

On Normative Ethics

Thomas Hurka

University of Toronto

I became interested in normative ethics in my last term as a philosophy undergraduate at the University of Toronto. Influenced by a traditional conception of the discipline, I'd till then studied mostly history of philosophy, with a special interest in, of all things, Hegel. But seeing the value of a balanced philosophy program, I enrolled in an ethics seminar in the winter of 1975. I'd studied the ethics of Plato, Leibniz, Hegel, and others in my history courses, but this was my first exposure to contemporary thought on the subject.

The seminar had only one other student, and since he was fairly quiet I had something close to a one-on-one tutorial with the instructor, Wayne Sumner – another contributor to this volume. I was immensely taken both by Wayne's teaching and by the seminar's topic, which was utilitarianism. Our study of it included both theoretical questions such as act- vs. rule-utilitarianism and applied ones such as our obligations to future generations. I found all these questions intellectually engaging and also humanly important, given their connection to live ethical issues – in short, I was hooked.

I was especially impressed by two non-utilitarian texts we read. One was the last chapter of G. E. Moore's *Principia Ethica*, meant to illustrate a non-

utilitarian or ideal consequentialist theory. I was attracted by Moore's general account of the good, which values states such as love and aesthetic appreciation alongside pleasure, and saw in his principle of organic unities a far clearer statement of ideas I had encountered in more pretentious form in Hegel. The second was Chapter 2 of W. D. Ross's *The Right and the Good*, which defends a pluralist deontology against consequentialism. I thought Ross's criticisms of consequentialism were right on the mark, and again clearly and precisely expressed. I had no idea these texts were largely *passé* in current discussions of the subject.

At this time I was applying to graduate programs in philosophy, and given my new interest decided to specialize in ethics. I chose Oxford as my university, and when I got there, spent some time studying metaethics, for example, emotivism and prescriptivism, with R. M. Hare. But the main activity in Oxford then was in normative ethics. There were joint seminars on normative topics each led by (some among) Ronald Dworkin, James Griffin, Jonathan Glover, J. L. Mackie, Derek Parfit, Amartya Sen, and Charles Taylor, as well as seminars by individual philosophers, such as Parfit's on what would become *Reasons and Persons*. They all used the same analytic methodology I'd been impressed by in the Toronto seminar, though on a wider range of topics. And I wanted to use it on a wider range still.

My thesis idea was to use this methodology to defend an ethical view that I initially called 'self-realizationist' and later 'perfectionist.' Understood broadly, this view tells us to promote a conception of the good that, like

Moore's, includes 'ideal' states such as knowledge and achievement alongside or instead of pleasure or desire-fulfilment. Understood more narrowly, it identifies these ideal goods as developing properties that are somehow fundamental to human nature, for example, by being essential to humans. This view wasn't being discussed much in Oxford or indeed anywhere in analytic philosophy, but it was one I'd encountered and been attracted to while studying Aristotle, Leibniz, Hegel, and other historical figures at Toronto and that I thought could be rehabilitated using ideas like Saul Kripke's account of essential properties. The idea was to put old perfectionist wine in new bottles.

It's said that the key to a successful thesis is to pick a 'sunrise' rather than a 'sunset' topic, one that other people have started to write on but that hasn't yet been worked to death. In this terminology mine was a pre-dawn topic, and in working on it I was largely on my own. This led to some floundering; for example, it took me a long time to see that perfectionism can tell each person to promote the development of everyone's human nature rather than just her own. But I worked away at the topic through a B. Phil. and D. Phil. thesis to, finally, in 1993, a published book called *Perfectionism*.

Over time my treatments became, I hope, less inadequate, but they also changed focus. The B. Phil. thesis was almost entirely devoted to the narrow perfectionist idea of grounding the good in human nature, with very little on the details of the moral view that would result. The D. Phil. was more evenly balanced between the two topics, while the book gave the bulk of its attention

to the detailed theory. After three chapters on human nature, it had eight about a broadly perfectionist theory that could be based on ideas about human nature but could also be defended apart from them; these chapters discussed, among other things, the aggregation and balancing of perfectionist goods, the parallels between knowledge and achievement as goods, and the political implications of a theory centred on them.

In the years since, my thinking has moved further in that direction, so I'm now more sceptical about grounding the good in human nature. It's not that I don't think the narrow perfectionist idea has intuitive appeal; I still think it does. But I'm less confident that a dispassionate examination of our best explanations of human behaviour really does pick out all and only those properties that are intrinsically worth developing. I'm more inclined to think that knowledge and achievement – *Perfectionism's* account of which I still stand by – are best valued on their own rather than as based in human nature, as are other states such as virtue that can't be accommodated in the narrow perfectionist framework.

