

NORMATIVE ETHICS: BACK TO THE FUTURE

Thomas Hurka

University of Toronto

The course of normative ethics in the 20th century was a roller-coaster ride, from a period of skilled and confident theorizing in the first third of the century, through a virtual disappearance in the face of various forms of skepticism in the middle third, to a partial revival, though shadowed by remnants of that skepticism, in the final third. The ideal future of normative ethics therefore lies in its past. It must entirely shed its traces of mid-century skepticism if it is to return to the levels of insight provided by G. E. Moore, Hastings Rashdall, J. M. E. McTaggart, W. D. Ross, C. D. Broad, and other early 20th century moral theorists.

These theorists shared several fundamental assumptions about ethics, many derived from late 19th century philosophers such as Henry Sidgwick. They were all moral realists, believing that moral judgements are objectively true or false. More importantly, they were non-naturalist realists, believing, as anti-realists also can, that moral judgements form a separate category of judgements, neither reducible to nor derivable from other judgements. For them the property of goodness is not identical to any physical or natural property, and no “ought” can be derived from an “is.” They therefore accepted a realist version of the autonomy of ethics: at the level of fundamental principles, moral judgements are independent of all other judgements.

These theorists also shared the general normative project of systematizing common-sense morality, or finding more abstract principles that can unify and explain our particular judgements about good and right. This project was not unique to them, but their approach to it was

distinctively shaped by their belief in the autonomy of ethics.

First, they trusted intuitive judgements as at least a reliable starting-point for moral inquiry. Precisely because they denied that moral claims can be derived from other claims, they thought direct intuitive insights, either their own or those of common sense, were the best available entree to the moral realm. They did not see these insights as infallible; they recognized that our intuitions can be distorted by self-interest and other factors. Nor did they think the disorganized and even conflicting collection of judgements that make up common-sense morality was in that form acceptable. This was their prime motive for theorizing common sense: only if its judgements could be systematized by a few fundamental principles would they be properly scientific. But one test of these principles was their consistency with everyday moral beliefs, and another was their own intuitive appeal. Many of these theorists considered intuitions about general principles more reliable than ones about particular cases, but at all levels of generality they thought the only route to moral knowledge was by some kind of intuitive insight.

Second, and because of their confidence in intuitive judgements, these theorists were in principle open to the whole range of moral views accepted in Western culture.¹ In practice their openness often had a limitation. Many of them were consequentialists, believing that what is right is always what will produce the most good, and as a result they did not say much about non-consequentialist views. Even Ross, who defended non-consequentialism, described more its general structure than its details. But Moore, Rashdall, and the others did collectively address a huge variety of views about the good: not only that pleasure is good, but also views valuing knowledge, aesthetic contemplation, virtue, love, and more. There was a golden age of value theory, in part because its theorists defended so many views about what is intrinsically worth

pursuing. In addition, they were prepared to explore the details of these views. They did not rest with general claims about their preferred values but produced subtle analyses of the elements of aesthetic contemplation (Moore), the specific character of love (McTaggart), and the forms of virtue and their comparative values (Ross).²

Finally, the analyses these theorists produced were what I will call structural rather than foundational. They described the underlying structure of common-sense judgements and in that sense unified and explained them. But they did not try to ground those judgements in others that concern a more fundamental topic and can therefore justify them to someone who does not initially accept them. Their analyses stayed within a circle of common-sense concepts rather than connecting them to others, either moral or non-moral, that they saw as more secure. For example, Sidgwick grounded utilitarianism in the principles that one should not prefer a lesser good at one time to a greater good at another or a lesser good for one person to a greater good for another. These principles make explicit two forms of impartiality central to utilitarian thinking; they also help to unify utilitarian claims. But they use similar concepts to those claims rather than relating them to others that are more fundamental.³ The same holds for Moore's formulation of the retributive theory of punishment in terms of his principle of organic unities, which says the value of a whole need not equal the sum of the values of its parts. In the case of punishment, Moore argues, a person's having a vicious character is bad, as is his suffering pain, but the combination of a vicious character and pain in the same life is good as a combination, and sufficiently good that inflicting the pain makes the overall situation better.⁴ This analysis again uncovers the structure of retributive claims and unifies them with others that involve organic unities. It also has important implications, for example, that while deserved punishment is good as deserved it is

bad as involving pain, so the morally best response to it mixes satisfaction that justice is being done with regret at the infliction of pain. But the analysis does not ground retributivism in some other, less contentious claim and will not persuade someone initially hostile to retributivism. Or consider Broad's treatment of what he calls self-referential altruism. This view holds, contrary to Sidgwick, that our duty to others is not to treat them impartially but to care more about those who are in various ways closer to us, such as our family, friends, and compatriots. Broad's analysis unifies a variety of common-sense claims about the demands of loyalty and invites further inquiry about exactly which relations make for closeness of the relevant kind. But it does not justify self-referentiality in other terms; instead, it assumes self-referentiality, saying that each person should care more about his family and friends because he should in general care more about those who stand in special relations *to him*.⁵

These theorists were aware of more ambitious normative projects: many of their contemporaries proposed deriving moral claims from associationist psychology, Darwinian biology, or Idealist metaphysics. But Sidgwick, Moore, and the others gave both general arguments against this kind of derivation and specific critiques of their contemporaries' views. For them the foundational approach to ethics was illusory and structural analysis the only profitable route to pursue. They did not address every important topic in ethics or leave nothing to be said about those they did discuss. But methodologically they provided a model for how normative inquiry should proceed. And then, in the middle decades of the century, normative ethics virtually disappeared from philosophy.

One cause of this disappearance was the replacement of the moral realism that had dominated the earlier period by crude versions of expressivist anti-realism, which held that moral

judgements are not true or false but only express simple pro- or con-attitudes, and which understood normative argument as an attempt to transmit these attitudes to others by a kind of emotional contagion.⁶ Another cause was a general conception of philosophy as a second-order discipline, which analyzes the logic or language of first-order disciplines but does not participate in them itself. Just as philosophy of science analyzes the logic of scientific confirmation but does not itself make scientific claims, so philosophical ethics should study the language of morals but leave substantive moralizing to preachers and poets.

Over time, however, the influence of these causes faded and the intrinsic interest of normative questions, both theoretical and particular, was able to reassert itself. The result was a revival of normative philosophizing in the last third of the century, stimulated especially by the 1971 publication of John Rawls's *A Theory of Justice*.⁷ Since then normative ethics has been a prominent part of the discipline and has produced much valuable work. But its revival has been only partial, held back by remnants of the mid-century's two dominant general attitudes to philosophy.

