

## THE STRUCTURE OF NORMATIVE ETHICS

Shelly Kagan  
University of Illinois at Chicago

If you open a typical textbook on normative ethics, it will have a discussion of what it thinks of as rival theories: there might be a chapter each on, say, utilitarianism, contractarianism, and virtue theory. The assumption seems to be that these three are alternative attempts to answer the same basic questions.

This seems to me exactly incorrect. At least as they are ordinarily understood, these three are addressing three different concerns in normative ethics, and in principle—although almost never in practice—they are completely compatible. Failure to see this, I believe, is due to our failure to have an adequate “map” of the structure of normative ethics. We lack an adequate account of what the various theories in normative ethics are trying to accomplish.

It is not that there is nothing at all like a received view concerning the nature of normative ethics: roughly, normative ethics involves the attempt to state and defend the basic principles of morality. It is concerned with determining which actions are right, which wrong, what is permitted and what forbidden. Similarly, it might be said to treat the basic moral rights, duties, virtues, and so on.

So far as it goes, there is nothing wrong with this account. But it does not go very far. It is like a map of a country which only displays its border. Such a map can serve to distinguish one country from another—as our account can help to set off normative ethics from metaethics—but it gives no significant detail about the internal features of the country. What we lack—what the received view does not give us—is a sense of the major regions of normative ethics and how they are related to one another. Lacking an adequate guide to the structure of normative ethics we can fail to recognize what we are doing when we compare and evaluate the specific normative theories that have been offered.

In this essay, then, I want to try to lay out the basic outlines of a more

adequate map. My goal is not to argue for a particular theory in normative ethics, but rather to try to give a perspicuous way of sorting out some of the most important theories and disputes. If I am successful, the account I offer should not strike you as a description of a new, exotic land or alien terrain. For although we may lack an adequate *map* of normative ethics, as moral philosophers we are not, of course, unfamiliar with the various internal doings of normative ethics itself. So the map should have a familiar feel to it, at least in the sense that we should recognize it as a plausible account of what we have been doing in normative ethics—one that does not distort too much of the field, and is not hopelessly idiosyncratic in its concerns.

As a first step, it seems to me useful to distinguish between two sorts of activities that go on in normative ethics. One thing we try to do is to articulate the various factors that are relevant to determining the moral status of an act. I'll say more about these factors later, but to fix the basic idea quickly, here is an example. Suppose that in order to save the life of someone drowning, I must row out to her in a boat. The fact that the act of rowing out and pulling her in would have a good result—saving a life!—is one morally relevant factor in determining the rightness or wrongness of the act. In saying this, I do not mean to be making the epistemological claim that we *appeal* to the presence of such a factor in *deciding* whether the act is right or not, although in fact I think that this is often the case. Rather I mean to be making the ontological claim, that the interplay of the various morally relevant factors is what makes it *be* the case that the given act is right or wrong. It seems quite plausible to think that goodness of outcome is one morally relevant factor in this sense.

It certainly need not be the only factor. Suppose that the only boat at hand is not mine, so I must steal it, if I am to rescue the drowning person. This might plausibly be thought to be a second morally relevant factor—that is, the fact that I must violate someone's rights, in particular, property rights. Whether, morally speaking, I should still take the boat or not depends on which of these two factors outweighs the other. And there might be other factors as well. The person drowning might be my wife, and so it might be thought that this gives me a special obligation, stronger than the general obligation to save a mere stranger. And so on.

Real life cases might get fairly complex. From a moral point of view, the complexity lies in the fact that several distinct morally relevant factors come into play. Whether a given act is forbidden, or required, and so on, is a function of these various factors. We might think of them, loosely, as right-making or wrong-making factors, although we should not assume in the absence of a worked out theory that an adequate classification will be that simple.

The task of articulation of the normatively relevant factors is an astonishingly complex one. There is considerable debate over which purported factors are of genuine relevance. Even when there is some agreement that a particular

factor belongs on the list, there remains the difficult and contentious task of demarcating its precise content and contours. And beyond this, there is the further task of determining how the various factors interact, and what outweighs what in cases of conflict.

Much of what takes place in contemporary moral philosophy is an attempt to work out answers to these questions. It is theory building at the level of normative factors. Many articles in normative ethics are primarily concerned with the precise specification of some part or parts of one or more normatively relevant factors, or with determining the outcome of the complex interactions of the given factors.

Such investigation of the normatively relevant factors is one of two major activities in normative ethics. The second activity is this. We offer rival theories of what I call the *foundations* of normative ethics. Roughly, we propose and evaluate alternative “devices” or “mechanisms” that purport to generate and thus explain the favored list of normatively relevant factors. Once again I will say more about this below, but to get a few examples on the table will be helpful: Contractarian theories claim that the correct list of moral rules, or morally relevant factors, consists of those that would be chosen by a group of rational bargainers, with various motivational and informational constraints. Contractarianism is thus a foundational theory, in my sense. (Strictly, it is of course a family of theories, but that won’t concern us for the moment.) An adequate contractarian theory would provide an explanation and justification for the specific normative factors generated by the theory.

Another foundational theory (or family of theories) is rule consequentialism, which claims that the correct rules or factors are those that would lead to the greatest amount of good, that is, the best results overall. Yet another theory is the ideal observer theory, which claims that the correct list is that which would be endorsed by a suitably characterized observer, for example, one with full information, and impartial benevolent concern for all parties.

