HENRY S. RICHARDSON   Specifying Norms
                        as a Way to Resolve
                        Concrete Ethical Problems

> We want to walk: so we need *friction*. Back to the
> rough ground!
>   Ludwig Wittgenstein

I. INTRODUCTORY

Starting from an initial set of ethical norms, how can we resolve concrete
ethical problems? We may try to *apply* the norms to the case, and if they
conflict we may attempt to *balance* them intuitively. The aim of this pa-
per is to show that a third, more effective alternative is to *specify* the
norms. The problems that I am concerned with are of a sort typically
ranged under the rubric of "applied ethics"—that is, relatively particular
questions about what should be done in individual cases (e.g., Baby
Doe's) or concretely described types of cases (e.g., Baby Doe cases). Al-
though some who have worked most fruitfully on concrete ethical prob-
lems have begun to chafe at the "application" label,[1] no alternative met-
aphor has taken hold.[2] I aim to develop an alternative metaphor—that of

  1. See, e.g., Albert Jonsen and Stephen Toulmin, *The Abuse of Casuistry* (Berkeley: Uni-
versity of California Press, 1988); Michael D. Bayles, "Moral Theory and Application," *So-
cial Theory and Practice* 10 (1984): 97–120; and Tom L. Beauchamp, "On Eliminating the
Distinction Between Applied Ethics and Ethical Theory," *Monist* 67 (1984): 515–31.
  2. In lieu of "applied ethics," some now employ "practical ethics." The latter, however,
does not suggest any particular methodological understanding, except insofar as it implies
a contrast, which I would reject, with "theoretical ethics." I would instead follow Aristotle
in thinking of ethics as a whole as a practical subject. Within that subject, concrete prob-
lems about what should be done are but a subclass of the important questions—as is em-
phasized by, for example, Edmund Pincoffs, "Quandary Ethics," *Mind* 80 (1971): 552–71.

specification—into a model that will promote a better methodological understanding than the currently prevailing models, those invoking only the operations of application and intuitive balancing. My purpose in working out this model is to reform the way many of those working on these issues understand what they are doing and to articulate more explicitly what has already been done by others.[3] While much of the best work that has been done undoubtedly already accords with the model of specification to some degree, making this alternative approach explicit should bolster these efforts. By stemming the skeptical doubts fostered by the failure of pure application while avoiding the excuses for inarticulateness that an excessive reliance on intuition affords, the model of specification should encourage a more fruitful approach to the difficult problems of concrete ethics.

By a "model of how to resolve concrete ethical problems" I mean a schema of what it would be to bring norms to bear on a case so as to indicate clearly what ought to be done. The deductive application of rules to cases and the intuitive weighing of considerations are the two cognitive operations usually thought central to this task. I seek to add specification as a third, even more important operation. A *pure* model would see only one of these operations as involved in settling concrete cases. Most models actually defended are, in fact, to some degree hybrids, giving some role to application and some to intuitive balancing. I will criticize these hybrids, arguing that instead of a hybrid view built upon these two commonly understood operations, what is needed is a model built around specification: a true third way, rather than just a mixture. To be sure, once the operation of specification has been adequately understood, it may then be admitted that it should be supplemented by application and balancing in a more complex hybrid model. For convenience, I will refer to this eventual hybrid centered on the idea of specification as "the model of specification." The main task of this paper, however, is simply to make out the nature and the promise of specification as a third way.

In concentrating on the form of reasoning involved in bringing norms to bear on cases, I largely set aside the heuristics of ethical discussion and deliberation. It is of course true that discerning morally relevant features of a case, marshaling facts, analyzing arguments and ideas, citing

3. My aim is not, for instance, to capture the a priori structure of concrete ethical reasoning or the logic of "practical inference."

examples, and comparing the case at hand with a range of related cases ("the case method") are generally indispensable means to productive discussion and deliberation about concrete ethical issues. Each of the models I will consider recognizes this.[4]

Let me begin by describing the pure operations from which the common hybrids are built, namely, application and intuitive balancing. I will hijack the term *application* to name a quite pure deductivist approach to ethical problems. It is appropriate to start by considering a deductivist model, for the problems with which we are concerned all pertain to individual cases—sometimes as sui generis but more likely as occurring regularly under some description. Either way, the crucial question is how ethical norms reach down to individual cases. The answer given by the pure model of application, as I am defining it, is: by deductive inference that subsumes the case under a rule.[5]

This said, it becomes obvious that a natural train of thought leads many to shift from the pure model of application to a hybrid including intuitive balancing, as follows. When working with a strictly individual case, the pure model of application must work along the lines of a "Peripatetic syllogism" of the following form:

(i)   For all actions x, if Ax then x is (is not) permitted (obligatory).
(ii) Aa.

Therefore,

(iii) Action a is (is not) permitted (obligatory).

In these formulas, A stands for a description of an action that might be complex, and might include information about the circumstances or occasion of the action, the agent, or the persons affected, in addition to the type of action to be performed or avoided. For example:

(1) It is wrong for lawyers not to pursue their clients' interests by all means that are lawful.
(2) In this case of defending an accused rapist, it would lawfully promote the client's interest to cross-examine the victim about her

---

4. Thus I agree with many of the positive suggestions about ethical heuristics put forward by Jonsen and Toulmin, Bayles, and Beauchamp.

5. This subsumption, of course, need not be seen as a mechanical process, but instead may (as Kant emphasized) call upon the faculty of judgment: see Onora O'Neill, "The Power of Example," *Philosophy* 61 (1986): 5–29.

sex life in such a way as to make sexist jurors think that she con-
sented.

Therefore,

(3) It would be wrong not to cross-examine the victim in this way.

Yet in addition to the norm stated in (1) there is another that is relevant
and is no less true, that may even be subscribed to by the lawyer in ques-
tion, and that leads to an opposite conclusion:

(4) It is wrong to defame someone's character by knowingly distorting
their public reputation.
(5) To cross-examine the victim about her sex life in such a way as to
make sexist jurors think that she consented would be to defame
her character by knowingly distorting her public reputation.

Therefore,

(6) It would be wrong to cross-examine the victim in this way.

Here we have at least an apparent moral dilemma, potentially generated
even out of the moral convictions of one central actor and certainly out
of principles common in our culture (with its adversarial legal system),
that is not easily brushed aside. Even apart from whether this dilemma
amounts to an intolerable logical contradiction, it frustrates the purpose
of a model for resolving concrete ethical problems, for the very least we
can say is that (3) and (6) give conflicting advice.[6] In the face of such a
quandary, one natural response is to try to "balance" the ethical consid-
erations represented by (1) and (4)—no longer seen as peremptory or
absolute—by assessing how much weight they are to be assigned in
these circumstances. In this way, the pervasiveness of ethical conflicts—
together with their centrality to our notion of an ethical "problem"—
leads many from the pure model of application to a hybrid model includ-
ing balancing. The question for balancers, in turn, will be how their
weightings are to be explained or justified. To anticipate my arguments
in a later section, we will see that to the extent that the balancing is

---

6. Logical contradiction might be avoided by, for instance, severing the connection be-
tween "wrong" and that moral sense of "ought" that invokes the principle that " 'ought'
implies 'can'." On these issues see the useful collection *Moral Dilemmas*, ed. Christopher
Gowans (New York and Oxford: Oxford University Press, 1987).

genuinely distinct from application it affords no claim to rationality, for to that extent its weightings are purely intuitive, and therefore lack discursively expressible justification.

The third model I shall develop and defend here, that of specification, aims to be more flexible, realistic, fruitful, and attainable than that of application without sacrificing its claim to discursive rationality. While the model of specification will be more precisely defined in Section III, its leading ideas are simple enough and may be briefly described at once. The model of specification concurs with the balancing approaches in seeing a need to qualify our commitments, but insists that this be done not by a quantitative weighting or discounting but instead by qualitatively tailoring our norms to cases. Thus, one is urged not merely to reflect and change one's mind in a way that resolves a conflict in an acceptable way, but to revise one's normative commitments so as to make at least one of them more specific. For instance, in our example of cross-examining the rape victim, one might get a narrowed version of (1):

(1′) It is wrong for lawyers not to pursue their clients' interests by all means that are *both* lawful *and ethical.*

By itself, of course, this amendment does not settle the issue; but it does focus attention on the strength of the role-excuse provided by the lawyer's place in the adversary system. In particular, it forces one to reconsider the breadth of (4). On the one hand, publicly undercutting a witness's character is a stock in trade of cross-examiners, and to that extent may appear to be an ethical tactic; but on the other, there are reasons to think that a rape victim requires special social protection for her allegations, just as the accused is due the adversarial protections of the presumption of innocence.[7] Accordingly, we might reaffirm a narrowed version of (4):

(4′) It is always wrong to defame *a rape victim's* character by knowingly distorting her public reputation.