The first of these was the technical, scientizing attitude of logical positivism and successor views such as W. V. O. Quine's naturalism. This attitude was hostile to common sense and philosophies that take it seriously, holding that the everyday view of the world is riddled with errors. In the linguistic terms that were popular in this period, it held that ordinary language is inadequate for understanding reality and needs to be replaced by a more scientific language, such as first-order logic. The second attitude, which arose in reaction to the first, informed the ordinary-language philosophy of Ludwig Wittgenstein, J. L. Austin, and others. It had a high regard for common sense, which it took to be the repository of centuries of human learning. But

it was resolutely anti-technical and anti-theoretical. It is a mistake, its partisans held, to think that whenever a single word applies to a set of objects there must be some single property they share; there may be only a loose set of “family resemblances.” The attempt to capture that property in an abstract principle does not deepen the insights of common sense but gets them entirely wrong.

Even apart from any general skepticism about philosophical ethics, this pair of attitudes left little room for the earlier project of systematizing common-sense morality. On one side was a view friendly to abstract principles but hostile to common sense; on the other was a view friendly to common sense but hostile to abstract principles. And these two views have continued to influence normative ethics since its re-emergence.

For its part, the scientizing attitude has encouraged philosophers to reject many common-sense moral views as confused or in some other way unacceptable. The result, especially early in the normative revival, was that theoretical ethics considered only a small number of options: utilitarianism, Kantian deontology, and little else. Ideas about desert, natural rights, and virtue were commonly if not universally set aside or analyzed in other, allegedly less suspect terms; either way, the distinctive approaches to ethics they express were ignored. Another object of skepticism was the topic most discussed at the start of the century: intrinsic value. Many philosophers found the idea that there are goods a person should pursue for herself other than pleasure or the satisfaction of her desires deeply problematic, and they were equally hostile to proposed goods that span people’s lives, such as distributions proportioned to their merits, or that lie outside them, such as complex ecosystems. This meant that consequentialism, which in the earlier part of the century had encompassed a wide variety of positions, was in the later period mostly equated with simple-minded versions of utilitarianism. This again cut philosophical ethics

off from everyday moral thought. Just as many non-philosophers use desert and rights as basic moral concepts, so many think there are goods worth pursuing other than pleasure or satisfaction. If a person is compassionate or generous, that in itself makes her life better; likewise if she has deep personal relationships or accomplishes difficult goals. (Consider the longtime recruiting slogan of the U. S. armed forces: “Be all that you can be.” It clearly appealed to widespread views about the intrinsic value of developing one’s talents.) In rejecting all such views philosophers were again rejecting much of their culture’s moral life.

In some cases the grounds of this rejection were conceptual, with philosophers urging, for example, that claims about intrinsic value are simply unintelligible. These particular arguments had little merit. If morality can tell people to pursue others’ happiness regardless of whether this will satisfy their own desires, as most philosophers allow, why can it not also tell them to pursue knowledge? And if the worry was that claims about intrinsic value presuppose a suspect realism, that too was groundless. Even if moral realism is problematic, as is by no means clear,⁸ claims about value can always be understood in an expressivist way. In fact, sophisticated versions of expressivism allow virtually any moral view to be accommodated in a scientizing or naturalistic picture of the world. These versions take the attitudes expressed by moral judgements to have the logical form of categorical imperatives, so they are directed at acts or states of a person in a way that is not conditional on that person’s having any attitudes. Consider the judgement that malice is evil. According to sophisticated expressivism, someone who makes this judgement expresses a negative attitude to all malice, whatever the malicious person’s attitude to malice. If she contemplates malice in someone who has no negative attitude to malice, her own attitude to the malice is negative; if she contemplates a possible world in which she herself has no negative

attitude to malice and is malicious, her attitude (from this world) to her malice is negative.⁹ As so formulated, expressivism makes virtually no difference to the study of normative questions. Both realists and expressivists can accept almost any substantive moral view and argue for it in the same way: by appealing to intuitive judgements, formulating abstract principles, and so on. The realists will interpret the judgements as providing insights into moral truth, the expressivists as expressing attitudes they hope others do or can come to share. But viewed on their own, their moral positions will be indistinguishable.

In other cases the grounds for rejecting common-sense views were normative: that people must be free to determine the content of their own good, as desire theories allow, or that all goods are, substantively, goods of individuals. But these arguments were often accompanied by a general sense that the views they targeted were too extravagant to be taken seriously. It is part of the self-image of scientizing philosophers to be tough-minded debunkers of confused folk beliefs, and in ethics this meant rejecting all views with more than a very austere content.¹⁰ But it is again hard to see a persuasive rationale for this approach. Whether moral judgements report a distinctive kind of truth or merely express attitudes, why should their content be limited in this way? Whatever the merits of conceptual parsimony in other domains, it is hard to find one here.¹¹

Relatedly, and especially when ethics was first re-emerging, philosophers tended to confine their attention to structurally simple views rather than recognizing the complexities their predecessors had noted. Whereas Broad had analyzed the structure of self-referential altruism, several prominent works took the main views needing discussion to hold that people should either care only about their own good or care impartially about the good of all.¹² And a common form of argument assumed that if a moral factor such as the difference between killing and

allowing to die makes no difference in one context, it cannot make a difference in any context.¹³

But the point of Moore's doctrine of organic unities had been precisely that the difference a factor makes can vary from context to context, depending on what other factors it is combined with.

These scientizing influences have been countered but also complemented by those of the anti-theoretical, Wittgensteinian attitude. Its adherents are not debunkers; they are open to many common-sense moral views, especially about the virtues and vices. But they deny that these views can be systematized or captured in abstract principles. Many cite Aristotle's remark that in ethics one should not seek more precision than the subject-matter allows,¹⁴ which they take to imply that theorizing about ethics is fundamentally misguided. The most extreme formulation of their view holds that moral knowledge always concerns particular acts in all their specificity; it does not generalize to other acts, cannot be codified in general principles, and is a matter only of trained moral insight.¹⁵ People with the right moral character can "see" what is right, just, or virtuous in a given situation, but they cannot express that vision in other terms or communicate it to those who lack it.