In calling theories of this sort “foundational” I do not mean to be claiming that beliefs about such theories are more certain than intuitions about the normative factors, nor do I mean to suggest that these theories provide some sort of bedrock and do not stand in need of further justification. Claims of these sorts are common in theories of epistemology or justification that often go under the label “foundationalism”. I have no sympathy for them in ethics, or elsewhere. I simply mean to be claiming that those who offer foundational theories in my sense typically take their favored foundational theory to offer an explanation and justification of their favored list of normative factors.

In principle, a foundational theory—at least if worked out—will yield a list of normatively relevant factors. If the foundational theory itself has any independent plausibility, then the fact that it yields the specific factors that it does provides something of a justification for the claim that those factors are indeed the genuinely relevant factors. But of course the very fact that

the foundational theory yields plausible factors also provides support in turn for the plausibility of the foundational theory. So the support can go in both directions. That is one of the reasons that I think that the various epistemological foundationalist views should be rejected.

My aim here, however, is not to discuss the metaethical questions of moral epistemology, but simply to distinguish the distinct sense in which I think of theories such as contractarianism or rule consequentialism as foundational: a contractarian thinks that the rules selected by the specified bargaining process are the morally relevant rules, and he presumably thinks that it is the very fact that they would be so selected that explains or is the ground of their being the correct or genuinely relevant rules. Similarly, a rule consequentialist thinks that it is the optimal rules that are the genuinely relevant rules, and she thinks that what makes them so is the very fact that they are indeed the optimal ones in terms of their results. And so on for the other foundational theories. Each foundational theory not only yields a favored set of normative factors, it purports to explain and justify the relevance of those factors.

By providing a mechanism which generates or selects some candidate normative factors and not others, a foundational theory strives to settle the debate over which factors have genuine moral relevance, and which are spurious, or of merely derivative significance. Ideally, an adequate foundational theory would go further. It would help to settle disputes over the precise content and contours of the favored list of factors. And it would illuminate how the various factors interact in determining the moral status of an act, explaining which factors outweigh the others in cases of conflict, and why. It would in effect provide and vindicate the tradeoff schedule in complex cases involving conflicting factors.

Of course in practice foundational theories are virtually never worked out in this kind of detail. Often in fact there is little more than a hand or two waved in the direction of showing that the given foundational theory will indeed yield anything like the list of normative factors that we are independently inclined to accept. But in principle, at any rate, one of the advantages of having an adequate foundational theory is that it would provide us with guidance in our articulation of the normatively relevant factors.

As a first approximation, then, normative ethics has these two different kinds of components. There are alternative theories about the morally relevant factors. And there are alternative theories about the foundations of normative ethics.

One of the things that is of interest is the extent to which these two activities can be—and are—conducted largely independently of one another. People frequently attempt to fine-tune our understanding of the normatively relevant factors without appealing to any particular foundational theory. (They may appeal, instead, to intuitions about various hypothetical examples.) And at

least in part, people often evaluate alternative foundational theories by considering what independent rationales, if any, can be provided for the various theories.

This is not to say that people never attempt to settle questions about factors by considering the implications of some favored foundational theory. And I certainly do not mean to deny that people frequently evaluate foundational theories by considering their implications at the level of factors. But for all that, the two activities have a surprising degree of independence.

The relative independence surfaces in a second way. Two people might agree about the list of normative factors, and yet disagree about the correct foundational theory. After all, it might be the case that two distinct foundational theories would generate the same list of factors. Or, at any rate, two people who agreed about the relevant list might still disagree over which foundational theory would in fact generate that list. Similarly, two people might agree that a given foundational theory is the correct one, while still disagreeing over which factors are generated by that theory.

This distinction between discussion of normative factors and discussion of foundational theories seems to me one of the most important divisions in normative ethics. But I said that thinking of normative ethics in terms of these two activities is only a first approximation. For it seems to me helpful to distinguish as well a third activity, involving debate over what I will call *evaluative focal points*. However, I want to postpone a discussion of these focal points for the time being, since what I have in mind is best brought out in the context of a more careful look at the foundational theories. And before turning to that, I first want to say something about the basic types of normative factors.

As I see it, the intuitively plausible normative factors fall under four main heads. (Perhaps it would be more accurate to say that I think that they can be usefully grouped in this way; other classificatory schemes are certainly possible.) The four major types of factors are as follows: (1) overall goodness of results; (2) general constraints; (3) special obligations; and (4) options. Although I do not intend to go into great detail about any one of these groups, I want to say enough—if only by way of offering some examples—to give a sense of what I have in mind for each.

As I have already indicated, it seems quite plausible to suggest that one morally relevant factor in determining the moral status of an act is the outcome of that act—that is, whether the act leads to good results overall, or better results than those of the alternative acts available to the agent. But this factor is itself a function of more than one subfactor. (Incidentally, I think it is mostly a matter of convenience, in any given case, whether we speak of “groups” of factors or, alternatively, of various “subfactors” coming together to determine some larger “single” factor.)

To begin with, most would agree that the overall goodness of an outcome

will turn, at least in part, on the levels of well-being of the various individuals who are part of the given outcome. So one thing needed by normative ethics is an adequate theory of individual well-being. Clearly, however, since typical cases involve more than one individual, there will also have to be some principle governing the aggregation of facts about the welfare of the various relevant individuals into overall evaluations of outcomes with regard to welfare. Here we need to settle whether what matters is the total amount of well-being, or the average level, or something captured by a still more complex formula.