And it looks as if (1′) and (4′) settle the conflict in favor of restraint.

As I shall show, there are at least four reasons to build an alternative model around the idea of making a norm more specific. The first, and

7. I am here abbreviating the sophisticated casuistical discussion of this case in David Luban, *Lawyers and Justice: An Ethical Study* (Princeton: Princeton University Press, 1988), pp. 150–52.

most obvious, is that it helps us imagine how our ethical precepts, many of which are very general and abstract, can reach concrete cases without generating unacceptable implications. A second important reason for this requirement of greater specificity is that it helps ensure (without merely stipulating) that the reasonable motivation behind the initial, unqualified norm is still captured by what one ends up with. As we shall see, a third is that the notion of specification, conceived thus as a relation between two norms, allows us to understand in a precise way how a "mid-level" norm can serve as a bridge between a general precept and a concrete case.[8] Fourth, this same notion can also explain how a moral theory can remain the subject of a more or less stable attachment despite the sort of revision that moral conflicts engender. It is only concurrence on such a relatively stable moral theory that can assure that different people will reach the same conclusions in employing the model of specification; the same, of course, is true of the model of application.

It will require some work to show that the model of specification is truly distinct from the other two, for there is always a strong tendency to redescribe the unfamiliar (specification) in terms of the familiar (application and balancing). Accordingly, in the next section I set out the models of application and balancing and explain their principal defects. Specification is then defined, in Section III, in a way that makes clear how it is a genuine alternative to the model of application. The fourth section undertakes to show that there are rational constraints on specification that preserve its claim to discursive justification and hence distinguish it from an intuitive balancing model; and Section V gives some examples that exhibit this possibility concretely. The final section of the paper will compare the three models, arguing in general terms for the superiority of the model of specification.

II. Deficiencies of the Traditional Models

The two dominant pictures of how to bring ethical norms to bear on concrete questions of practice—application and intuitive balancing—can seem to exhaust the possibilities.[9] One abstract explanation for this false appearance is the following. Recall that in characterizing what the dif-

---

8. The notion of "mid-level bridging principles" is given prominence in Bayles, "Moral Theory and Application," but the bridging relation receives little analysis in his treatment.

9. See, for example, F. H. Bradley, "My Station and Its Duties," in his *Ethical Studies*, 2d ed. (Oxford: Clarendon Press, 1927), pp. 193–99.

ferent models are models *of*, I said that we may suppose that the delib-
erators or discussants start with a set of ethical norms to which they are
in some important way initially committed. For our purposes, it does not
matter where these norms come from, whether they change over time,
or how they are grounded, if at all. Suppose that someone tries to apply
them to concrete issues using deductive means to produce answers or
decisions—not expecting a complete decision procedure, but hoping to
avoid inconsistency. There seem to be two possibilities: either the norms
are orderly enough to allow deductive application to avoid dilemmas or
they are not. If, say, the norms can be lexically ordered according to
higher-order priority rules, then the model of application may succeed.
If, by contrast, there is no way to order the norms so as to preclude ap-
parent dilemmas, then one seems forced to fall back on an intuitive bal-
ancing of the clashing norms—for by hypothesis there are no higher-
order rules that could settle the conflicts in a nonintuitive way.

  The idea of balancing, *sans phrase*, could arise under each of these
possibilities, depending upon whether or not the "weights" are taken to
reflect criteria that exist antecedently to and independently of the act of
weighing. Antecedent weighting principles might exist, for example, if
the "weight" of a norm were a homogeneous and objectively measurable
feature of it, or else if the weights derived from a set of implicit priority
rules that ranked the norms in question. Under either of these supposi-
tions, balancing could be assimilated to the pure model of application. In
the former case, the single highest-order principle would be to choose
the option with the greatest net weight of reasons in its favor.[10] In the
latter case, the implicit priority rules could be made explicit, if not in the
context of deliberation, then at least for the purposes of justifying some
choice after the fact. For example, one might grade students' philosophy
papers rather intuitively or impressionistically, yet be able to reconstruct
one's implicit reasons afterward by noting that one ranks making argu-
ments above expounding texts. Or again, one might justify breaking a
promise to a friend in order to save a drowning person on the grounds
that saving a life is more important than keeping a promised social en-
gagement.[11] In each case, one has invoked a priority rule, however
vague. This then returns one to the question of the orderliness of these

  10. For grounds for doubting that the use of comparison cases is a good way of bringing
out antecedent, implicit weights, see Shelley Kagan, "The Additive Fallacy," *Ethics* 90
(1988): 5–31.
  11. I am here indebted to the comments of an anonymous reviewer for this journal.

priority rules: if the implicit rules are consistent and sufficiently com-
plete, then they could be deductively applied. If not, the need for a more
intuitive sort of balancing may arise.

If the balancing of norms is thought not to be getting at antecedent
weighting principles, but to be creating or supplying them, or merely to
be providing particular judgments that might inductively support some
weighting principle, then it is what I am calling "intuitive balancing." In
intuitive balancing, any weighting principle—and therefore the notion of
weight as a homogeneous property, as discussed in the last paragraph—
is consequent upon the weighing.[12] An analogy would be the case of Von
Neumann–Morgenstern utility functions, which are consequent upon
(induced by) a consistent set of preferences. In teleological versions of
intuitive balancing, values are balanced.[13] In nonconsequentialist ver-
sions—I will get to Ross in a moment—principles are balanced.[14] There
are various labels for the faculty exercised in intuitive balancing: intui-
tion, judgment, perception. The balancing might occur within a theory
that had reduced all moral considerations to two heads, assessing only,
for instance, the relative weight of autonomy and beneficence in a given
case; or the balancing might be more open-ended, declaring the "bal-
ance of reasons" tentatively to favor one option over another.

It is widely if not universally recognized that the two pure models in-
volving only application or balancing are each seriously deficient. The
pure model of application depends upon an ideal achievement of moral
theory that is beyond our grasp. The claims of certain versions of natural
law theory notwithstanding, we are unable to systematize our common-
sense norms sufficiently to preclude widespread and serious conflicts

12. One can therefore understand Alan Donagan's complaint in *The Theory of Morality*
(Chicago: University of Chicago Press, 1979), p. 23, that to call this intuitive assessment
"weighing" or "balancing" is "fraudulent"; but this metaphorical use of these terms is well
entrenched in Western culture—embodied, as it is, in the figure of blind Justice—and is a
convenient one for the comparative intuitive assessment of competing considerations.

13. Although there are elements of teleological balancing in G. E. Moore and Hastings
Rashdall, the implications of this for concrete ethics are obscured by their overall utilitar-
ianism. A less ambiguous statement of this sort of view is found in Thomas Nagel, "The
Fragmentation of Value," in his *Mortal Questions* (Cambridge: Cambridge University
Press, 1979).

14. One nonconsequentialist balancing view that is purer than Ross's is that of Bradley
in "My Station and Its Duties." See also Robert M. Veatch, *A Theory of Medical Ethics*
(New York: Basic Books, 1981), pp. 303–4, and Baruch Brody, *Life and Death Decision
Making* (New York and Oxford: Oxford University Press, 1988), esp. pp. 77–79.

among them. Utilitarianism would introduce system, but at too high a
price, requiring too great a departure from norms to which we are firmly
committed. This is no contingent limitation of our current ethical knowl-
edge or sophistication. The inherent particularity and variability of the
subjects with which ethics has to deal will prevent us from ever being
able to present a system of rules that takes account of all the needed
qualifications and distinctions.[15] Since it is inherent to the nature of the
subject, this limitation is not merely contingent upon the present stage
of progress in ethical theory. Consider the parallel issue in the law. All
operating legal systems have accepted severe limitations on treating law
as a deductive system, and have instead developed case-oriented and
precedent-bound approaches that make room for equity, as described by
Aristotle (in the *Nicomachean Ethics*, bk. V) and as familiar in English
common law, namely, scope for the judge to modify the rules to fit the
case at hand.[16] Given the role of law in the public legitimation of the state
and in grounding stable expectations for commerce and society, there is
every reason to strive for a rule-bound, deductive approach to adjudica-
tion. Since even the law is forced to give up on the pure deductive ideal,
it is hardly likely that ethics, where the motivations for deductive trans-
parency are much weaker, could succeed in living up to it. The pure
model of application is doomed to be hamstrung by dilemmas and
strained by the effort to accommodate the diversity of moral experience.