In my view an anti-theoretical position is properly open only to those who have made a serious effort to theorize a given domain and found that it cannot succeed. Anti-theorists who do not make this effort are simply being lazy, like Wittgenstein himself. His central example of a concept that cannot be given a unifying analysis was that of a game,¹⁶ but in one of the great underappreciated books of the 20th century Bernard Suits gives perfectly persuasive necessary and sufficient conditions for something's being a game. (Roughly: in playing a game one pursues a goal that can be described independently of the game, such as directing a ball into a hole in the

ground, while willingly accepting rules that forbid the most efficient means to that goal, such as placing the ball in the hole by hand.) With an exemplary lightness of touch, Suits mentions Wittgenstein only once:

‘Don’t say,’ Wittgenstein admonishes us, “‘there must be something common or they would not be called ‘games’” -- but *look and see* whether there is anything common to all.’ This is unexceptionable advice. Unfortunately, Wittgenstein himself did not follow it. He looked, to be sure, but because he had decided beforehand that games are indefinable, his look was fleeting, and he saw very little.¹⁷

Similarly, ethical anti-theorists have decided beforehand that there can be no unifying account of, say, the human good and therefore do not try seriously to construct one. More specifically, they typically consider only simple-minded ethical analyses and take the failure of those to demonstrate the impossibility of all analyses. But, for example, the fact that not all pleasure is good, because sadistic pleasure is bad, does not refute all general theories of the value of pleasure. An only slightly complex theory can say that sadistic pleasure, while good as pleasure, is bad as sadistic, and more bad than it is good; a more complex theory can say that when pleasure is sadistic it loses its goodness as pleasure. Far from abandoning generality, these analyses use it to illuminate values in a way anti-theorists never could.

Despite their differences, scientizers and anti-theorists share a common assumption: that any acceptable moral theory must be simple in its content and form. Believing in theory, the scientizers confine their attention to simple views; finding those views unacceptable, anti-theorists abandon theory. But both are hostile to the systematic analysis of complex moral views

that had been the hallmark of the earlier period. Consider, for example, the topic of how different moral considerations weigh against each other. Scientizers are suspicious of such weighing, especially if it rests on intuitive judgements. They want a mechanical and even empirically implementable procedure for weighing values, as there would be if all values reduced to a single one.¹⁸ By contrast, anti-theorists embrace a plurality of values but insist they are “incommensurable,” which they take to imply that nothing systematic can be said about how they compare.¹⁹ The early 20th century theorists avoided both these extremes. While recognizing that different values cannot be weighed precisely, they insisted that we can make rough comparative judgements about them, such as that an instance of one good is much, moderately, or only a little better than an instance of another. They also pursued structural questions, such as whether the complete absence of one good, say, knowledge, can always be compensated for by a sufficient quantity of another. There was the intermediate approach of partly theorizing the partly theorizable, but it is excluded by both the scientizing and anti-theoretical attitudes.

Among some ethicists the influence of these attitudes is now fading.²⁰ In the last decade or so there has been more sympathy for views that would earlier have been rejected as extravagant, such as perfectionist accounts of each person’s good,²¹ as well as detailed explorations of non-personal goods such as equality²² and desert.²³ There has also been greater awareness of the complexity moral views can have; thus, the point that a factor’s importance can depend on its relations to other factors has been made and is widely accepted.²⁴ Scientizing and anti-theoretical attitudes have by no means disappeared. Some philosophers still reject any normative claims not derived from desires; others still look askance at all formal analysis. But the two attitudes are becoming less dominant, and as they recede, ethics is coming to theorize a

wider range of views. But there is another, more indirect effect of the scientizing attitude that continues to retard moral study.

Because they are skeptical of appeals to moral intuition, scientizers are dissatisfied with what I have called structural analyses, ones that relate a moral view to principles that are more abstract but use similar concepts. To them these principles appeal to the same basic intuitions as the original view and so cannot properly justify it. There are many illustrations of this dissatisfaction.

In a recent book on desert, George Sher says that Moore's defence of retributivism in terms of his principle of organic unities is "inconclusive," since it "merely restates what the retributivist needs to explain." Though intuitive appeals like Moore's are not worthless, they are at best a "prologue" that should lead to "independent justifications of our beliefs about desert."²⁵ Similarly, Rawls allows that a pluralist view that weighs its values intuitively can describe the structure of its comparative judgements, for example, on indifference graphs, but adds that since these graphs give no "constructive criteria" establishing the judgements' reasonableness, what results is "but half a conception."²⁶ Or consider non-consequentialism. A great contribution of recent ethics has been to show that non-consequentialist views have a self-referential or "agent-relative" structure. When they forbid, for example, killing, they tell each person to be especially concerned that *he* does not kill, even if the result is a somewhat greater number of killings by other people.²⁷ But rather than being seen to support non-consequentialism, by clarifying its structure, this analysis has been taken to generate objections against it. How can it be rational, some have asked, to avoid one act of killing if the result is more killings overall? If killing is bad, should we not try to minimize killing by everyone?²⁸ Here it is not taken to be a sufficient answer

to point to the intuitive appeal of an agent-relative prohibition against killing or even of the abstract structure it embodies. What is demanded is a justification of agent-relativity of some deeper, more philosophical kind.

This dissatisfaction has led many contemporary philosophers to search for foundational justifications of moral views, ones that relate them to other concepts they see as somehow more secure. These justifications have taken many forms. In some the foundational principles appealed to are still moral. Thus, a prominent justification of retributivism appeals to ideas about fairness, saying it is unfair if the majority of people have restrained their self-interest by obeying the law but a criminal has not. He has gained the benefit of others' restraint without paying a similar cost himself, and punishing him removes that imbalance.²⁹ Here the proposed foundational claim is at a coordinate level to those being justified; fairness is not a more abstract concept than desert. But other justifications start from very abstract principles, for example, that everyone should be treated with equal respect and concern³⁰ or that those acts are wrong that are forbidden by rules no one could reasonably reject.³¹ They then claim that the best interpretation of these principles, guided by normative views but not by ones directly favouring a given moral position, can support that position and therefore justify it philosophically. The grand exemplar of this approach is Rawls, who says the correct principles of justice are those that would be chosen by rational contractors in a specified initial position. Rawls's general contractarian claim is best understood as a moral one, and moral judgements also guide his specification of his initial position. But since these judgements do not directly concern the principles the contractors choose among, the contract provides an "Archimedean point" for justifying specific claims about justice.³²

These last analyses shade into ones that are more explicitly ambitious, claiming to derive

specific moral views from the logic of moral language or the definition or purpose of morality. R. M. Hare claims that the language of morals, properly understood, allows only utilitarianism as a fundamental moral view; all other views misuse the language.³³ More vaguely, others claim that since the purpose of morality is to satisfy human wants and needs, any acceptable principles must concern wants and needs.³⁴ This seems to rule out perfectionist views of the good by definitional fiat, and certainly rules out goods not located within people's lives, such as distributions proportioned to their merits and the existence of ecosystems.