And my shift in attitude has been more general. Normative ethics tries to systematize our everyday moral judgements by relating them to principles that are somehow more explanatory. But there are two different ways of doing this, or two ends on a continuum of ways.

One relates everyday judgements to principles that are more abstract but at the same time continuous with those judgements, because they use similar concepts and concern a similar topic. An example is Moore's account

of retributive desert using his principle of organic unities: while vice and pain are both evil, he says, the combination of vice and pain in the same life is good as a combination, and sufficiently so that inflicting the pain is on balance an improvement. This account reveals the underlying structure of retributive judgements, connects them to others involving organic unities, and suggests further questions, such as how the goodness of deserved pain compares with the evil of vice. But it doesn't justify the judgements in other terms and so won't persuade anyone not already sympathetic to retributivism. Another example is the claim that we have stronger duties to promote the happiness of our children, friends, and other intimates because we have in general stronger duties concerning people who stand in certain special relations to us. This again reveals the structure of associative duties and suggests further questions, such as what exactly the relevant relations are, but it doesn't justify the duties in non-associative terms. I call this style of normative theorizing *structural*.

Many contemporary philosophers find this approach unsatisfactory. They say structural principles are too close to the everyday judgements they're meant to explain: they in effect assume what they're meant to justify and so don't justify at all. That requires relating a moral judgement to claims that use different concepts and concern some other, more fundamental topic, one that does yield justifications. The resulting *foundational* style of theorizing has several variants. Some of these try to ground moral judgements outside morality itself, for example in scientific theories about psychology or

evolutionary biology, in metaphysical theories of the self, or in semantic claims about the meanings of the moral words. Others relate everyday judgements to claims within morality, but on some other, allegedly more fundamental topic. Thus, contemporary neo-Aristotelians say the reason we should promote others' happiness is that this will express benevolence on our part, where virtues like benevolence are essential for our flourishing, which of necessity is our ultimate aim. Another exemplar of this approach is John Rawls's *A Theory of Justice*, which assumes that claims about liberty and equality can't be affirmed just on their own. They're properly justified only if we can show that they would be chosen by rational contractors in what Rawls calls the 'original position.' This position isn't independent of morality, since we use moral judgements in specifying its features. But it's far removed from everyday thinking about political values, which makes no reference to rational contracting.

The narrow perfectionism discussed in my book was to a considerable extent foundationalist. Unlike many versions of similar ideas, it didn't try to ground the human good directly in non-evaluative claims about human nature; instead, it affirmed a substantive moral principle valuing whatever properties that nature turns out to involve. But ideas about human nature are fairly distant from everyday thought about the value of understanding or achievement, and I now extend my scepticism about those ideas to foundational theorizing in general. However impressive it may seem, and

however exciting the justifications it promises to deliver, it rarely if ever succeeds.

This can't be shown across the board; individual foundationalist arguments must be examined individually. But time and again they turn out not to yield the conclusions they're meant to, or not without tacitly assuming what they're meant to prove. (The latter was the foundationalist charge against structural analyses, but they don't claim to explain on a wholly different basis.) Consider the neo-Aristotelian idea that the virtues are those traits a person needs in order to flourish or live well. To yield the intended kind of explanation, it needs an independent understanding of what human flourishing consists in, which many give in terms of rationality. But if rationality is understood in a not morally loaded way, it can surely be exercised in malicious acts as much as in benevolent ones. And if it's instead taken to include a power to know and be guided by true practical principles such as 'Promote others' happiness,' then it assumes its intended conclusion.

And there's a further problem: many foundational arguments distort the phenomena at issue by turning the duties everyday thought recognizes into something they're not. Consider again the neo-Aristotelian view. To the question, 'What's the ultimate reason why you should promote another's happiness?' it answers 'Because that will make your own life better,' and that's not the right answer. The right answer is 'Because it will make the other person's life better.' Here the foundational argument has turned an other-regarding duty into an egoistic one, and the same occurs in Rawls's theory.

He thinks fairness requires his contractors not to know their particular conception of the good, or the specific values they're pursuing. But he then assumes that they'll try to maximize the resources they'll need to pursue whatever conception of the good they have. And that makes their primary concern to live, not a life that's in fact good, but whatever life they think is good, so they care about their conception of the good not because it's true but because it's theirs.