Intuitive balancing, by contrast, is all too easy. As preference utilitari-
ans are fond of reminding us, we do it all the time. The problem with
intuitive balancing is not its unattainability but its arbitrariness and lack
of rational grounding. By whatever name the balancing faculty is graced,
its operations seem intrinsically beyond the pale of justification. There is
no faculty of ethical intuition (or perception or judgment) whose discur-
sively unjustifiable deliverances carry their warrant with them. Since in-
tuitive balancing does not proceed by measuring any objective feature of
the world or of our system of reasons, what it does do remains mysterious

15. This Aristotelian theme is nicely developed by a series of writings by Martha Nuss-
baum, including *Aristotle's De Motu Animalium* (Princeton: Princeton University Press,
1978), essay 4; *The Fragility of Goodness: Luck and Ethics in Greek Tragedy and Philos-
ophy* (Cambridge: Cambridge University Press, 1986), esp. chap. 10; and "The Discern-
ment of Perception," in *Proceedings of the Boston Area Colloquium in Ancient Philosophy*,
vol. 1, ed. J. Cleary (Lanham, Md.: University Press of America, 1986).

16. If this is right, then it betrays a misconception of the law to speak of the model of
application as "ethical legalism."

or, at best, subjective. The balance we strike can be taken as a primitive, as a *datum*; but that is again just to admit that we cannot give it a rational grounding.[17]

As I have noted, most influential models of concrete ethics are hybrids that combine elements of application with elements of balancing. What these hybrid approaches have in common is the thought that although moral conflicts frustrate a purely deductive model, there is nonetheless considerable room for deductive application. A common metaphor for a hybrid approach is that of core and penumbra. Ethical principles are taken as having a core zone in which deductive application may proceed confidently, and a penumbral zone in which there can be no secure application—whether or not there are conflicts with other norms. In this picture, intuitive judgment is called in more to settle the uncertain meaning of a single norm than to adjudicate a conflict between norms.[18] But since norms will conflict, we should consider hybrid models that account for this possibility. The most influential and interesting of these is W. D. Ross's account of prima facie duties. A prima facie duty is a norm that does serve as an adequate basis for deducing actual duty in a given situation so long as it does not come into conflict with any other norm.[19] When prima facie duties do conflict, however, none of them is any longer allowed to serve as a basis for deducing actual duty.[20] In cases of conflict,

17. For a frank casting of preference formation as primitive in this way, see James Griffin, *Well Being: Its Meaning, Measurement and Moral Importance* (Oxford: Clarendon Press, 1986), pp. 36, 103.

18. The distinction between core and penumbra is recently familiar from H.L.A. Hart's jurisprudence: see "Positivism and the Separation of Law and Morals," in Hart's *Essays in Jurisprudence and Philosophy* (Oxford: Clarendon Press, 1982), pp. 63–64. A development of the metaphor of core and penumbra for moral purposes is found in classical Catholic casuistry, as expounded by Jonsen and Toulmin in *The Abuse of Casuistry*, chap. 16. Central to this conception of casuistry is the idea of a paradigm case and progressive departures from it until things become muddy, when "discernment" is needed.

19. W. D. Ross, *The Right and the Good* (Oxford: Clarendon Press, 1930; repr. Indianapolis: Hackett, 1988), p. 19. On pp. 30–34 Ross develops the claim that every possible act is the subject of conflicting moral considerations. In describing Ross's view as a hybrid model in the text, I have been supposing that this is not necessarily the case. If it is, then Ross's view obviously collapses back into the pure model of balancing.

20. My reading of Ross's notion of "prima facie duty" focuses on the contrast with "actual duty," thus coming close to John Searle's helpfully critical interpretation, in "*Prima Facie* Obligation," in *Practical Reasoning*, ed. Joseph Raz (Oxford: Oxford University Press, 1978), pp. 81–90, of what he calls Ross's "official view." Because he focuses solely on conflict situations, Searle gives insufficient weight to Ross's insistence that a prima facie duty is one that will be an actual duty so long as no other duties conflict with it.

Ross says, quoting Aristotle, "the decision rests with perception."[21] Accordingly, on Ross's view, the model of application prevails so long as there are no conflicts; and where there are, we turn to intuitive balancing.[22]

These hybrid views fare no better than their purer counterparts. In conceding that the ideal of pure application is beyond our grasp, and therefore relying considerably on intuitive balancing, they undercut the former model without explicating the latter. Perception's decision of conflicts between prima facie duties, like judgment within the penumbra, remains mysterious. It undercuts the idea of application by the admission that it is not the only model of practical reasoning and by declaring certain norms not to "apply" in the penumbral or conflicted cases. The universal norms appealed to by a pure model of application get much of their support from their claim to introduce a deductively consistent and complete hierarchical order into our moral thinking. Since the hybrid models largely give up on this ambition, they tend to weaken the rationale of the norms with which they start without providing an alternative picture of systematization to replace that of deductive hierarchy. In this respect, these hybrid models appear as application manqué, as deductivism thrown back on intuitionist resources by the harsh realities of moral conflicts.

Given these difficulties both with the pure models of application and balancing and with hybrids that mix them, it is not surprising that philosophers have been casting about for alternative ways of relating ethical norms to concrete issues of practice. A common tack is to accept Ross's basic framework while looking for a supplement that would provide for a more discursive resolution of conflicts among the prima facie duties.[23] A

21. Ross, *The Right and the Good*, p. 42, quoting Aristotle, *Nicomachean Ethics* 1109b23; cf. 1126b5.

22. Instead of using the Peripatetic syllogism to represent the deductive aspect of a Rossian hybrid view, it might be more accurate to follow G. H. von Wright's suggestion of thinking of obligatoriness, permittedness, and so on *both* as predicates of individual actions *and* as operators on propositions applied to types of action ("On the Logic of Norms and Actions" in *New Studies in Deontic Logic: Norms, Actions, and the Foundations of Ethics*, ed. Risto Hilpinen [Dordrecht: Reidel, 1981], pp. 3–35). It would still be the case, however, that *every* act that fell under a forbidden (or permitted) category of action and no other morally relevant category would be forbidden (or permitted).

23. The view that Ross's position is all right as far as it goes but needs a significant supplement has been embraced by moral philosophers of as diverse positions as Albert Jonsen, one of the coauthors of *The Abuse of Casuistry*, for whom the supplement is an

surprising number of these attempts follow Ross in recalling Aristotle's protean account of practical wisdom (*phronêsis*) without doing the philosophical work necessary to explain how discursive reasoning figures in this account. One thoughtful critic who has emphasized how the Aristotelian ideal of practical wisdom can both help us understand how we may rationally cope with moral conflicts and serve as a prop to obscurity finishes by saying that although judgment is essential, no theory of practical judgment is possible.[24] Despite this warning, I will use the notion of specification to begin to develop one that replaces Ross's framework rather than merely supplementing it.

III. Specification as an Alternative to the Dominant Models

The reason that the models of application and balancing do not exhaust the field is that they each suppose that the set of norms invoked in ethical discussion and deliberation is held fixed. As has been emphasized especially strongly by Deweyan ethical pragmatism, however, and as is also implied by the Rawlsian idea of wide reflective equilibrium, our norms are subject to revision. The model of specification starts from this recognition of revisability, but reestablishes a kind of constancy or stability not implied either by the general pragmatist approach or by the idea of reflective equilibrium. This stability is essential to the claim that the initial norms are in some way *brought to bear* on concrete cases by means of more specific norms.

It will be useful to see how the model of specification both incorporates aspects of the pragmatist approach—which is currently enjoying something of a revival[25]—and adds to it a crucial missing element. In the face

---

ordered appeal to paradigm cases (personal communication); and R. M. Hare, for whom the supplement is "critical level" thinking, which is allegedly utilitarian. See Hare's "Comments" in *Hare and Critics: Levels of Moral Thinking*, ed. Douglas Seanor and N. Fotion (Oxford: Clarendon Press, 1988), p. 223.