Yet another approach appeals to metaphysical facts, especially about the person, that it says an acceptable moral view must reflect. Rawls says that utilitarianism fails to take seriously "the separateness of persons," and that doing so leads to a more egalitarian view like his own;³⁵ Robert Nozick says the same facts about separateness support agent-relative prohibitions against using some as means to benefit others.³⁶ In a similar vein, Samuel Scheffler argues that what he calls "the independence of the personal point of view" justifies agent-relative permissions. Apart from its deontological prohibitions, common-sense morality does not require people always to produce the most good impartially calculated. It permits them to give somewhat more weight to their own interests and therefore, within limits, to fall short of doing what is impersonally best. This permission can be justified, Scheffler argues, by noting that people do not rank outcomes only from an impersonal standpoint but also, and independently, in terms of their own preferences and values.³⁷

Yet another foundational approach tries to demonstrate the rationality of certain moral principles, so that failing to act on them is a species of irrationality. The most clear-headed such argument, that of David Gauthier, claims that accepting certain moral constraints is in each

person's self-interest, where self-interest is understood austere, as involving the satisfaction of her pre-existing desires.³⁸ A more high-minded view understands self-interest as involving an ideal of "flourishing" that may or may not be what a person actually wants. Here the justification of morality is that flourishing involves the moral virtues as essential components, and these virtues require one to act morally.³⁹ An especially grand form of argument, proposed by Alan Gewirth and Christine Korsgaard, tries to derive moral demands from the presuppositions of rational agency. The beliefs implicit in any rational action, it claims, commit one logically to certain moral principles, typically Kantian ones, so failing to act on those principles involves logical inconsistency.⁴⁰

Despite their differences, these arguments share the general project of trying to ground moral views in something more than direct intuition, through what I have called foundational analyses. Now, the difference between these and structural analyses is only one of degree, not of kind. It turns on whether the concepts an analysis uses are similar to those in the view it is analyzing, and similarity admits of degrees. There can even be analyses that straddle the structural-foundational divide, so they are neither clearly the one nor clearly the other. But a difference of degree is still a difference, and other analyses do clearly fit one model rather than the other. Thus, Rawls's and Hare's arguments are clearly foundational, while Moore's and Broad's are structural. There can also be similarities in how the two types of analysis are defended. Structural analyses cannot argue just from particular intuitive judgements to principles that fit them; if the principles are to explain the judgements, they must be independently plausible. Conversely, foundational analyses need not argue just from the general to the particular. They can and often do take the fact that an abstract idea coheres with and can explain

particular judgements to be important evidence for it, as Rawls, for example, does. Both approaches can therefore use coherentist reasoning, treating all parts of a moral theory as justified by their relations to other parts and none as immune to revision or the source of all others' warrant.⁴¹ But even when they share this coherentism the two approaches differ in the type of moral theory they generate and, more specifically, in the type of moral explanation they give. Structural analyses assume that a concrete moral view can be explained by principles that express similar ideas in a more abstract way, using similar concepts at a higher level of generality; those who demand foundational analyses require a genuine explanation to relate the view to ideas that are different and more fundamental, because they concern some more fundamental topic.

It would be impressive if any of these foundational arguments succeeded, but in my view critical discussion has repeatedly shown that they do not. It would be going too far to state categorically that no such argument will ever succeed; each must be assessed on its merits. But time and again it turns out that to yield the specific views they are meant to justify, they must tacitly appeal to the very intuitive judgements they are meant to supplant. Let me give a few illustrations.

Arguments from metaphysical facts about the person need not violate the autonomy of ethics if they make the general moral claim that the correct normative principles for a thing must reflect that thing's nature.⁴² But in most cases the facts they start from either are not specific enough to support the particular moral views they are intended to or, if they are redescribed to do so, presuppose those views or something very close to them. For example, it is undeniable that persons are separate, in the sense that there are connections between states within a life that do not hold across lives. But many moral views reflect this fact other than the egalitarian ones it is

usually taken to support. Consider a pluralist view that compares goods such as knowledge and achievement in a way that favours balance or well-roundedness within lives but not across them. It treats lives as morally significant units but has no distributive implications whatever. And even if separateness does have distributive implications, why must they be egalitarian? What is wrong in this respect with a desert principle requiring people's degrees of happiness to be proportioned to their degrees of virtue, or even with an antiegalitarian principle like Nietzsche's that directs everyone to promote the excellence of the few most excellent individuals? Of course, one can redescribe the separateness of persons so it does support only egalitarian views, for example, by saying that unit gains in less happy lives count for more than similar gains in happier ones. But then egalitarianism is being built into the supposedly metaphysical facts, which no longer give it independent support. A similar point applies to Scheffler's justification of agent-relative permissions. The fact that people have a personal ranking of outcomes that can differ from the impersonal one seems perfectly well captured by a view that combines each person's personal ranking with the impersonal one, say by averaging them, and then requires him to maximize the resulting combined value. To support permissions in particular over this alternative, the metaphysical facts must be redescribed so they favour giving each person with an independent point of view the further independence to decide which point of view he will act from on particular occasions. But then the redescription again assumes something very close to the conclusion it is meant to justify.⁴³

A similar difficulty faces arguments from abstract moral principles such as ones about equal respect or rational agreement. In some cases these principles will be widely accepted, but only because they are open to different interpretations that support very different substantive

views; in other cases they do favour one view, but have contentious features that cannot be independently justified.⁴⁴ The former is true of arguments from equal respect. Their foundational principles are unexceptionable if they say only that a moral view must give all people equal standing, but this claim has no substantive implications. In particular, it does not imply that people should have, even initially, equal resources or happiness; it is perfectly satisfied by views that give people equal rights to acquire resources (which they may then exercise very unequally) or say they should receive different rewards according to their different deserts. The abstract claim does not even distinguish between consequentialist and deontological views, which each interpret respect in their own way.⁴⁵ The contrary problem arises at many points in Rawls's contractarian argument. Its basic design excludes perfectionist, entitlement, and many other views -- so it is hardly a neutral device -- and even views it initially allows, such as utilitarianism, are eventually excluded by stipulations that have no persuasive contractarian rationale. Or consider the suggestion that agent-relative prohibitions are justified because, by making persons "inviolable," they give them higher moral status than they would otherwise have.⁴⁶ It may be true that these prohibitions give each person higher status by denying that she may legitimately be sacrificed whenever this will give greater benefits to other people. But they also give her lower status by denying that other people may be sacrificed to give benefits to her, and there is no independent way of saying whether the gain or loss here is greater.⁴⁷

These abstract arguments share a further difficulty with some that try to ground a moral view in another of a coordinate degree of abstractness, such as the fairness justification of retributivism. These latter arguments are difficult to generalize about and may sometimes prove illuminating. But consider the claim of the fairness justification that punishment removes an

imbalance whereby some people have benefited from others' restraint without paying a similar cost themselves. If cost is understood in a standard way, say in terms of pleasure and pain, this claim does not yield the right results about particular punishments: many crimes that intuitively call for severe punishments, such as murder, are not ones people pay a high cost in avoiding, since they are not ones they are strongly tempted to commit. So the justification must revise its understanding of cost, and in doing so make some finely balanced claims about the relative positions of criminals and law-abiding citizens. And we can then ask whether those claims really are explanatory of retributivism. On one side is a simple judgement many find intuitively compelling, about the intrinsic appropriateness of punishing the guilty; on the other are some precarious claims about the balance of costs and benefits. Is it really an advance to relate the former to the latter? I am not suggesting that everyone must accept retributivism; many may on reflection reject it. But their reflection will be most productive if it considers retributivism in its own right rather than detouring through a view that seems too convoluted to underlie it.⁴⁸ A similar point applies to more abstract arguments. Rawls tries to ground his egalitarian claim that it is most important to improve the condition of the worst-off by arguing that his contractors will not know the probabilities of their occupying different positions in society and, lacking that knowledge, will prefer those principles whose worst outcome for them is least bad. Critics have questioned whether either of these stipulations is reasonable, in my view persuasively. But even if they are reasonable, how much is gained by relating an egalitarianism many find compelling in itself to these finely balanced claims about rational choice?