This leaves the alternative, structural style of theorizing, and while it doesn't yield as grand explanations as foundationalism, it does yield explanations. It connects everyday judgements to principles that are more general and therefore more explanatory and that are also often intuitively appealing in themselves, or apart from their implications. To illustrate this, consider the very different, non-Aristotelian account of virtue discussed in my second book, *Virtue, Vice, and Value*.

This account takes virtue to consist in morally fitting attitudes to other, previously given goods and evils. If something is intrinsically good, then loving it for itself, or desiring, pursuing, and taking pleasure in it for itself, is also intrinsically good and a virtue. So if another person's pleasure is good, then desiring, pursuing, and taking pleasure in his pleasure is good and, more specifically, benevolent. Similarly, if another person's pain is evil, then hating his pain, or desiring its absence and being pained by its presence, is intrinsically good and compassionate. But loving another's evil is the evil of malice, while hating his good is envy.

This account remains reasonably close to everyday thought, which would agree that benevolence involves caring positively about what's good for other people and compassion being against their evil. But it expresses these ideas in a more abstract and therefore explanatory form. If we have the general idea of loving the good, we can collect under it not only the virtue of benevolently desiring another's pleasure but also those of disinterestedly pursuing knowledge, if knowledge is good, desiring just distributions, and more. Going further, we can unify all the virtues under the general heading of attitudes appropriate to their objects, because their orientation matches their objects' value, either positive to positive or negative to negative. This type of unification is explanatory, especially when the abstract principles are intuitively appealing in themselves, as I believe they are. It's attractive just as an abstract idea that loving a good is another good and loving an evil an evil.

At the same time, the account suggests further questions of detail. If attitudes can be more or less virtuous, how does their degree of value depend on their intensity and the degree of value of their object? What about attitudes to evaluatively neutral objects, or indifference to good or bad ones? And how do the different forms of virtuous love, such as desire for absent goods and pleasure in present ones, compare with each other? If we can answer these questions, and in a systematically connected way, that will again deepen our understanding.

This account of virtue is by no means my invention. It was widely accepted among moral philosophers of the late 19th and early 20th centuries,

among them Moore and Ross. And in fact this period's philosophers consistently favoured structural over foundational theorizing. They rejected the more ambitious approach of ancient ethicists such as Aristotle, in part because it was obnoxiously egoistic, and likewise rejected Kant's attempt to derive specific duties from the bare idea of consistent rational willing. I think the closeness of Moore's and Ross's analyses to everyday moral thought, and their resulting clear relevance to that thought, is part of what attracted me to them as an undergraduate and sustained my admiration for them even when I knew they were out of philosophical fashion. In now more explicitly favouring structural analysis I'm returning to that initial attraction.

In recent years there's been much valuable writing in the structural style. In the theory of value there have been more elaborate analyses of equality, desert, and other goods than any given by Moore or Ross, and there have been similar advances in understanding deontological ethics. Though Ross held that there's a more stringent duty not to harm other people than to benefit them, he said little about what the difference between the two consists in. The recent voluminous literature on the doing/allowing, intending/foreseeing, and other distinctions, though not conclusive, has greatly illuminated the philosophical options in this area. It's likewise been emphasized that deontological ethics is structurally agent-relative, so it doesn't tell people to minimize the total number of, say, lies told, but to care specially that *they* don't lie. And despite his effort to remain true to common-sense morality, Ross held that, when other duties are silent, we're required to

maximize the good impartially. Against this, more recent writers have shown that everyday thought is less demanding, allowing agents to care somewhat more about their own good than about other people's and therefore sometimes not to do what's impartially best.

At the same time, however, immense energy has in recent decades gone into foundational projects, which probably have the greatest prestige among philosophers. They include, alongside neo-Aristotelian and Rawlsian views, a grandiose Kantian argument that commitment to specific moral principles is implicit in the presuppositions of all rational agency. Even some of the structural discoveries mentioned above have generated foundational demands: philosophers have required the agent-relativity of deontological duties to be given some separate and deeper justification, or have tried to ground the permission to favour one's own good in metaphysical claims about the self and its points of view.

If, as I have argued, foundational theorizing rarely succeeds, these efforts are unlikely to yield positive results. But they also distract philosophers from what I think would be more fruitful work. Someone who's seeking the deep foundation of morality is less likely to look at the details of what morality says; her attention is somewhere else. And there's a more insidious effect. The more complex morality is, and the more subtle its features, the harder it will be to justify on conceptually separate grounds, just because the justification has to justify more. There's therefore a temptation for foundational theorists to simplify the content of morality, so their

justifications have less to do. (The alternative is to leave the details to ‘true practical principles’ or ‘what the virtuous person perceives,’ but that’s simply to give up on explaining what can be explained by structural means.) In both ways, more interest in grand foundational theorizing leads to less interest in the details of our actual moral scheme.