24. Charles E. Larmore, *Patterns of Moral Complexity* (Cambridge: Cambridge University Press, 1987), pp. 18–20.

25. See, e.g., Andrew Altman, "Pragmatism and Applied Ethics," *American Philosophical Quarterly* 20 (1983): 227–35; and James D. Wallace, *Moral Relevance and Moral Conflict* (Ithaca: Cornell University Press, 1988). An uncritical and antitheoretical version of the pragmatist approach—one hardly hospitable to the ideal of wide reflective equilibrium—is criticized in Bayles, "Moral Theory and Application." In the present article, I abstract from the consequentialism that Altman takes to be central to the pragmatist account of justification. My approach is particularly close to Wallace's, and I would endorse many of the

of conflicts, Dewey wrote, deliberation will be reasonable if it can devise a "*way* to act," or light upon a conception of the object of action, "in which all [competing tendencies] are fulfilled, not indeed in their original form, but in a 'sublimated fashion,' that is, in a way which modifies the original direction of each by reducing it to a component along with others in an action of transformed quality."[26] Since an explanation of the way in which the action reflects the "original" norms would be facilitated if the principle of the action were explicitly stated rather than simply left to the observer to infer from the action's features, we might think the following reformulation of Dewey's idea by James Wallace to be an advance: "[When moral considerations conflict,] the aim [of deliberation] must be to modify one or more considerations so that it applies, so that its original point is to some degree preserved, and so that one can live with the way [of proceeding] so modified."[27] In fact, without necessarily buying into broader pragmatist assumptions about the nature of truth, the model of specification will take this general description of practical reasoning as its point of departure. The key unanswered question, however, for both Dewey's and Wallace's versions of this pragmatist account is, What licenses us to call a modification or sublimation of an original norm still in some significant sense the *same* norm that we started out with? Why is it not a self-contradiction to speak of modifying a consideration so that *it* applies? Is not what "applies," in the end, just a different norm or consideration?[28] If so, then so far as *concrete* ethics is concerned, it would look as if the pragmatic approach posed no alternative to the hybrids we have already discussed: changes in norms are arrived at intuitively, and once the change has occurred, the new norms can be deductively applied.

The idea of specification aims to complement this general pragmatic approach by laying down conditions on the relation between the initial norm or norms and their modifications that explain how the original norms *are* being respected (in a "sublimated fashion"). Answering this question of stability serves two purposes. First, in this way, the idea of

points made in his stimulating book. It is the idea of specification that is lacking from his account, and that is needed to fill out the general pragmatist approach that he sets out.

26. John Dewey, *Human Nature and Conduct* (New York: Henry Holt, 1922), p. 194.

27. Wallace, *Moral Relevance and Moral Conflict*, p. 86.

28. I am indebted to Nathan Salmon for having pressed me for an answer to this question.

specification converts the general pragmatist approach into a truly distinct model for concrete ethics. Second, we want some such notion of constrained change both to allow the development of a stable moral theory and to give us some assurance that the commitment that underlay the initial norm is being appropriately honored. Yet there are three reasons not to simply stipulate that an acceptable specification is one that captures this initial commitment and leave it at that: (1) We may not be aware just what the contours of this initial commitment are. The model of specification lets these emerge on reflection. (2) If the underlying rationale of a norm being specified can be laid out, then this just means that we have looked through it to some other norms, which are then what should be in play. (3) To honor the initial "point" or source of commitment is easier said than done. For all these reasons, in defining the notion of specification I will seek to delineate a notion that will not merely posit stability but will explain it.

For this sort of stability in the course of revision to be possible, it must be the case that the norms being specified are not "absolute" in logical form the way (i) is, that is, are not strictly universal with respect to the domain of possible acts. Instead of being, in this sense, prefaced by an "always," they must be seen as implicitly beginning with a "generally speaking." To see why the latter form is necessary to stability, note that there are two possibilities: either the more specific norm that results from deliberation replaces the initial norm it specifies or else it stands alongside it. If it stands alongside, it would mark an expansion of the set of norms. If it replaces, that would be a true revision of the set of norms.[29] If the initial norm were strictly universal, then a specification that stood alongside it would be otiose, since it would already be implied in the initial norm, and could be omitted as an unnecessary step in a deductive argument to a practical conclusion. For example, if we began with "it is always wrong to lie," then "it is wrong to lie to someone who has a right to the truth" would be redundant. As Kant held, the right to the truth would be irrelevant to the permissibility of lying. If the more specific norm replaced the one it specifies, however, the result would be an implied exception that would be logically incompatible with the initial norm's universal command, making it difficult to see any stability. Accordingly, to conceive of a kind of stability over the course of a path of

---

29. For the terminology of revision and expansion, see Peter Gärdenfors, *Knowledge in Flux: Modeling the Dynamics of Epistemic States* (Cambridge, Mass.: MIT Press, 1988), sec. 3.1.

specification that does useful work, one must suppose that the norms
being specified are not absolute in the fashion of (i).

   This supposition poses no difficulty, for the norms to which we are
commonly committed are not plausibly viewed as formally absolute in
this way. Rather, they are typically qualified, at least implicitly, by var-
iants of "generally" or "for the most part." This sort of looseness is a com-
mon feature of our norms as we find them, whether they be prohibitions,
positive duties, or ends. As T. M. Scanlon has written, "common-sense
moral principles . . . are not simple self-contained rules. . . . Qualifica-
tions having to do with intent, justifications, excuses, and so on, while
not always explicit in any formulation of the principle, are part of the idea
referred to (though their exact boundaries are never clear)."[30] That blan-
ket prohibitions on lying and killing are not universally applicable is
widely recognized. The kind of looseness our norms allow is often
thought of as making implicit room for exceptions. Another way of re-
garding it is to use Kant's notion of "latitude." This idea is different in-
sofar as it suggests that the exact extent and nature of the duty may
require further specification. Thus, in presenting the positive duty of be-
neficence, Kant presents it as not specifying when, exactly, one must
help others, what one must do to help them, or to what degree one must
sacrifice one's own welfare in doing so. Nonetheless, there is an "imper-
fect" duty to do *something* for the sake of helping others.[31] I believe that
our common beliefs about beneficence are similarly latitudinarian; and
our ends and values are even more obviously qualified in this way. To be
sure, many philosophical debates begin by supposing they are not: we
may think, for instance, that a political regime must do everything it can
to promote liberty *and* everything it can to promote equality, and end up
having to cope with the resulting conflict.[32] Yet most ends, including
these political ideals, fall short of requiring that we do everything possi-

   30. T. M. Scanlon, "Levels of Moral Thinking," in *Hare and Critics: Levels of Moral
Thinking*, ed. Seanor and Fotion, p. 134. Despite this talk of reference to an idea, Scanlon
states clearly that he does "not mean to suggest that there is a fixed set of moral ideas" to
be simply discovered (p. 137). Hare, responding to Scanlon in ibid., p. 263, poses the chal-
lenge that I am now addressing of how a norm can be seen as "the same" before and after
revision.
   31. Immanuel Kant, *The Metaphysics of Morals*, pt. 2 (*Tugendlehre*), Introduction, secs.
7–8.
   32. For a sketch of the importance of the idea of specification for liberal political theory,
see sec. 8 of my essay "The Problem of Liberalism and the Good," in *Liberalism and the
Good*, ed. Gerald M. Mara, R. Bruce Douglass, and Henry S. Richardson (New York: Rout-
ledge, 1990).

ble that can be done in their service. For instance, a serious commitment
to the end of liberty does not imply that one must shield pornographers
or reduce government regulation to an absolute minimum.