It was not with pleasure that the early 20th century theorists said, as they often did, that the resolution of a fundamental moral issue comes down to each person's intuitive judgement about

the question when it is clearly framed. They would have been delighted if more satisfying justifications were available. But if they are not available, the widespread search for foundational analyses in recent ethics has been quixotic. This is not a decisive objection; many philosophical projects are quixotic, and the discipline still learns from their failure. But in this case the pursuit of grand justifications has in several ways skewed the field's priorities.

First, it has helped keep the focus of normative ethics narrow; as well as rejecting moral views for being extravagant, philosophers can dismiss them for lacking deep foundations. In some cases this dismissal has not had such chilling effects. Thus, the charge of some that agent-relative prohibitions are irrational has not stopped others from exploring their details. But, in another case, the novel entitlement theory of justice described by Robert Nozick has received less constructive discussion than it deserves in part because of the complaint that Nozick's treatment leaves it "without foundations."⁴⁹ This might be a reasonable complaint if competing views did have such foundations; it is not when they do not.

Second, the search for foundations has diverted philosophers' attention from the details of moral views. Faced with a set of ideas about value or duty, a theorist can go in either of two directions. She can look into the view, to discover its structure and the way it generates specific claims, or she can try to connect it to larger ideas about, say, the nature of morality or the status of persons as free and rational. And the more she does of the latter the less she will do of the former. In fact, a climate of foundational inquiry can lead even theorists sympathetic to a given view to spend more time answering others' abstract objections to it than examining its specific contents. But in ethics as in many fields, God (or the devil) dwells in the details. Our common-sense ideas about any moral topic are at best loosely organized and demand systematization.

Sometimes this task cannot be completed: the ideas turn out to conflict, and we can only retain some by abandoning others. But at other times theorizing a view can show it has hidden strengths. For example, what seemed at first a loose collection of judgements can turn out to have a sophisticated underlying structure, even a mathematical one, that everyday moralists follow though they are unaware of doing so. Or the view can have a surprising unity. It may include, say, a pair of plausible intuitive judgements that look independent, so we can wonder whether they are even consistent. But then analysis reveals that, perhaps given a third plausible judgement, the two entail each other, so the view has an impressive internal integrity. Or the view can be surprisingly related to several others, all of which turn out to instantiate the same more abstract idea. Or, examined closely, it can generate fascinating new questions and offer intuitively satisfying answers to them -- sometimes just one answer, sometimes a set of competing ones. Any of these developments counts in favour of a view's acceptability, as internal contradictions count against it. But none will emerge if philosophers are busy batting around abstract conditions and chasing foundational dreams.

In fact, the search for moral foundations can lead to high-minded distortions of the moral phenomena. Consider the increasingly popular view that most desires are for objects thought of as good or as somehow supported by reasons. It has led one writer to say that thirst involves a recognition that we have reason to drink,⁵⁰ but surely no one actually gets thirsty in this highly intellectualized way. As Plato said, what we want when we are thirsty is not a good drink, but just a drink.⁵¹ And even Plato was too restrictive in limiting this type of desire to the bodily appetites. Anything we can desire on the basis of an evaluative judgement we can also desire apart from such judgements; as well as pursuing knowledge because we believe it is intrinsically

good, we can pursue it from simple curiosity. Moore, Rashdall, and Ross recognized this, holding that alongside forms of virtue directed at the right and the good as such are forms that involve no evaluative thoughts, such as simple compassion.⁵²

Other distortions are more disturbing. A persistent temptation for foundational views, especially ones that affirm the rationality of morality, is to ground moral requirements in facts about the self: in how acting on them promotes the agent's self-interest or flourishing or relates to a "practical identity" or conception she has of herself as a specific type of person.⁵³ But if they take this line, these views imply that my ultimate reason for benefiting another person, say, by relieving his pain, concerns how this will affect my life. And this is not the right ultimate reason, which concerns how my action will affect the other's life: my fundamental reason to help him is to help *him*. In a trivial sense morality has to concern me, since it has to tell me how to act. But it does not follow that the ground of its requirements must always be located in me, and views that place it always there are in that respect distortions. This is reflected in the fact that people who were motivated by what these views say are their ultimate reasons -- people whose ultimate motive for helping another was to promote their own flourishing or to live up to their own practical identity -- would be acting in an objectionably narcissistic way. It is not admirable but morally self-indulgent to relieve another's pain primarily from concern to be a virtuous reliever of pain.⁵⁴ Now, a moral view can say that people will only act successfully on their ultimate reasons if they do not try consciously to do so. Thus, consequentialist views often say that people will only produce the best consequences if they do not aim at those consequences directly but instead follow certain simple rules. But it would be very odd for the foundational views I am discussing to make this type of claim. That would involve their saying that acting on what they

say are our ultimate reasons is objectionable in itself, and not just because of its effects, and surely a moral view cannot so directly condemn its own acceptance. That is a sign that something has gone deeply wrong -- that the search for philosophical foundations has led theorists to place the source of our reasons where it intuitively cannot be.

It is in fact striking how many strands in contemporary ethics, of Kantian, Aristotelian, and other inspirations, approach the subject in an essentially self-centred way. They see the fundamental task of the moral agent as being to reflect on her own desires, or care about her own flourishing or integrity, or give moral laws to herself and so achieve her own autonomy. On all these views the agent's fundamental orientation is toward herself. There are, of course, exceptions; Thomas Nagel says very clearly that another person's pain by itself gives me reason to relieve it.⁵⁵ But a great deal of currently influential writing seems to assume that moral reasons must be not only trivially but also substantively about the self. The causes of this are partly philosophical: disbelief that there could be underivative other-regarding reasons and distrust of the intuitions that say there are. But there may also be social causes. May the self-centredness of so much contemporary ethical theorizing be the predictable product of a narcissistic age, offering moral foundations for the me-generation?