But matters obviously do not stop there, for it is often suggested that the *distribution* of well-being is itself a relevant factor in determining the overall goodness of an outcome. Egalitarians will claim that it matters directly whether the welfare is distributed equally, and will presumably sometimes be willing to settle for a smaller total amount of well-being in exchange for a more egalitarian distribution. And this will be rejected by those—the welfarists—who think that distribution has no intrinsic importance (although of course it is often instrumental to an increase in welfare). Furthermore, even if we agree that distribution matters, there is still room for disagreement over the correct distributive principle: for example, should our concern truly be with equality per se, or should we instead have a special concern for the worst off? And even among strict egalitarians there will still be the need to establish the precise tradeoff schedule between welfare and equality.

Still other good-making factors might be *impersonal* in the sense of not being functions of well-being or its distribution at all. Perhaps the existence of beauty, or of knowledge, might directly contribute to the overall goodness of a state of affairs (that is, above and beyond any contribution it might make to increasing welfare). And other more exotic candidates have been entertained as well.

There are of course many other factors that might plausibly be claimed to have direct weight in determining the overall goodness of an outcome, such as desert, fairness, or entitlement. But since all I am trying to do is to give a sense of what sorts of factors fall under each of the four major headings, I hope that the examples I have already discussed will suffice.

In trying to articulate a theory of the overall good we face the same three tasks here as elsewhere. First, there is disagreement over which of the proposed factors have genuine weight: for example, does distribution genuinely matter directly in determining the goodness of an outcome, or are we simply misled into thinking so by the fact that distribution typically affects the amount of well-being? Second, there is disagreement over the precise specification of the relevant factors: for example, assuming that distribution does have genuine weight, should our concern be with equality or with, say, maximin (which gives lexical priority to the worst off)? Third, there is disagreement over the tradeoffs between distinct factors: for example, assuming

that gains in equality can sometimes outweigh gains in total well-being, exactly how much of the former outweighs how much of the latter?

As I have suggested, theory construction and evaluation in this area— as with the other major groups of factors, and the theories of their interplay— can be conducted with a surprising degree of autonomy. Considerable debate over the nature of good outcomes takes place without recourse to particular foundational theories. At the same time, I find it implausible to think that the debate can be concluded autonomously. Appeal to an adequate foundational theory may well be necessary to answer questions about the precise content or weight of various proposed factors; and at any rate it seems inevitable if we are to have a satisfying explanation of why the normatively relevant factors are indeed morally relevant in the first place.

Now most of us believe that although the value of the outcome of an act is indeed one relevant factor in determining the moral status of that act, it is not by any means the only relevant factor. Intuitively, certain acts seem to be forbidden even if the results of performing those acts would be good overall. This brings us to the next group of factors, that of general constraints.

By constraints, I have in mind the types of features that have recently been discussed under such headings as side-constraints, deontological restrictions, and agent-relative prohibitions. Most of us believe, for example, that there is a prohibition against harming people. The fact that my act will harm someone is generally taken to be a morally relevant factor, a factor that generates moral reason not to perform the act. This factor is typically thought to be so powerful that it virtually always outweighs other relevant factors, in particular the possibility of promoting good consequences. I cannot kill an innocent person, say, even if this is the only way to save two other innocent people from being killed. When we say that there is a prohibition against harming, we are expressing the belief that this factor generally outweighs any opposing factors.

It is worth noting that one can give weight to a constraint against harming without viewing that weight as absolute. One might hold that if *enough* is at stake—for example, the avoidance of nuclear war—then promotion of the overall good will indeed outweigh the constraint against harming. In short, constraints can have finite thresholds. And in principle, these thresholds might be fairly modest; the prohibition against killing might be overridable, say, even in order to save as few as five lives. But in fact it seems likely that anyone who hopes to capture our intuitive beliefs on these matters will hold the thresholds to be extremely high. Some—the absolutists—would argue that the threshold is infinitely high, so that harming the innocent is never permitted. But one can be more moderate on this score without abandoning the view that this constraint is indeed a genuinely relevant factor, and outweighs other factors in many or most cases.

As always, the precise specification of the constraint against harming is

a controversial matter. For example, exactly what type of act is forbidden by the constraint in question? Is it the *doing* of harm that is especially morally offensive (as opposed to merely allowing harm)? Or is it rather the *intending* of harm as a means (as opposed to merely foreseeing harm as a side-effect)?

Another familiar debate concerns exactly which setbacks constitute harms: Should mental discomfort or offense be included along with bodily injury? What about invasions of privacy? Or appropriation of one's property? To a certain extent some of this debate is merely terminological. Nothing prevents us from listing distinct constraints for different types of setbacks. But much of the debate can be understood as involving the substantive issue of which setbacks are genuinely protected by a constraint at all, and for those that are protected, whether the grounds of that protection are of fundamentally the same sort.

This last question reemerges when we consider plausible constraints that are typically not subsumed under the rubric of a constraint against harm. We might well want to include a constraint against lying, or a constraint against paternalistic interference with another's autonomy. To my mind, at least, there are interesting questions about whether or not these are best understood as simply being particular forms of harm after all. And if they are best understood in this way, then the question arises whether there are indeed any general constraints other than various versions of the constraint against harming.