While the model of specification begins with norms that allow for lati-
tude, it also ends with them; and this feature is crucial to its status as a
distinct model. It is not realistic to view the goal of concrete ethical dis-
cussion or deliberation as coming up with adequately limited ("mid-
level") though still strictly universal norms. To be sure, if we never arrive
at norms of this form, then since deductive application along the lines of
the Peripatetic syllogism must begin from a norm of the form of (i), it
will prove impossible. That it is impossible was argued in the previous
section, precisely on the grounds that the complexity of the moral phe-
nomena always outruns our ability to capture them in general norms.
Fortunately, there is no need to settle individual cases deductively in or-
der to settle them on rationally defensible grounds. The central assertion
of the model of specification is that specifying our norms is the most
important aspect of resolving concrete ethical problems, so that once our
norms are adequately specified for a given context, it will be sufficiently
obvious what ought to be done. That is, without further deliberative
work, simple inspection of the specified norms will often indicate which
option should be chosen.[33] A conclusion supported by considerations that
hold only "for the most part" is, of course, rebuttable by further deliber-
ation; but that is the way our moral reasoning goes. The ability to bring
norms to bear on cases even while leaving them nonabsolute is a distinc-
tive feature of the model of specification.

Having thus sketched the model of specification, I am now in a posi-
tion to define the specification relation more precisely. Doing so will re-
quire some preliminary definitions of component notions. Although the
model of specification supposes that the norms to which we are actually
committed are typically not absolute, it will be convenient to define the
relation of specification by reference to artificially tightened versions of
these norms that, unlike a norm that begins with "generally" or "for most
actions," can have well-defined and presumably bivalent conditions of

---

33. As Aristotle would put it, once the norm (end) is sufficiently specific, there is no need
for further deliberation, and it is "perception" that must supply the "premise" that a cur-
rently possible action satisfies the norm. See, e.g., *De Motu Animalium*, chap. 7, and the
discussion thereof by David Charles, *Aristotle's Philosophy of Action* (Ithaca: Cornell Uni-
versity Press, 1984), p. 96.

satisfaction. Without undue violence, then, the typical looseness of ordinary norms could be represented as taking the following form:

(i*) For *most* actions x, if Ax then x is (is not) permitted (obligatory).

The *absolute counterpart* of a norm of the form of (i*) is one that restores it to the form of (i) by replacing the hedging "for *most* actions" with the absolute "for *all* actions."[34] The detailed way in which the satisfaction of a norm is to be understood will presumably depend upon the *type* of norm involved, that is, upon whether it is an end, permission, requirement, or prohibition. In general, we can say that an *instance* of an absolute norm is an alternative that satisfies it. An instance of an end is an example of what amounts to its achievement or actualization; an instance of a permission is an example of what is permitted; and an instance of a prohibition is an example of avoiding what is prohibited. Given the generality of my account, it would be wrong to try for too much precision in the notion of satisfaction that underlies this extended usage of the term *instance*.[35]

On the basis of these other definitions, I now define *specification*, considered as a relation between two norms, as follows:

Norm p is a *specification* of norm q (or: p specifies q) if and only if

(a) norms p and q are of the same normative type;[36]
(b) every possible instance of the absolute counterpart of p would count as an instance of the absolute counterpart of q (in other words, any act that satisfies p's absolute counterpart also satisfies q's absolute counterpart);
(c) p qualifies q by substantive means (and not just by converting universal quantifiers to existential ones) by adding clauses indicating what, where, when, why, how, by what means, by whom, or to whom[37] the action is to be, is not to be, or may be done or

---

34. If a norm is already absolute, then its absolute counterpart is itself. I remind the reader that the absoluteness I have in mind is a matter of logical form, not epistemic basis.

35. See the instructive worrying of this notion of satisfaction in Ludwig Wittgenstein, *Philosophical Investigations*, 3d ed. (New York: Macmillan, 1958), pt. 1, secs. 437–39.

36. Ideally, this restriction on cross-type specification should be overcome; but I have thought it better to avoid the complications that would come with the attempt.

37. This clause uses the traditional list of "circumstances" developed by the casuists: see Jonsen and Toulmin, *The Abuse of Casuistry*, p. 253. The list goes back to Aristotle's theory of the "predicables" in the *Categories*.

the action is to be described, or the end is to be pursued or con-
ceived; and

(d) none of these added clauses in p is irrelevant to q.

Several comments are in order. First, by referring to the absolute coun-
terpart of the specified norm, clause (b) enables specification to be de-
fined in terms of the ordinary notion of containment without requiring
that specification proceed from absolute norms.[38] It can therefore rule
out making an exception by disjunction. That is, if the original norm says
"when you have received great benefits from someone that were not sim-
ply your due, you should generally express your gratitude to him or her,"
this cannot be specified to read "when you have received great benefits
from someone that were not simply your due, you should generally *either*
express your gratitude to him or her *or surreptitiously aid his or her
child.*"[39] An act that satisfied the second norm by aiding the child could
fail to satisfy the first norm's absolute counterpart. Second, clause (c)
implies that the sense in which a specification is more "specific" than
what it specifies goes beyond the "subset" requirement of clause (b), rul-
ing, for instance, that a move from "torture is always wrong" to "torture
is sometimes wrong" is not a specification. Specification proceeds by set-
ting out substantive qualifications that add information about the scope
of applicability of the norm or the nature of the act or end enjoined or
proscribed. Sometimes it will do both. For instance, "euthanasia is gen-
erally wrong" might be specified by "it is generally wrong to shut off the
respirator of a patient in a potentially reversible coma." Third, while
clause (b) rules out specification by disjunction, clause (d) rules out
some forms of specification by conjunction. For instance, it blocks taking
"to promote the health of my patients and to write a great opera" as a
specification of the end of promoting the health of one's patients. The

38. Specification might begin from an absolute norm—and for this reason some in-
stances of deductive application are also instances of superfluous specification—but it need
not.

39. This restriction on the logical means available in specification is rather narrowly
drawn. One might achieve much the same effect by putting the "exception" into the cir-
cumstances: "When you have received great benefits from someone that were not simply
your due, and when you do not have an opportunity to aid their child surreptitiously, you
should express your gratitude." This norm, however, is of narrower scope than that in
which the exception is by disjunction; in particular, the former does not say what one's
duty is in those circumstances in which one does have this opportunity, whereas the latter
explicitly leaves the agent two options in those cases.

element referring to opera is (presumably!) irrelevant to the patients'
health.[40]

Finally, note that in the definition of specification as a relation be-
tween two norms, there is no mention of any temporal or justificatory
priority between the two. To be sure, my thesis is that in the attempt to
resolve a concrete issue, the most important thing will be to make our
norms more specific; but it is also important that we sometimes *revise*
an end or principle, for what we consider to be good reasons, in a way
that will not count as a specification of it. Sometimes we move to a *less*
specific formulation of a norm, for instance. Even here, however, the
notion of specification can be useful, depending upon the nature of the
grounds for claiming that this change is a rational one. In many of these
cases, our appeal in changing a norm is to a deeper or more general norm
that underlay it, which we now claim to understand better, and for which
we provide a new specification. An example of this kind will be given in
Section V.

In light of my definition of specification, I can now elaborate on some
of the advantages of the model of specification briefly mentioned in the
first section. The first is that it allows for a strong reply to the objection
that a pragmatist view underestimates the seriousness of our commit-
ments to our initial norms. Although we recognize that they cannot be
regarded as absolute, if we suppose that they need to be and are specified
further we can nonetheless genuinely claim to be appealing to *them*.
This is because the definition of specification assures as nearly as is pos-
sible without stipulating it that the commitments expressed in the initial
norm would be honored in the satisfaction of the specified norm. While
the basic notion of extensional narrowing in clause (b) is important to
this result, clauses (c) and (d) reinforce this tendency. Although this will
not guarantee the preservation of commitment, it does rule out many of
the likely threats to it.[41]

Second, by giving us a way to articulate how a specific norm is mean-
ingfully related back to a more abstract one, the model of specification
helps secure a role for a stable ethical theory that seems as elusive on

40. There is no pretense, here, that any account of relevance can be innocent of substan-
tive ethical presuppositions.
41. My thinking on these issues, and on the model of specification in general, is much
influenced by David Wiggins, "Deliberation and Practical Reason," *Proceedings of the Ar-
istotelian Society* 76 (1975–76): 29–51; repr. in *Practical Reasoning*, ed. Raz.

the general conception of a change in view as it is unattainable conceived as a deductive hierarchy. System is more attainable by specification because the norms need not be taken as fixed or as formally absolute, and so will conflict less readily and adjust more easily than the sort of norms postulated by a pure model of application. Nonetheless, the connection back to an initial norm afforded by the notion of an instance of a general norm's absolute counterpart enables one to set out clearly what has remained the same in the course of specification. By making clear what remains constant despite modifications that are occurring, the model of specification allows one to distinguish the progressive refinement of a theory that remains the same in essentials from the mere shifting from one holistic equilibrium to another.

