics*, seventh edition. Indianapolis, IN: Hackett Publishing Company, 1981.
- Singer, Peter 1974: 'Sidgwick and Reflective Equilibrium'. *The Monist*, 58, pp. 490–517.
- Sinnott-Armstrong, Walter 1992: 'An Argument for Consequentialism'. *Philosophical Perspectives*, 6, pp. 399–421.
- Smart, J. J. C. 1973: 'An Outline of a System of Utilitarian Ethics'. In Smart and Williams 1973, pp. 1–74.
- Smart, J. J. C. and Bernard Williams 1973: *Utilitarianism: For and Against*. Cambridge: Cambridge University Press.
- Sprigge, T. L. S. 1965: 'A Utilitarian Reply to Dr. McCloskey'. *Inquiry*, 8, pp. 264–91.
- 1988: *The Rational Foundations of Ethics*. London: Routledge & Kegan Paul.
- 1989: 'Utilitarianism and Respect for Human Life'. *Utilitas*, 1, pp. 1–21.
- Tersman, Folke 1993: *Reflective Equilibrium: An Essay in Moral Epistemology*. Stockholm: Almqvist & Wiksell International.
- Williams, Bernard 1973: 'A Critique of Utilitarianism'. In Smart and Williams 1973, pp. 75–150.

—— 1988: 'The Structure of Hare's Theory'. In Seanor and Fotion 1988, pp. 185–96.

Woollard, Fiona 2008: 'Intricate Ethics and Inviolability: Frances Kamm's Nonconsequentialism'. *Ratio*, new series, 21, pp. 231–8.