In contrast to such general constraints, special obligations involve the particular duties and responsibilities that we have by virtue of our own individual circumstances and history. The most familiar and relatively uncontroversial examples of such special obligations involve moral requirements generated by acts freely undertaken in the past: promises generate special obligations to keep those promises, and so one factor relevant in determining the moral status of an act is the fact that the act is the keeping (or breaking) of a promise. Similarly, the principle of fair play may generate special obligations to do my fair share in sustaining a joint practice from which I have knowingly and willingly benefited in the past. And I may have special obligations toward my infant daughter, generated in part by the fact that I am directly responsible for her existence.

But mention of this particular family obligation brings to mind the thought that not all role-related obligations can be traced to previous free acts of the agent. Even within the family it might be argued that one has special obligations toward one's parents, although the relation to one's parents is not one freely undertaken. Similarly, it might be argued that one's determinate social role generates a host of specific obligations, even though one may not be completely free with regard to one's social role. (Political obligations might sometimes work like this.) Now it can of course be argued that many of these

involuntarily generated special obligations are best explained as instances of the obligation of gratitude; but since the obligation of gratitude itself can have this involuntarily entered character, the basic point remains. Against all this, there is of course considerable contemporary scepticism about whether there are any genuine special obligations not grounded in free acts of the agent.

There is one other obligation that I think merits our attention. Many of us feel that there is a fairly strong obligation to rescue particular individuals when we find ourselves facing the immediate opportunity to do so. The fact that we do not consider ourselves under a similarly strong obligation to aid unnamed persons who are, say, starving at the other end of the globe, provides some evidence that we are indeed describing what we take to be a special obligation, rather than merely an instance of the general (“weak”) obligation to promote better consequences overall (for example, by saving lives). Even more evidence emerges when we consider that the opportunity to directly rescue a known individual is often taken to generate an obligation which must be met even if more good could be done elsewhere by diverting the resources necessary for the rescue.

It is perhaps worth observing that the obligation to rescue appears to be a special obligation that is not grounded in voluntary acts of the agent. Of course general constraints have had this “involuntary” character all along, so friends of the thesis that all special obligations must be “voluntary” might suggest that the obligation to rescue should actually be classified as a general constraint.

Now I am not in fact convinced that there is any deep distinction to be found between the category of general constraints and that of special obligations. For some purposes it might be more illuminating to collapse the two into one. Even a constraint against lying, or against harming, for example, will only yield a specific determinate obligation with regard to the performance of a specific act in the right kinds of circumstances; so the distinction between “general” constraints and “special” obligations may only be one of degree, rather than marking a sharp break. Or perhaps we should simply move directly to marking the voluntary/involuntary distinction, or some other distinction (such as that between negative and positive obligations). But these are questions I will not attempt to pursue here.

The last major type of factor is that of options, or what have elsewhere been called agent-relative permissions, or agent-centered prerogatives. Most people believe that even if an act would have good results (one morally relevant factor) and even if the act would not involve the violation of any general constraints or special obligations (two other types of morally relevant factors) the act might nonetheless *not* be morally required. For agents have the option of pursuing their own personal projects and interests, even if greater good overall could be done in some other way. Many morally praiseworthy

acts are above and beyond the call of duty—meritorious, but not required.

What exactly is the factor at work here? The matter is not uncontroversial, but I think the most likely account is that the relevant factor is an appeal to the cost to the agent that can be involved in sacrificing the pursuit of her projects for the sake of the overall good. If the cost to the agent is significant, this is one more normatively relevant factor in determining the moral status of the act: an act might be optional (permitted but not required) if it would involve a significant sacrifice on the part of the agent, even though an act with similar results might well be required if the cost to the agent would be slight.

Here too the need remains for considerable further discussion. We need to specify which other factors can be outweighed by an appeal to cost, and to what extent. The size and nature of the options that we have will be a function of this tradeoff schedule. Most of us believe, for example, that although we have options to allow harm, we do not have similar options to do harm: I can devote my income to my family, rather than to famine relief, but I cannot *kill* to provide a comparable income for my family. Thus cost appears to outweigh the possibility of attaining good results, but is itself outweighed by general constraints. Even the option to allow harm, however, is not unlimited: when enough is at stake, the appeal to cost to the agent may indeed be outweighed by a consideration of results.

These, then, are the four basic types of factors that come into play in determining the moral status of an act: results, constraints, special obligations, and options.

It may be worth explaining, parenthetically, why the concept of rights has not appeared in my discussion of the basic types of normative factors. The main reason is simply that I think that talk of rights frequently lacks sufficiently determinate content to be helpful in analysis. A related reason is that talk of rights is often meant to bring in more than one of the types of factors that I have been distinguishing. When we say that someone has such and such a right, often we mean to be drawing attention to what I have called constraints or special obligations. But sometimes we mean—in addition, or instead—to be drawing attention to the existence of an option. A right may involve one of these elements without the other, and so it seems to me clearer to lay out the possible elements directly.

The picture of the normative factors that I have sketched is not, I trust, an unfamiliar one. Whether an act is required or not turns on such matters as what the results of the act would be, whether it involves keeping or breaking any general constraints or special obligations, and what the cost to the agent would come to. As I have explained, one of the major activities of normative ethics involves the specification of the various factors and subfactors, debates over which of these have genuine weight, and attempts to determine which genuine factors take precedence in cases of conflict.