Third, and by the same token, the notion of specification provides a clear sense to the notion of a "mid-level bridging principle" which might otherwise be lacking. There is no trouble with "mid-level," understood loosely in terms of a rough sense of degrees of generality: the difficulty is in explaining the "bridging" relation. A mid-level norm that specifies a general one and thereby helps mediate the latter to a concrete case serves as a bridge in a quite definite sense—one across which, as I have just claimed, the discussant's or deliberator's commitment will likely travel. A series of progressively more specific norms would provide a bridge with multiple spans.

In addition, the model of specification can explain how it is that contingently conflicting norms "hang around" even after a conflict of norms has been resolved, sometimes giving rise to what Ross called "compunction" and to a duty to make reparations.[42] As John Searle has argued, Ross's own explanation of this phenomenon conflicts with his "official" view that in cases in which prima facie obligations conflict, they do not state actual duties.[43] On the model of specification, however, the general but nonabsolute norm, understood as stating an actual obligation, can stand alongside the specified version that averts the particular conflict at issue.[44] Take the example of a clash between a requirement to respect

42. Ross, *The Right and the Good*, p. 28. Regarding this paragraph, I have benefited from the friendly skepticism of Arthur Applbaum and Amy Gutmann, in addition to that of others named above.

43. Searle, "*Prima Facie* Obligation," p. 83. See also Bernard Williams, *Ethics and the Limits of Philosophy* (Cambridge, Mass.: Harvard University Press, 1985), pp. 176–77.

44. The sense in which nonuniversal norms can conflict will be explicitly explored in the next section.

299   *Specifying Norms*
as a Way to Resolve
Concrete Ethical Problems

one's promises and a prohibition on hurting others, in which you have promised to a person now dead to use their estate to build a research center on a certain plot of land you own, and only later find out that doing so would seriously disrupt the town water supply. Suppose that the initial conflict is averted by specifying the prohibition on breaking promises in a way limited in scope: "When the promisee is dead, respect for one's promises requires fulfilling the spirit of the promise but need not bind one to its precise terms." As the scope limitation allows, suppose further that the original norm has not been thoroughly rejected. Thus, if you subsequently find out that the town has switched to a new reservoir (and you have not yet spent the money), then there will be strong reason both to follow the letter of the promise and to add a further clause limiting the original qualification: ". . . need not bind one to its precise terms *when following them would be seriously disadvantageous to the public*." Here, the initial specification is rebutted (in practical effect) by a further specification that derives its force partly from the importance of the initial, general prohibition on promise-breaking, which has been retained.

    A similar argument could explain how the persistence of this initial prohibition—which is still seen as stating an actual duty—is relevant to accounting for the appropriateness of reparations. Thus, if no alternative water supply turns up, the appropriate secondary specification, instead of referring to public disadvantage, might require paying deference to the deceased's wishes in some more symbolic fashion. To be sure, when a more general norm is denied application to a particular case by a specification (via a scope restriction, for example), whether the specification replaces or stands alongside what it specifies, and in the latter case, whether it gives rise to a duty of reparations are questions that must be addressed on a case-by-case basis. As I shall argue in the next section, all such questions are to be answered in terms of the overall coherence and mutual support of the whole set of norms. If the initial general prohibition remains among them, then coherence grounds may support setting a threshold for its violation, requiring that the "least restrictive alternative" be chosen, or requiring reparations.

    So far, then, we have seen that the model of specification is genuinely distinct from the model of application, since a specification will in general not follow deductively from what it specifies. This alternative model escapes the apparently exhaustive choice between deductively applying absolute norms and qualifying norms solely by "weights" by denying that

the norms with which one starts are always, or even usually, absolute. Nonetheless, the notion of specification has enough structure to represent a significant improvement over the bare pragmatist idea of a change in view, and offers a number of other advantages besides. To show that the model of specification is also a genuine alternative to intuitive balancing, however, I must go further, and explain how specification can be rational—and in particular, how it can be discursively justified.

IV. Rational Specification

Are there rational constraints on specification? Unless there were, and unless the superiority of one specification over another could thereby be defended, it would be hard to see how specification could be anything but a special employment of intuition. I will propose a coherence standard for the rationality of specification. This standard in effect carries the Rawlsian idea of "wide reflective equilibrium" down to the level of concrete cases.[45] The power and interest of this coherence standard is underestimated if it is seen simply as requiring an absence of logical contradictions. In the last section, we saw how the relation of specification allowed for a stable development of theory despite revisions; now I suggest that developing a theory—even if only one quite limited in scope—is crucial to the rational defense of any specification. A (successful) theory importantly makes intelligible logical connections among the norms to which one is committed that do not merely demonstrate that they are logically compatible with each other, but also explain some of them in terms of others. It is this kind of argumentative support, and not mere lack of incoherence, that can justify.[46] Note that although there is no meaningful measure of the strength of argumentative support, it remains—unlike the grounds invoked by the intuitionist—fully subject to discursive statement and criticism.

There are two reasons why building explanatory or intelligible argumentative connection is vital to the rationality of the model of specifica-

45. The notion of wide reflective equilibrium was introduced, although without the label, in John Rawls, *A Theory of Justice* (Cambridge, Mass.: Harvard University Press, 1971), p. 49. The label is applied in Rawls's Presidential Address, "The Independence of Moral Theory," *Proceedings and Addresses of the American Philosophical Association* 47 (1974–75): 8.

46. On the importance of mutual support to justification, see Rawls, *A Theory of Justice*, pp. 21, 579.

tion. The first, as Gilbert Harman has emphasized, is just that it is important to the rationality of any "change in view."[47] The second is that developing a theory is an important means of avoiding contingent practical conflicts. These matter because the enterprise of deciding what ought to be done in a concrete situation (if not all of ethics) is a practical enterprise. This fact puts it in the same camp as intending or willing: if we believe it impossible—even contingently impossible—to do both x and y, then we have reason to revise the set of norms leading to the conclusion that we ought to do both. (This contrasts with wishing and desiring, for which jointly incompatible objects are perfectly normal and apparently acceptable.)

When a conflict of norms is the result of some contingent fact about the world, it is of course an important question whether we should not try to avert the conflict by changing the world rather than by changing our norms.[48] Often, there will be good reason to do both. For example, suppose that contingent limitations of technology now prevent us from being able to protect hospital populations from the spread of a potentially fatal disease except by severely limiting the work of surgeons discovered to be carriers of this disease. We cannot well serve both the epidemiological interests of the hospital population, the professional interests of the infected surgeons, and the interests of the patients potentially benefited by their surgical talents. In such a situation, we surely have reason *both* to seek improved technologies to block the spread of the disease *and* to attempt to resolve this contingent clash between the ends involved in a way that takes for granted the limitations of current technology. In doing so, it will be important to arrive at a finer-grained specification of the aims that guide us in this situation. For instance, the "interests of the surgeons" might be factored into financial security and the liberty to pursue their calling. The former interest might be satisfied by a special indemnity for those found to be carriers, while the latter might be further specified in a way that makes plain that it cannot be consistently achieved in a way that involves inflicting harm on patients, even unwittingly, by spreading disease. The fact that some such specification of the

47. Gilbert Harman, *Change in View: Principles of Reasoning* (Cambridge, Mass.: MIT Press, 1986), pp. 32–33 and chap. 7.
48. See Ruth Barcan Marcus, "Moral Dilemmas and Consistency," *Journal of Philosophy* 77 (1980): 121–36, repr. in *Moral Dilemmas*, ed. Gowans; and the criticism of Marcus in Alan Donagan, "Consistency in Rationalist Moral Systems," *Journal of Philosophy* 81 (1984): 291–309, repr. in *Moral Dilemmas*, ed. Gowans, p. 281.

ends involved could avoid the contingent conflict between the surgeons' interests (unspecified) and those of the patients provides one reason for specifying in this way.

Ross and F. H. Bradley were right, then, that such practical conflict among our norms is ubiquitous, but wrong in thinking that this fact left us with no alternative but intuitive and ad hoc balancing.[49] The owl of Minerva can do better than that. Because of the pervasiveness of potential conflicts, however, the only way to have any grounds to hope that a specification does not simply avert one local conflict at the cost of giving rise to worse ones elsewhere is to begin to develop a relatively stable moral theory.