But those details are fascinating, surprising, and eminently worth study. Our everyday intuitions often make fine-grained distinctions, about either value or duty, whose underlying rationale it’s illuminating to identify. And those intuitions can have more integrity than appears: sometimes an attractive moral claim turns out to entail others on the same topic that initially seemed independent of it. Or there can be unexpected connections between topics, with our judgements about one having the same structure as about another, so the two sets of judgements run in parallel. In all these cases structural analysis can show everyday thought to be more coherent than first appears.

Of course we may not always find this coherence. For some fine-grained judgements, for example about when one may cause harm to prevent greater harm to others, a systematizing principle has proved very hard to formulate. And if in fact there’s no such principle, that may be grounds to question the judgements’ reliability. Or if there is a principle but it has no independent appeal – its only merit is to yield the right particular judgements – that too may be grounds for doubt. And of course we will often encounter disagreements about particular judgements. Everyday moral thought isn’t a

monolith. It contains different elements, which different people emphasize to different degrees, leading to disputes about what's right and wrong in particular cases. So everyday thought contains not just one view but a plurality of partly competing ones. In all these cases, however, structural analysis remains the best response. If our everyday judgements can't be systematized, a serious effort to do so is the only way to establish that fact. And when people disagree, a similar analysis can pinpoint the exact crux of their dispute, increasing the possibility of resolution and at least making clear what the ground of difference is.

It may be objected that structural theory is inherently conservative, confined to systematizing existing moral beliefs rather than changing them in fundamental ways. But foundational work is also often conservative, either in practice or also, as in Rawls, in its explicit intent. And the structural approach is no less able to play a reforming role. Discovering hidden connections between moral judgements can lead us to extend ideas about one topic to another, as when the valuing of human pleasure leads, by analogical reasoning, to a similar valuing of animal pleasure. Structural theory can even identify wholly new abstract principles, intuitively appealing in themselves and demanding radical changes in our moral outlook. That this approach often starts with everyday moral judgements doesn't mean it's tied forever to them. It can introduce new principles, so long as they imply and explain a set of particular moral claims and use concepts not drastically different from those in the claims.

In pursuing the structural approach we should attend to the richness not only of everyday thought but also of the normative theories philosophers have developed. Philosophy often reduces the main contending views on a subject to a small number: in normative ethics typically just utilitarianism, Kantian deontology, and (a recent addition) virtue ethics. But there are vastly more options than those: non-utilitarian consequentialisms such as Moore's, non-Kantian deontologies such as Ross's, and rival accounts of virtue like the one described above. We should also attend to full richness of our philosophical tradition. When contemporary philosophers align themselves with historical figures they tend to pick one from a small group: Aristotle, Hobbes, Kant, and maybe Hume or Sidgwick. But again there are vastly more writers worth reading, including especially, I would say, those early 20th-century figures such as Moore and Ross who developed moral theories starkly opposed to Aristotle's and Kant's.

A recommendation of structural over foundational theorizing is, I recognize, unlikely to have much effect; the latter is too attractive to philosophers. And it would be going too far to say foundational work is never of value. Even when its projects fail, the efforts devoted to them can in various ways be instructive. But surely even those drawn to it need to do structural work first. If they're to give deep justifications of our moral judgements, they must first know exactly what those judgements say. And if there are structural principles that partially explain our judgements, a more

foundational theory should adopt them and incorporate their claims. So even those drawn to the other approach should see structural theory as a necessary propaedeutic.

It doubtless looks self-serving: a philosopher recommends as the most valuable approach to moral theory the one he himself has pursued. But in some cases the causation runs in the opposite direction: a philosopher first feels, even in the inchoate way I did in my undergraduate seminar, that a certain style of theorizing is most fruitful, and is then drawn to practise it by that feeling. And my recommendation is by no means unique to me. I close by quoting a passage from Friedrich Nietzsche, urging a similar approach while directing appropriate barbs at Kant and his followers:

One should own up in all strictness to what is still necessary here for a long time, to what alone is justified so far: to collect material, to conceptualise 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 one has not 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 (Friedrich Nietzsche, *Beyond Good and Evil*, trans. Walter Kaufmann (New York: Vintage, 1966), sec. 186).