Some of these debates are legitimately seen as in-house debates over details. This might be true, for example, for disagreements over the correct version of maximin, or disagreement over the precise location of a finite threshold for a constraint against doing harm. In other cases the disagreement will be far more radical. The most familiar example of a large-scale debate is that between consequentialists and deontologists.

Consequentialists hold that the only normative factor that has any genuine weight in its own right is that of the overall goodness of the results. (This still leaves room for debate among consequentialists over the nature of the good. If welfare is taken to be the only good-making feature of outcomes, then the resulting normative view is utilitarianism.) Other factors may have derivative significance, insofar as attending to them rather than directly to consequences can often contribute to making the results better or worse; but no other factor has any direct significance. At this level of normative factors, then, a *deontologist* should be understood to be anyone who believes in the direct significance of further factors, in particular general constraints or special obligations.

What this means is that rule utilitarianism is *not* from this point of view a utilitarian or even a consequentialist theory! Rule utilitarianism (and, more broadly, rule consequentialism) is, rather, a *foundational* theory which attempts to explain why other normative factors than the goodness of results have direct normative relevance in determining the status of an act. Thus at the level of the normative factors themselves, rule utilitarianism does not support a utilitarian theory, but rather a deontological one.

Similarly, although contractarianism—another foundational theory—is typically thought to underwrite a deontological normative view, at a sufficiently general level of description this cannot be assumed. Some versions of contractarianism end up supporting utilitarianism at the level of the normative factors.

This debate over who is a deontologist and who a utilitarian or consequentialist takes a rather different cast if we turn to the foundational theories. For consequentialism reappears as a family of foundational theories as well. At the level of normative factors, as I have just explained, consequentialism is the theory that goodness of results is the only directly relevant factor in determining the status of an act. But as a foundational theory, consequentialism is the view that whatever the genuinely relevant normative factors may be, the *ground* of their relevance ultimately lies in their connection to the promotion of the overall good. As a foundational theory, that is, consequentialism holds that an explanation of the significance of the normative factors must ultimately be in terms of an appeal to the good.

Classical consequentialists at the normative level often justified their views by appealing to consequentialism at the foundational level. For some, the conception of the good has been metaphysical or platonic; for others it has

taken a more naturalistic form. But it is important to recognize that one can be a consequentialist at the factoral level without favoring consequentialism at the foundational level. As I have just noted, one might prefer instead to ground one's factoral consequentialism in a foundational contractarianism, or for that matter in an ideal observer theory, or some other alternative.

And it is equally important to recognize that one can be a consequentialist at the foundational level without believing that this supports consequentialism at the factoral level. Rule utilitarians (or, more broadly, rule consequentialists) should be understood in just this way. They ground the moral relevance of the various normative factors in terms of the good—saying that those factors that would (if acted on) lead to best results have direct weight—but they do so in a manner that generates a nonconsequentialist theory at the factoral level. So there is a sense in which rule utilitarians deserve to be called utilitarians after all—namely, at the foundational level. But this is compatible with insisting that there is another sense in which they should be called deontologists—namely, at the level of normative factors.

This prompts the question of whether there is any useful notion of being a deontologist at the foundational level. The thought that initially suggests itself is that foundational deontologists are those who think that an adequate foundational theory will select normative factors not simply with an eye to the promotion of the good, but rather (in addition, or instead) on the basis of features that are somehow more “intrinsic”. But I am not sure how to spell out the second half of this suggestion. After all, many would take contractarianism to be a deontological theory at the foundational level, but I am hard pressed to think of any interesting sense in which contractarian approaches turn on especially “intrinsic” properties—whether of acts, or of the normative factors. So I am inclined to think that deontology at the foundational level should be understood simply in terms of the negative, first half of the suggestion. That is, deontological foundational theories are those that select the normative factors on the basis of something other than, or more than, the promotion of the good.

We would come close to this by saying that foundational deontologists are those that offer a foundationally nonconsequentialist theory. But this would be too wide a definition. For consequentialists appeal, in particular, to the promotion of the *overall* good, and we presumably also want to exclude as nondeontological foundational theories—such as egoism—that appeal instead solely to promotion of one's *personal* good, or self-interest. So we may prefer the following. Call a foundational theory *teleological* if it appeals solely to promotion of the good, however this is construed—whether personal good, or overall good. We can then say that deontological foundational theories are those that are nonteleological.

I will not attempt to give a more fine-grained system of classification for the foundational theories than this. Unlike the four major divisions of the

normative factors, I do not know of any typology for the foundational theories that is comparably useful. At any rate, the task of producing an illuminating classificatory scheme for the foundational level will be made more difficult by the fact that foundational theories can combine with one another, mixing their elements.

Consider, for example, universalization theories, which claim that the correct normative factors are those that can be appropriately universalized, that is, those for which it is possible to have a world in which everyone acts on the proposed factors without something “going wrong”. Now in pure versions of such theories, the thing that can “go wrong”, and that universalization tests for, is this: for illegitimate factors, the concept of such a world of universal compliance entails a logical or practical contradiction. But some advocates of universalization theories think that it is insufficient to rule out only those factors that when universalized generate a contradiction. They suggest, instead, that a proposed factor should also be ruled out if its universalization would yield unacceptable results.