These extensions of the idea of consistency of norms in the direction of practical or contingent consistency and mutual argumentative support explain the possibility of a conflict between norms that are not strictly universal. Even if such norms cannot be logically inconsistent with one another, they can be practically inconsistent insofar as they guide choice in opposite directions and disturb our attempts at theoretical systematization. Thus, consider the difficulties on both counts that have been posed by the clash between the views that it is generally wrong to lie and that it is important to prevent others from coming to harm.

A specification is rationally defensible, then, so long as it enhances the mutual support among the set of norms found acceptable on reflection. Typically, the removal of a conflict will enhance mutual support, but not always. In such matters, it is vain to strive after an ephemeral ideal of the total absence of conflict—especially since sometimes, on reflection, we see no acceptable way to rationalize away a given practical conflict. Still, there will be clear cases in which a shift in specification yields a more coherent overall view by acceptably removing a given conflict.

V. Examples of Specification

To reinforce my claim that specification can be discursively pursued and justified, I will present two schematic examples of it. One is set in a medical context in which the ruling model is that of a hybrid deontology

---

49. For Bradley's claim that conflicting considerations are pervasive, see "My Station and Its Duties," pp. 196–97 (footnote). For Ross's, see *The Right and the Good*, pp. 30–34. For Bradley's Hegelian view that philosophy cannot guide, but can only understand, see "My Station and Its Duties," p. 193.

along the lines of Ross, and the other is a case that would be treated as ethical only by certain value-maximizing or value-balancing ethical views.

First, then, let us consider a hypothetical and simplified course of deliberation about whether to withhold nutrition and hydration from a severely malformed newborn so as to let it die. I ask the reader to suspend any dissent from the particular norms mentioned, and to concentrate upon the relations among them. Our hypothetical agent will be imagined to hold certain views, not only about the case in question (which I will not describe exactly) but also about a range of related cases with respect to which—seeking rationality in specification—she aims to work out a consistent view. In the debate over the Baby Doe cases, it appears that at least three main principles are in contention, each of which our deliberator begins by accepting in nonabsolute form: (1) a prohibition on directly killing innocent persons (here, the newborn), (2) a duty to respect the reasonable choices of parents regarding their children (suppose that in this case the mother and father want to let their baby die), and (3) a duty to benefit the persons over whom one has responsibility (here, from the point of view of the medical personnel, the patients—i.e., the infant and the mother).[50] Assuming that the prohibition on directly killing innocent persons is not to be evaded, in the case in question, either on double-effect grounds (i.e., on the grounds that the killing is not "direct") or on act-omission grounds (i.e., on the grounds that withdrawing nutrition and hydration is merely "letting die" and not killing), there is a conflict between the first principle and the other two. An analysis of the concept of personhood will not help much here, for our deliberator quails at the implications *either* of classing these newborns as persons and sticking by the prohibition on killing *or* of classing them as nonpersons and thereby treating that prohibition as irrelevant. She finds the first option terrible from the point of view of benefit—cui bono? The second option she deems unacceptable partly on its own account and partly because of its implications by analogy for the treatment, say, of the elderly and the mentally retarded. The suggestion of the model of specification, by contrast, is that we must not assume that the principles mentioned are all fixed in their content, leaving us only to understand the terms of

50. Cf. the principles invoked by Laurence B. McCullough and Catherine Myser, "Recent Developments in Perinatal and Neonatal Medical Ethics: A US Perspective," *Seminars in Perinatology* 11 (1987): 216–23.

the principles precisely and to describe the case accurately so as to determine correctly which principle applies. I want to illustrate the idea of specification at work in two ways in a possible course of deliberation: first, in helping understand a revision of a norm that is not itself a case of specification, and second, in a specification proper.

Suppose, as a first step, that our deliberator is led by this practical conflict to examine her reasons for accepting the general prohibition on killing, and suppose that these revolve around the notion of respect for persons. What is it about persons, she wonders, that demands respect? These difficult cases involving severely defective infants force this question upon her. Suppose that she decides, on reflection, that the prohibition on killing is, in effect, a specification of a more general principle requiring respect for self-conscious life, and that she is led to revise the prohibition, replacing it with one that specifies the underlying norm differently. Suppose, also, that the present practical conflict leads her to qualify the revised prohibition on killing further by reference to the principle of benefit. As a result of these deliberations, she specifies the prohibition on directly killing innocent persons to read "it is generally wrong directly to kill innocent human beings who have attained self-consciousness, and generally wrong directly to kill human beings with the (genetic?) potential to develop self-consciousness who would not be better off dead, but it is not generally wrong directly to kill human beings who meet neither of these criteria." Although its explicit exception prevents this norm from being a specification of the original prohibition (1), it can still claim some support as a specification of the underlying norm of respect for life. It can also be defended more particularly by reference to the fact that it resolves many conflicts among the norms governing the treatment of defective newborns. As always, of course, if resolving these conflicts provokes other ones, this defense of the specification may be rebutted. To reflect this, and taking to heart the maxim that "hard cases make bad law," the deliberator may wish to preface her specifications with a scope limitation, such as "in all cases involving defective newborns. . . ."

Let us suppose, however, that the particular case facing our agent is one in which it does not appear that the infant in question would be better off dead. Of course, one might at this stage seek a resolution by specifying further the tremendously vague notion of "benefit" as it appears in the principle she arrived at in the preceding paragraph; but let us suppose, instead, that she first works with the principle of respect for

parental choice, since there remains a clash between the revised prohibition and the wishes of the parents. In this case, there is a specification readily available in the ethical tradition which will lessen the practical conflict she is facing by restricting the range of the principle of respect for parental choices to choices that themselves express respect for their children. (Compare Kant's version of respect for autonomy.) Accordingly, she may specify this principle as requiring "that one respect the reasonable choices of parents regarding their children so long as they respect the children's rights." She might, further, decide that this narrowed version of respect for parental choice ought to be taken to *replace* the initial, unqualified version, thereby freeing herself from a requirement to respect the wishes of the parent if they would go against the infant's right to life.

While the course of this hypothetical path of deliberation has been highly controversial and considerably oversimplified, let me highlight the two features that give rise to its claim to be a rational process. First, the specifications and respecifications offered do manage to avert a practical conflict among principles—a conflict that in itself seems rationally unacceptable. Since they therefore enhance the fit among our principles, at least within this limited domain, these specifications mark an improvement in coherence that might be counted as being, in itself, strong (though hardly conclusive) grounds for the rationality of this change. Second, although each of these specifications is controversial, their grounds are not some kind of private and nondiscursive perception or intuition. Rather, they rest on grounds open to rational public debate and assessment, such as those arguments resting on the underlying theory of respect for persons. For instance, the specification of the prohibition on killing offered above could be rationally rejected on the basis of an argument to the effect that the grounds of the prohibition are quite independent of its specifying the ideal of respect for persons.

For a second example, I will shift from the tragic to the merely distasteful, and discuss a choice that is the personal counterpart of what is now a major policy question.[51] It is a choice faced by a committed environmentalist—one convinced that he should, within reason, live his life

---

51. In this essay I have generally tried to avoid getting into the complications that arise when the decision to be made is, in an important way, joint, public, or interpersonal. While one might claim that an advantage of, say, some forms of preference utilitarianism is that they leave controversial questions of value specification to individuals, I believe that this is a misleading dodge.

so as to minimize his adverse impact on the ecosystem. And now the momentous question is: Should he use disposable diapers for his baby or cloth ones? He is enough of an environmentalist that we may ignore questions of convenience and the baby's well-being (allergies, dryness, and so on), and focus solely on the relevant environmental values.[52]

Here is one approach our environmentalist might take: he might make a list of all of the kinds of environmental damage relevant to this choice. In using disposables there is tree loss, topsoil lost and species disturbed at the logging site, air pollution from logging vehicles, water pollution from the pulp plant, oil consumed and air polluted during plastics manufacture, energy consumed in disposable diaper production, the bulk of disposable diapers as a strain on available landfill space, and the biological hazards posed by disposing of soiled diapers. In using cloth diapers there are pesticides used in farming cotton, air pollution from the farming vehicles, energy consumption in manufacturing cotton diapers, energy consumption in transporting and washing cotton diapers, the incremental strain on the sewage system of emptying the cloth diapers into the public sewers, and water pollution from the detergent and bleach used in cleaning the cloth diapers. Being compulsive enough to develop such a list, our deliberator is certainly not going to rest content with an intuitive balancing of this complex set of pros and cons. In order to be more systematic, he instead may try to develop an "environmental impact index" that (1) develops a measure of each of these different types of effect and (2) assigns that measure a weight. But how is such a weighting to be rationally defended? The task seems hopeless, and its value merely heuristic. If this factoring into pros and cons is just a more complicated version of intuitive balancing, it might be easier just to go by his gut reaction to the overall choice.