Finally, there is an issue of philosophical style. The early 20th century theorists wrote about ethics simply, directly, and even (one thinks, for example, of Rashdall) with wit. But much contemporary ethical writing has a tone of high intellectual earnestness. Sidgwick thought the progress of ethics had been impeded by the desire to edify; today the greater danger is the desire to be philosophically profound. At the very least, the pursuit of grand foundational projects does not encourage an engaging lightness of touch.

Any assessment of the current state of ethics must address the mixed influence of Rawls. *A Theory of Justice* is a great book, with many novel ideas and an original proposal for uniting them. And in a lecture given shortly after that book's publication, Rawls made methodological recommendations similar to those of this essay: that at least for now philosophers should set aside questions of moral truth and examine the structures of the moral views people actually hold.⁵⁶ But these recommendations have had much less influence than Rawls's practice or what was taken to be his practice in *A Theory of Justice*. Many even of those who do not think that book's contractarian argument succeeds think it is the right type of argument, aiming to justify a moral view not by direct intuitions but by something more philosophically ambitious. And Rawls's later work has reinforced this tendency, with its emphasis on grand abstractions such as persons' "higher-order interest" in forming, revising, and pursuing a conception of the good and the ideal of "public reason." Rawls's writing in moral philosophy has had many effects, but one has been to encourage a style of argumentation that is more high-flown than it is productive.

How, then, should normative ethics proceed? I take it that on most views its starting-point is the moral judgements people actually make. And it is worth emphasizing that people do make these judgements. Despite the claims of some cultural analysts, categorical moral evaluations are not disappearing in the face of moral relativism. No one says that rape, racism, and genocide are morally all right for you if you believe they are all right; everyone simply condemns them as wrong. (Though some describe themselves as relativists, their moral practice repeatedly belies that claim.) And though the content of these judgements may change somewhat over time, their core remains the same. There is less concern today than fifty or a hundred years ago with the ethics of consensual sexual activity, and more with discrimination and intolerance. But the

central moral prohibitions against coercion and wilful harm remain essentially unchanged, as do the central recommendations of compassion and fidelity. There are metaethical questions about how these prohibitions and recommendations are to be understood: in a realist way, as making claims that can be true or false, or as merely expressing attitudes. But these questions do not affect, except at the margins, either the content of these judgements or their importance.

Whatever metaethics decides, the practice of making moral judgements will continue essentially as before.

In their initial form, however, these judgements require philosophical analysis, both to understand and to assess them. In my view this analysis is best done by addressing them in their own right, by seeking to identify more abstract principles that, while using similar concepts, can organize and explain them, and by investigating their details. In fact, these two tasks often go together, since trying to describe a view's structure comprehensively can force one to address specific issues one might otherwise not have thought of. Much valuable work of this kind has been done in recent ethics. The clear description of agent-relative prohibitions and permissions, which require us to care especially that *we* do not kill and allow us to give more weight to *our* interests, is, when separated from bogus demands for deep justifications, an immense contribution to our understanding of common-sense and deontological views. Other writers have shown how, in formulating their prohibitions, these views can use either or both of the distinctions between harms one causes and those one merely allows, and between harms one intends as an end or means and those one merely foresees; they have also discussed how these two distinctions might combine with each other. This is not a topic where a completely satisfactory analysis has yet emerged. Perhaps that analysis requires a more subtle application of

the principle of organic unities than has yet been considered; perhaps it is simply not possible. Even so, the identification of the two distinctions has greatly improved our understanding of this family of views.⁵⁷ There have also been subtle suggestions about how a deontological view can weigh its prohibitions against the overall good an action will cause without aggregating that good in a simple additive way. Thus, the view can make it permissible to kill one innocent person to save some large number of other people from being killed or otherwise seriously harmed, but not to save any number of people from mild headaches.⁵⁸

There have been similar insights into views about the good. Thus, it has been shown that these views need not require everyone to maximize the good. They can instead require agents only to aim at outcomes that in one of several senses are “good enough”; this identification of “satisficing” consequentialism has again widened the range of moral views philosophers can consider.⁵⁹ There have also been discussions about whether the onerousness of one person’s duty to pursue good outcomes should be affected by other people’s failure to fulfil their duty, and if not, why not.⁶⁰ And there have been illuminating analyses of particular values, such as a distinction between two kinds of broadly egalitarian view⁶¹ and a searching examination of the complexities of one of them,⁶² as well as the beginning of a sophisticated structural analysis of judgements about desert.⁶³ There remain many moral views that have not received the structural analysis they deserve, but on others significant progress has been made.

In some cases this analysis will contribute significantly to assessing the views. Thus, it may be that while one of two competing views resists structural analysis, because it contains ineliminably contradictory elements, another can be beautifully unified by an attractive general principle. This result supports the second view against the first. But in other cases even complete

analysis may leave us with difficult judgements to make. If each of two views can be successfully unified, the choice between them comes down to an overall intuitive assessment that we may find difficult to make and that different people may make differently. Here there may be neither intuitive certainty for one person nor agreement among persons. This may also be the result in the more likely case where each view can be unified to a considerable degree but not entirely, so there is a plausible general principle that captures many of the intuitive judgements associated with it but requires others to be abandoned. But this is just the situation we face if, as I have argued, there are no magic-bullet arguments in normative ethics and intuitive judgements of some kind are all we have to go on. We can play such judgements off against each other, including judgements at different levels of generality. If we think some judgements are more likely to be distorted we can give them less weight, though if we do the correction will always come from other intuitive judgements. And if we think some judgements are especially reliable, we can place more trust in them. This methodology leaves philosophical analysis only some comparatively modest tasks: to clarify the issues our judgements must address and to help us understand the views that resolve them differently. But those are not insignificant tasks; on the contrary, and unlike the high flights of foundational analysis, they can actually yield results.

The main themes of this essay were anticipated by Friedrich Nietzsche. In *Beyond Good and Evil* he urged moral philosophers to own up

to what alone is justified so far: to collect material, to
conceptualize and arrange a vast realm of subtle feelings of value
and differences of value which are alive, grow, beget and perish ...
all to prepare a *typology* of morals.

To be sure, so far no one has been so modest. With a stiff seriousness that inspires laughter, all our philosophers demanded something far more exalted, presumptuous, and solemn from themselves as soon as they approached the study of morality: they wanted to supply a *rational foundation* for morality -- and every philosopher so far has believed that he has provided such a foundation. ... How remote from their clumsy pride was that task which they considered insignificant and left in dust and must -- the task of description -- although the subtlest fingers and senses can scarcely be subtle enough for it.⁶⁴

Subtle fingers and senses, yes: normative ethics will do best if it does more of the work Nietzsche recommends and less of the kind that would have him rolling on the floor.