Here too, of course, there is room for disagreement as to the type of results that are relevant: some ask whether universalization would lead to bad results overall; others whether it would lead to bad results in terms of the agent’s personal good, and so on. For our purposes, however, the point to note is this. However “bad results” is understood, a theory of this latter sort seems to combine the idea of universalization with teleological elements. It is a hybrid—combining elements from two “pure” theories to form a new theory. (Is this new type of theory a deontological one? It depends on whether assessing results from the standpoint of universalization in this way means that justification is no longer simply a matter of the promotion of the good. I think that people’s intuitions will differ on this.)

This possibility—of combining elements from more than one foundational theory to form a distinct foundational theory—should not be confused with pluralism at the foundational level. And pluralism at the foundational level should in turn be distinguished from pluralism at the factoral level.

This last point is by now a familiar one. A pluralist at the factoral level believes that more than one normative factor has weight in its own right. Factoral consequentialists deny this; but most others accept it. As we have already seen, however, one can be a pluralist in this sense and still believe that there is a single foundational theory that justifies and explains the relevance of this plurality of normative factors. Rule consequentialists and contractarians, for example, typically hold this view.

But a pluralist at the factoral level need not believe that there actually is a single foundational theory that generates all of the relevant normative factors. (Indeed, scepticism on this score sometimes leads people to deny that there exists any foundation at all for the normative factors.) For there is another possibility: one could be a pluralist at the foundational level as well—

holding that certain of the factors are grounded in a particular foundational device, while other factors are grounded in some second device, and still other factors are grounded in yet another device. On such a view pluralism at the factoral level would not be underwritten by monism at the foundational level; the justificatory picture would be more like a patchwork or a quilt.

There is of course a sense in which any given version of foundational pluralism will itself be one more foundational theory—one compounded out of the original foundational devices. Nonetheless it seems to me useful to distinguish between this case—where combination is largely a matter of mere conjunction—and other cases (like that of the hybrid of teleology and universalization) where the elements being combined are “synthesized” so as to form a foundational theory that intuitively possesses some internal unity.

If we do have a case of pluralism at the foundational level, the pluralism may well be ultimate and irreducible. But I think that even here it need not be. On some views, at least, there may be—at a “deeper” foundational level still—some single foundational theory that grounds and explains the legitimacy of the plurality of more “superficial” foundational theories.

Indeed, this possibility that a given foundational theory might be underwritten or supported by a distinct foundational theory seems to me to arise even in cases where there is only a single foundational theory at work on any given level. Thus, rather than directly generating the various normative factors, the machinery of a given foundational theory might generate—and thus support—one of the other foundational theories. Presumably, this process could be repeated. For example, egoism might generate contractarianism, which might in turn generate rule consequentialism before, finally, generating the relevant normative factors. There are in fact theorists that have held such multi-leveled foundational views, and it can be illuminating to see them in this light.

Rather than pursue such exotic possibilities any further, however, I want to return to an earlier point. As we have seen, consequentialism at the foundational level can support deontology at the factoral level, if the foundational theory takes the form of rule consequentialism. We thus need to distinguish at the foundational level between rule consequentialism and act consequentialism (a distinction familiar, of course, from the debate between act and rule utilitarians). Act consequentialists think that the given act should be evaluated directly in terms of *its* consequences; thus it emerges at the factoral level that the *only* factor with genuine weight is that of goodness of results. Rule consequentialists in contrast think a given act should be evaluated in terms of its conformity to the optimal set of rules (where the rules are themselves evaluated in terms of their results) rather than directly in terms of its consequences; this is what generates the possibility that other factors may have genuine weight at the factoral level beyond goodness of results.

The distinction between act and rule consequentialism might be put this

way: both share the foundationally consequentialist thought that moral justification must ultimately be in terms of promotion of the good; both use impact on the promotion of the good as the method for evaluating their particular favored objects, whether acts or rules. But the two approaches differ as to what kinds of objects provide the *primary* evaluative focal point. Rule consequentialists select *rules* as their primary evaluative focal point; they then evaluate acts in a secondary or derivative way, in terms of the directly evaluated rules. In contrast, act consequentialists select *acts* as their primary evaluative focal point, evaluating them directly. (This last claim is potentially slightly misleading, but will do for the moment.)

Although it has not been widely appreciated, this debate concerning the proper primary evaluative focal point is not at all limited to foundational consequentialists. It seems to me, in fact, that for all, or almost all, of the major types of foundational theories there are distinct versions of that foundational theory which differ in terms of their choice of primary evaluative focal point.

Consider, for example, the ideal observer theory. The basic idea of such a theory is that moral justification is ultimately in terms of the choices or preferences of a suitably ideal observer. An act version of this approach would hold that a given act is morally right provided that the ideal observer would approve of or favor this act. The foundational machinery of the theory—the device of the ideal observer—is here being applied directly to the evaluation of acts. (Incidentally, it might still turn out, of course, that the ideal observer is sensitive to a number of features of acts, and so nothing here yet rules out the possibility that at the level of normative factors there will be several factors with genuine weight. Until the details of the theory are filled in, there is simply no way to tell.)

The rule version of this approach does not apply the device of the ideal observer directly to the evaluation of acts, but rather to the evaluation of rules. It holds that a given act is morally right provided that the act conforms to the various rules favored by the ideal observer. (Once more, it may be worth noting that until the theory is filled in, we simply cannot tell what would emerge at the level of normatively relevant factors; in particular nothing yet indicates whether a rule ideal observer theory would give genuine weight to more than one normative factor.)