This abortive effort at weighting these different types of environmental harm, however, will not be in vain if it leads our deliberator to reflect on the way he would specify his guiding norm. We started by supposing that his single overriding relevant principle was to protect and preserve the

52. Since the following paragraphs were written, Arthur D. Little Inc., a respected consulting company, published a "life cycle analysis" of the two types of diaper, much along the lines suggested in the text. It concludes that neither type of diaper is clearly better for the environment. Although the study was funded by one of the makers of disposables, the spokesman for one national environmental organization complained only that the study ignored the pesticides used in growing the cotton for cloth diapers. See John Holusha, "Diaper Debate: Cloth or Disposable?" *New York Times*, 14 July 1990, late edition, p. 46.

environment: but what shall he mean by this? Preserve it in what state, from what dangers? Reflection on the various pros and cons can be of heuristic value in helping him further specify his central principle. Thus, while many of the harms on either side seem to cancel out, in a rough way, there remains a salient difference between the two options—one which might be captured by the notions of material flow versus energy expenditure. Disposable diapers come from the forests and the oil reserves and end up in landfills. Cloth diapers stick around, but require a lot of pollution-generating energy consumption to do so. This overall contrast suggests that it will be important, for the purposes of making this choice, whether our environmentalist specifies his leading norm in one of two broadly familiar ways. First, there is the more old-fashioned conception of the conservationist: the supporter of wilderness areas, hiker (or perhaps NRA member), and reader of John Muir, whose notion of preserving the environment centers on the idea of keeping those parts of nature not yet touched by man from becoming disrupted. Second, there is the newer, more urban-oriented and liberal environmentalist who focuses on those parts of the earth where man's influence is already noticeable, is concerned largely with human health, and seeks to minimize the pollution of populated areas. The first specification would lend differential support to the use of cloth diapers, the second to the use of disposables (at least if the biohazard can be contained). How will our deliberator decide which specification of his guiding norm to adopt (supposing he is initially unsure)? Here is where his effort at ranking the particular harms involved, though incomplete and insufficiently precise to yield a single-valued index, will nonetheless help, for it is likely that his pattern of ranking reflects one of these specifications more than the other. This differential could draw his attention to the way he would specify his guiding norm. Once he focuses on this question, it is likely that this more finely specified version of his environmental end will fit with and will help make sense of a broad range of his policy positions. This fit would justify or explain his adopting that specification. If he can specify his guiding norm more finely in one of these two ways, then which diaper option he should choose will become relatively obvious.

To be sure, once he has settled upon one of these specifications—or perhaps upon a more complex one that crafts a compromise—this would provide a better basis for constructing an index. Obviously, however, the resulting weighting would not be one that was antecedent to the choice.

Furthermore, since other choices might call for a still finer grained spec-
ification, or one that qualified along a different axis, it would be unwise
to place too much confidence in the weighting that seems implied by this
particular choice.

In this second example, as in the first, the claim to rationality for the
specification stems from the possibility of setting out discursively the
reasons for holding that it yields a better fit than its rival. Here, there is
no conflict averted; but where it had seemed that the choice would have
to be made purely intuitively, we have instead an argument that is open
to assessment. The claim, in this example as in the prior one, is that one
can see how the specifications involved count as courses of reasoning—
not that one can see that rationality dictates a unique answer.

VI. CONCLUSION

I have argued that the model of specification is not only distinct from the
pure models of application and balancing and their hybrids, but also su-
perior to each of them. It deals constructively with moral conflicts rather
than being stymied by them, as is the pure model of application; and it
has a claim to proceeding by discursive rationality that intuitive balanc-
ing cannot share. It benefits from a considerable degree of casuistical
flexibility without sacrificing a potentially intimate tie to guiding theo-
ries; and it is able to proceed from norms looser and hence more accept-
able than the completely universal ones required by the deductivist to
reach a conclusion through a Peripatetic syllogism. By showing, without
relying upon universal norms, how a theory might remain stable in the
face of conflicts, the model of specification also undercuts the argument
for ethical skepticism that departs from the fact that norms will conflict.

The decisive advantage of the model of specification lies in its attitude
to conflicts of norms. Whereas the pure model of application is bedeviled
by conflicts and the pure model of intuitive balancing sails through them
untouched, the model of specification learns from the conflicts it faces,
exploiting their friction to push off toward a more concrete and definite
understanding of the relevant norms. If we assume that our norms can-
not be absolute, then they will naturally be qualified in the course of
resolving concrete ethical problems. The model of specification searches
out qualifications that are specific to the content of the norm being spec-
ified (recall the relevance requirement), tailored to the situation being
addressed, and articulate as to their rationale.

While the superiority of the model of specification's attitude to con-
flicts over those of the pure models of application and balancing is per-
haps obvious, its superiority in this respect in comparison to a hybrid
such as Ross's, which combines these two operations, needs more elab-
oration. I have already commented upon the superior ability of the model
of specification to account for the phenomenon of general norms "hang-
ing around" even though they are not followed in a given case. Beyond
this, however, is a deeper difference in attitude. This hybrid model has
served the very important and constructive function of convincing many
that we need not cast aside a nonconsequentialist attachment to moral
principles simply on account of counterexamples to each principle taken
absolutely—for these can all too readily be generated from our consid-
ered judgments about hypothetical and real examples. Shifting from an
"absolute" to a "prima facie" interpretation of moral principles allows the
nonconsequentialist to survive in the face of such objections. But in the
field of medical ethics, at least, it would seem that this hybrid model has
won this battle. It turns out, however, that the prima facie principles that
have been put forward in this field settle few questions by themselves:
rather, one gets a sense of a pervasive conflict between, say, autonomy
and beneficence. Because of the nature of the hybrid model's framework,
there is little theoretical incentive to refine these prima facie principles
in a way that expands the constructive role of discursive reasoning. This
lassitude derives from the very feature that gave Rossian intuitionism
such appeal in the first place, namely, its ability to brush aside counter-
examples. A Rossian intuitionist need not worry about presenting a very
vague and overly inclusive interpretation of "autonomy" as "self-deter-
mination," for instance, because when a concrete case arises in which
the demands of autonomy so interpreted seem unacceptable, it is highly
likely that this unacceptability can be explained in terms of a clash with
another prima facie principle, therefore depriving any principle of
straightforward application. Combine this universal stopgap with the
claim that "the decision rests with perception" when the various prima
facie principles conflict, and you can see that this hybrid approach offers
no logical requirement to go any further than listing a few prima facie
principles, the exact contours of which one need not worry much about.
If one does start to worry about these—wondering, for instance, why au-
tonomy in the control of one's body is more important than autonomy in
the control over the use of one's surgical capacities (say, when there is a
conflict over whether the surgeon should amputate a patient's leg), then

one has started down the road of specification.[53] Many of the most thoughtful writers on concrete ethics do just this. They should then recognize that Ross's hybrid model no longer adequately describes what they are doing.

In contrast to its alternatives, as I have stressed, the model of specification typically uses practical conflicts, in which unacceptable implications of norms arise, as occasions that give one *pro tanto* grounds for more finely specifying these norms. This feature of the model of specification was apparent in the hypothetical deliberation concerning malformed newborns, in which a more fine-grained interpretation of autonomy was generated in response to the case at hand. It is this sort of further specification that is now needed, I suggest, in relatively well developed areas of concrete ethics such as medical ethics. Unless we can learn from our ethical conflicts in this way, our prospects for a reasonable and rational treatment of the problems of concrete ethics are dim indeed.

53. I doubt that this differential between the body and the use of professional capacities may be deduced from the concept of autonomy. It seems, rather, to reflect an independent set of norms, of the sort that are embodied in the law of torts.