Notes

1. They should also have been open to ideas accepted in other cultures, but tended to believe, like many in their time, in the higher moral development of the West.

2. G. E. Moore, *Principia Ethica* (Cambridge: Cambridge University Press, 1903), pp. 189-202; J. M. E. McTaggart, *The Nature of Existence*, vol. 2 (Cambridge: Cambridge University Press, 1927), pp. 147-61, 436-39; W. D. Ross, *The Right and the Good* (Oxford: Clarendon Press, 1930), pp. 155-73.

3. Henry Sidgwick, *The Methods of Ethics*, 7th ed. (London: Macmillan, 1907), pp. 379-84.

4. Moore, *Principia Ethica*, pp. 214-16.

5. C. D. Broad, "Self and Others," in *Broad's Critical Essays in Moral Philosophy* (London: George Allen & Unwin, 1971).

6. See A. J. Ayer, *Language, Truth, and Logic* (London: Gollancz, 1936), Ch. 6; and Charles L. Stevenson, *Ethics and Language* (New Haven: Yale University Press, 1944).

7. John Rawls, *A Theory of Justice* (Cambridge, MA: Harvard University Press, 1971).

8. For a recent defence of realism see Thomas Nagel, *The View From Nowhere* (New York: Oxford University Press, 1986), Ch. 8.

9. This point is made in recent defences of expressivism such as Allan Gibbard, *Wise Choices, Apt Feelings: A Theory of Normative Judgment* (Cambridge, MA: Harvard University Press, 1990); and Simon Blackburn, *Ruling Passions: A Theory of Practical Reasoning* (Oxford: Clarendon Press, 1998).

10. Compare Bernard Williams's claim that utilitarianism's popularity rests on its being a

“*minimum commitment morality*”; see his *Morality: An Introduction to Ethics* (New York: Harper & Row, 1972), p. 91.

11. Note that austerity is not just a matter of theoretical simplicity. Robert Nozick’s view that all values are instances of “organic unity” (see his *Philosophical Explanations* [Cambridge, MA: Harvard University Press, 1981], pp. 415-28) is theoretically simple but would be rejected by scientizers as hopelessly extravagant.

12. See, e.g., Thomas Nagel, *The Possibility of Altruism* (Oxford: Clarendon Press, 1970).

13. For the most famous illustration of this argument see James Rachels, “Active and Passive Euthanasia,” *New England Journal of Medicine* 292 (Jan. 9, 1975): 78-80.

14. Aristotle, *Nicomachean Ethics*, trans. W. D. Ross and J. O. Urmson, in *The Complete Works of Aristotle*, ed. Jonathan Barnes (Princeton, NJ: Princeton University Press, 1984), 1094b12-28. Note that, despite this remark, Aristotle gives a mathematical formula for distributive justice at *Nicomachean Ethics* 1131a30-b16, saying that justice requires the ratio of person *A*’s award to *B*’s to equal the ratio of *A*’s merits to *B*’s. He is nothing like the consistent anti-theorist some of his contemporary admirers take him to be.

15. See, e.g., John McDowell, “Virtue and Reason,” *The Monist* 62 (1979): 331-50; and Jonathan Dancy, *Moral Reasons* (Oxford: Blackwell, 1993), Chs. 4-6.

16. Ludwig Wittgenstein, *Philosophical Investigations*, trans. G. E. M. Anscombe (Oxford: Basil Blackwell, 1972), sec. 66.

17. Bernard Suits, *The Grasshopper: Games, Life and Utopia* (Toronto: University of Toronto Press, 1978), p. x. Note that Suits’s analysis of games is structural; once it is in hand,

Wittgenstein's fussing about the differences between games that use boards and cards or that are and are not amusing is embarrassingly superficial.

18. Rawls expresses a moderate version of this view in *A Theory of Justice*, saying that pluralist views that weigh principles intuitively are less satisfactory than ones that give some principles lexical priority over others, so the former's demands always take precedence over the latter's (pp. 40-45).

19. See, e.g., Martha C. Nussbaum, *The Fragility of Goodness: Luck and Ethics in Greek Tragedy and Philosophy* (Cambridge: Cambridge University Press, 1986), pp. 107-17, 294-98; and Michael Stocker, *Plural and Conflicting Values* (Oxford: Clarendon Press, 1990), Chs. 7-8.

20. In philosophy generally the scientizing attitude remains very strong; in fact, successor views to Quine's naturalism are now the dominant methodological views in the discipline. The anti-theoretical attitude lost its position in general philosophy decades ago and now survives almost solely in ethics.

21. See, e.g., Derek Parfit, *Reasons and Persons* (Oxford: Clarendon Press, 1984), pp. 493-502; Thomas Hurka, *Perfectionism* (New York: Oxford University Press, 1993); and George Sher, *Beyond Neutrality: Perfectionism and Politics* (Cambridge: Cambridge University Press, 1997), Ch. 9.

22. Derek Parfit, "Equality or Priority?," The Lindley Lecture, (Lawrence, KS: University of Kansas Press, 1995); and Larry S. Temkin, *Inequality* (New York: Oxford University Press, 1993).

23. See, e.g., George Sher, *Desert* (Princeton, NJ: Princeton University Press, 1987); and Shelly Kagan, "Equality and Desert," in Louis P. Pojman and Owen McLeod, eds., *What Do We*

Deserve? Readings on Justice and Desert (New York: Oxford University Press, 1999), pp. 298-314.

24. Shelly Kagan, "The Additive Fallacy," *Ethics* 99 (1988): 5-31; and F. M. Kamm, *Morality/Mortality, Vol. II: Rights, Duties, and Status* (New York: Oxford University Press, 1996), Ch. 2.

25. Sher, *Desert*, pp. 72, 19.

26. Rawls, *A Theory of Justice*, pp. 39, 41.

27. Bernard Williams, "A Critique of Utilitarianism," in J. J. C. Smart and Bernard Williams, *Utilitarianism For and Against* (Cambridge: Cambridge University Press, 1973), 77-150, p. 99; Parfit, *Reasons and Persons*, pp. 143-48; Nagel, *The View From Nowhere*, pp. 152-56, 175-85.

28. Samuel Scheffler, *The Rejection of Consequentialism: A Philosophical Examination of the Considerations Underlying Rival Moral Conceptions* (Oxford: Clarendon Press, 1982), Ch. 4; and Shelly Kagan, *The Limits of Morality* (Oxford: Clarendon Press, 1989).

29. Herbert Morris, "Persons and Punishment," *The Monist* 52 (1968): 475-501; and Sher, *Desert*, Ch. 5.

30. Ronald Dworkin, "Liberalism," in Stuart Hampshire, ed., *Public and Private Morality* (Cambridge: Cambridge University Press, 1977), pp. 113-43.