Here is another example of a foundational theory where we face the same choice between act and rule as our primary evaluative focal point: egoism, or the self-interest theory, which holds that moral justification must ultimately be in terms of the promotion of the agent's self-interest, or personal good. Ethical egoism is almost always understood as having acts as the primary evaluative focal point. So construed—that is, as act egoism—it holds that an act is morally right provided that the given act best promotes the agent's self-interest. As has been widely observed, this is actually a rather unpromising

suggestion for the foundations of normative ethics; at the level of normative factors, no factor will emerge as having genuine weight except for the agent's own well-being. What has been largely (although not completely) overlooked, however, is the possibility of *rule* egoism, which holds that an act is morally right provided that it conforms to the rules, conformity with which on the part of the agent would best promote the agent's self-interest. This is actually a much more promising suggestion, for it seems at least possible that the optimal set of rules from this point of view will give genuine weight to many of the normative factors that we intuitively think of as having moral force.

Unlike egoism, where it is generally assumed that the primary evaluative focal point will be acts, when it comes to contractarianism people have generally assumed that the primary evaluative focal point will be rules. Standard versions of the theory hold that an act is right provided that it conforms to the rules that would be selected by the suitably specified bargainers. The contractarian machinery is used for the direct evaluation of rules; acts are only evaluated indirectly. But nothing rules out the possibility of an act contractarian approach, which would use the contractarian machinery to evaluate acts directly. Such a theory would hold that a given individual act is right provided that the hypothetical bargainers could agree to it

Acts and rules do not exhaust the list of plausible primary evaluative focal points. Motives belong on the list as well. Thus we have motive consequentialism, which asks which set of motives would lead to the best results overall, and then derivatively evaluates acts in terms of whether a person with the ideal set of motives would perform them. We might similarly consider the merits of motive contractarianism, motive egoism, and so on for the various other foundational theories.

Other possible primary evaluative focal points include institutions, norms, character traits, and intentions. For each of these, I think, we can construct a variant of any given foundational approach. Abstractly described, the general idea is this: the machinery of the given foundational theory is used to directly evaluate instances of the favored type of evaluative focal point, and these direct evaluations are used to provide derivative evaluations of instances of the other types of focal point.

This choice of primary evaluative focal point can have an influence—potentially, a rather significant one—on what exactly is generated at the level of the normative factors. For the same basic type of foundational device may well support different lists of normative factors, or different weights for those factors, depending upon the particular evaluative focal point chosen. For example, as we have noted, act consequentialism supports the view that the only normative factor with direct weight is the overall goodness of results, whereas rule consequentialism appears to support the conclusion that other normative factors have independent significance as well.

The choice of evaluative focal point will also influence the precise *form* that the favored normative factors will take. Roughly speaking, on the rule version of a given foundational theory, the normative factors that are selected will emerge—naturally enough—in the form of rules or principles that are to be conformed to. On the motive version, the normative factors will take the form of various motives that are to be had and acted upon. On the act version, the normative factors will appear as positive and negative “values”—features of acts that are to be displayed (or not displayed) in one’s actions. And so on, for the various other focal points.

It will be noted that I have been writing as though a fully specified theory at the foundational level must always choose one particular focal point and elevate it to the special status of *primary* evaluative focal point, evaluating the other focal points only indirectly, in terms of the primary one. But this is not in fact the only option. One could have a theory that refrained from selecting any of the focal points as primary, evaluating all of them, instead, directly.

Given this possibility, we need to distinguish between *act* theories—strictly so called—which directly evaluate only acts (and evaluate the other focal points only indirectly), and what might be called *direct* theories, which directly evaluate not only acts but also the other focal points as well. (I think, incidentally, that what many people have in mind by the position that is generally referred to as “act consequentialism” is actually direct consequentialism. Of course, not being aware of the distinction, most people have not had either theory determinately in mind.) The differences between these two types of theories can be rather subtle, and I won’t explore them here. But it should, I think, be borne in mind that the choice of focal point includes this possibility of taking a direct approach to all of the focal points, and not only the possibility of selecting one particular focal point for primacy.

Now from one point of view, of course, the specification of the evaluative focal point is simply one way among many in which any given basic type of foundational theory needs to be elaborated if it is to assume a more determinate form. After all, even if we know that we are dealing with act contractarianism, say, we still ultimately need to specify the motivation of the bargainers as well as what kinds of information they possess. Similarly, even if we know that we are dealing with a universalization theory, we still need to specify whether it is logical contradiction that we are testing for, or practical contradiction, or some particular form of “bad result” But for the most part, the particular details that need to be filled in vary from one basic type of foundational theory to the next.

In contrast, as I have been trying to bring out, the choice of focal point is one that each foundational theory faces. For this reason it seems to me helpful to think of the chosen focal point as one of two basic components—the other is the given type of foundational mechanism or device—that together

determine the fundamental character of the theory at the foundational level. (Of course, for all that, it remains natural to think of foundational theories that differ only in terms of their choice of evaluative focal points as variants of the same basic type of theory. Accordingly, in those contexts where a contrast is being drawn between the choice of foundational device and the choice of focal point, I will continue to use the expression “foundational theory” to refer to the basic foundational device itself.)

If the choice of evaluative focal point is indeed duplicated by each of the foundational theories, then it seems plausible to think that the focal points may deserve study in their own right. This is not to suggest that discussion should or could take place in complete isolation from discussion of the foundational theories, but only that a certain degree of autonomous investigation might be illuminating, just as it has proven to be in the case of the normative factors or in the case of the foundational theories.

For the most part, however, discussion of the various evaluative focal points has taken place only within the context of examining one or another foundational theory. What I am suggesting is that we would do well to bring the evaluative focal points into the philosophical light, and investigate them directly.

One obvious way in which there is some benefit to be had from thinking of the focal points in this way is that it suggests possible normative theories that might otherwise be overlooked. The example of rule egoism illustrates this point.

Secondly, at least some of the arguments that have been offered for running a given foundational theory in terms of one or another evaluative focal point seem to stand on their own, and thus may carry over to the choice of focal point for other foundational theories. Alternatively, investigation might reveal what it is about a given focal point that makes it a plausible choice for one particular foundational theory, but implausible for another.

Here is an example. There has been considerable debate over whether or not rule utilitarianism provides a genuine alternative to act utilitarianism. The thought is that since the rules are to be evaluated in terms of the goodness of results when the rules are acted on, the best rule will simply turn out to be act utilitarianism, which (after all) already directs agents to bring about the best results! If rule utilitarianism does indeed “collapse” into act utilitarianism, then this prompts the question whether for *all* foundational approaches, the given rule version similarly collapses into the corresponding act version. Direct investigation of the properties of rules and of acts at work in this argument should help clarify for which foundational theories, if any, this objection is sound.

Another possibility is this: it might turn out that the arguments for some particular evaluative focal point are sufficiently powerful and general that they succeed in establishing the superiority of that focal point—independently

of being committed to any particular foundational theory. This parallels a phenomenon we have observed previously, in which claims at a given level of normative ethics can apparently be made while maintaining neutrality concerning the other levels. The contractarian, for example, holds that the correct *foundational* theory is the contractarian one; in itself this claim is neutral on the question of the correct account at the level of normative factors. And utilitarianism, I have suggested, should typically be taken as a claim at the level of the normative factors, a claim that is in itself neutral about the correct foundational approach. Similarly, then, one might argue that a given focal point is the uniquely correct primary evaluative focal point, while at the same time maintaining neutrality concerning the correct foundational approach or the genuine normative factors.

Virtue theory, I suggest, is best understood in this way. Particular virtue theorists, of course, may well be committed to a host of claims concerning the foundational or factoral levels. But what virtue theories have in common, as a class, is the claim that it is virtues, rather than acts or rules, that should be our primary evaluative focal point. Virtue theorists thus provide the single most important exception to my earlier observation that the study of focal points has received little direct attention. The nature of a virtue, at any rate, has received considerable attention. And much of that discussion has been concerned at the same time with arguing for the relative superiority of virtues over the other potential primary evaluative focal points.

I have now, I hope, made good on my promise to explain why I believe that it is a mistake to think that utilitarianism, contractarianism, and virtue theory are all rival attempts to do the same thing. Let me quickly review the outlines of my answer.

It is helpful to think of normative ethics as having two major levels--foundational and factoral. But investigation at the foundational level can itself be further subdivided into discussion of foundational devices and of focal points. Foundational devices are combined with, or directed toward, a choice of evaluative focal point and thereby generate the favored list of normatively relevant factors. In principle, a complete theory of normative ethics will include all three components, and will show the interconnections between these elements.

But there is something to be learned from incomplete theories as well, and a theory might confine itself to claims involving only one of these three areas. On at least one plausible construal, utilitarianism is a theory about the normative factors, contractarianism is a theory about the foundational device, and virtue theory is a claim about the evaluative focal points. Far from being incompatible or rival views, one could in fact endorse all three!

This is not to deny that once spelled out, choices in one area may constrain one's choices in another area. Perhaps there is no plausible way to fill in the details of the contractarian approach, for example, and still hope to generate

goodness of results as the only genuinely relevant normative factor. But many claims in normative ethics are initially offered at a fairly high level of generality, and tensions and incompatibilities may only surface after considerable investigation. At any rate, the point remains that we might be prepared to make a claim concerning one of these areas, without yet knowing for sure where we stand with regard to the others.

Obviously enough, the picture I have been drawing of the structure of normative ethics is incomplete in many ways. It will be apparent, for example, that I have made no attempt at all to be complete in my listing of potential evaluative focal points. Indeed, in my own mind, at least, the list is at present rather open ended, and unstructured. As I have noted, there has been little systematic discussion of the focal points, and I am not at all sure how best to organize the candidates or even what all the most plausible candidates are.

In contrast, the examination of foundational theories is at least more well trod terrain, and I have somewhat greater confidence here that I have a handle on the most important proposals that have been made. Unfortunately, I have not had the space to discuss or even to introduce all of them.

Even for those foundational theories that I have discussed, I have said little or nothing about the various rationales that might be offered on their behalf. Of course, as I have noted, one kind of support that a foundational theory can have is the very fact that it succeeds in generating the various normative factors that we are independently inclined to accept. But foundational theories can be attractive in their own right, and it is always worth asking of a given approach why an approach of that *kind* should seem plausible.

As often as not, if we do start to explain why a given kind of foundational approach seems attractive, we will soon find ourselves appealing to alternative conceptions of the very nature of morality—its point, and its place in the world of persons, reasons, and things. In short, we will find ourselves appealing to alternative metaethical views. But this is just as it should be. For normative ethics is not, in the final analysis, independent of metaethics. Each leads inevitably to the other.