31. T. M. Scanlon, *What We Owe to Each Other* (Cambridge, MA: Harvard University Press, 1998).

32. Rawls, *A Theory of Justice*, p. 261. In later work Rawls has abandoned this strong "Archimedean" claim about his contractarian argument, but still holds that connecting his liberal-egalitarian political views to contractarian ideas gives them a kind of justification they would not

have if defended just by direct intuition.

33. R. M. Hare, *Moral Thinking: Its Levels, Method, and Point* (Oxford: Clarendon Press, 1981).

34. See, e.g., Philippa Foot, "Moral Beliefs," *Proceedings of the Aristotelian Society* 59 (1958-59): 83-104; G. J. Warnock, *Contemporary Moral Philosophy* (London: Macmillan, 1967), Ch. 5-6; and Warnock, *The Object of Morality* (London: Methuen, 1971).

35. Rawls, *A Theory of Justice*, pp. 27, 187.

36. Robert Nozick, *Anarchy, State, and Utopia* (New York: Basic Books, 1974), pp. 32-33.

37. Scheffler, *The Rejection of Consequentialism*, Ch. 3; for another version of metaphysical foundationalism see David O. Brink, "Self-Love and Altruism," *Social Philosophy and Policy* 14/1 (Winter 1997): 122-57.

38. David Gauthier, *Morals By Agreement* (Oxford: Clarendon Press, 1986). Because of its austerity, this argument cannot justify all the moral constraints most people find intuitively compelling.

39. G. E. M. Anscombe, "Modern Moral Philosophy," *Philosophy* 33 (1958): 1-19; and Rosalind Hursthouse, "Virtue Theory and Abortion," *Philosophy and Public Affairs* 20 (1991): 223-46.

40. Alan Gewirth, *Reason and Morality* (Chicago: University of Chicago Press, 1978); and Christine M. Korsgaard, *The Sources of Normativity* (Cambridge: Cambridge University Press, 1996)

41. Because of this, analyses that I call foundational need not involve foundationalism in the sense used in epistemology, where it connotes the view that some beliefs are self-justifying

and the source of justification of all other beliefs. Analyses that appeal to the language of morals or defend the rationality of morality involve something like this epistemic foundationalism, but others such as Rawls's do not. The latter are foundationalist only in my sense, which concerns explanation rather than justification.

42. Rawls, *A Theory of Justice*, p. 29.

43. In addition, most appeals to a metaphysical fact assume a certain valuation of it; thus, Scheffler assumes that the independence of the personal point of view is not simply regrettable, as strict impartialists might hold, but has positive worth. This is another point where these appeals come close to assuming what they are meant to explain.

44. I think most philosophers would now agree that a similar dilemma faces arguments from the language of morals or the definition of morality. Most would say its having substantive moral implications actually counts against a view on these topics; an adequate view should allow the disagreements that obviously exist by remaining neutral on substantive moral questions.

45. Consequentialist views say that respect for a person involves promoting her good, deontological views that it involves, more centrally, not sacrificing her good. This is a fundamental difference, but cannot be resolved by reflecting on the true meaning of "respect."

46. Kamm, *Morality/Mortality*, Vol. II, Ch. 10.

47. Kagan, "Replies to My Critics," *Philosophy and Phenomenological Research* 51 (1991): 919-28, pp. 919-20.

48. In his version of the fairness justification Sher equates the cost a criminal avoids with the strength of the moral prohibition he violates (*Desert*, pp. 81-82), but this brings the justification very close to the Moorean formulation it is meant to improve upon.

49. Nagel, "Libertarianism Without Foundations," *Yale Law Journal* 85 (1975): 136-49.

50. Scanlon, *What We Owe to Each Other*, p. 38.
51. Plato, *Republic*, trans. G. M. A. Grube (Indianapolis: Hackett, 1974), 438a-439b.
52. Moore, *Principia Ethica*, p. 179; Hastings Rashdall, *The Theory of Good and Evil* (London: Oxford University Press, 1907), vol. 1, pp. 119-24; Ross, *The Right and the Good*, pp. 160-63, 170-73.
53. For the latter view, see Korsgaard, *The Sources of Normativity*.
54. On moral self-indulgence see Bernard Williams, "Utilitarianism and Moral Self-Indulgence," in his *Moral Luck: Philosophical Papers 1973-80* (Cambridge: Cambridge University Press, 1981), pp. 40-53.
55. Nagel, *The View From Nowhere*, pp. 156-62.
56. Rawls, "The Independence of Moral Theory," *Proceedings of the American Philosophical Association* 48 (1974-75): 5-22. But Rawls includes in the "structure" of a moral view an associated "conception of the person," which sounds more foundational than the more purely mathematical analysis of structure I have been urging.
57. See, e.g., Philippa Foot, "The Problem of Abortion and the Doctrine of Double Effect," in her *Virtues and Vices* (Berkeley, CA: University of California Press, 1973), pp. 19-32; Judith Jarvis Thomson, "Killing, Letting Die, and the Trolley Problem," *The Monist* 59 (1976): 204-17; Warren S. Quinn, "Actions, Intentions, and Consequences: The Doctrine of Doing and Allowing," *Philosophical Review* 98 (1989): 287-312; Quinn, "Actions, Intentions, and Consequences: The Doctrine of Double Effect," *Philosophy and Public Affairs* 18 (1989), 334-51; Jeff McMahan, "Killing, Letting Die, and Withdrawing Aid," *Ethics* 103 (1993): 250-79; Kamm, *Morality/Mortality, Vol. II*, Chs. 1-7.
58. Samantha Brennan, "Thresholds for Rights," *Southern Journal of Philosophy* 33

(1995): 143-68; Kamm, *Morality/Mortality, Vol. I: Death and Whom to Save from It* (New York: Oxford University Press, 1993), Chs. 8-10; and Scanlon, *What We Owe to Each Other*, pp. 229-41.

59. Michael Slote, *Common-Sense Morality and Consequentialism* (London: Routledge & Kegan Paul, 1985), Ch. 3; and *Beyond Optimizing: A Study of Rational Choice* (Cambridge, MA: Harvard University Press, 1989), Chs. 1-2.

60. Liam B. Murphy, *Moral Demands in Nonideal Theory* (New York: Oxford University Press, 2000); Brad Hooker, *Ideal Code/Real World* (Oxford: Clarendon Press, 2000).

61. Parfit, "Equality or Priority?"

62. Temkin, *Inequality*.

63. Kagan, "Equality and Desert."

64. Friedrich Nietzsche, *Beyond Good and Evil*, trans. Walter Kaufmann (New York: Vintage, 1966), sec. 186.