Journal of Medicine and Philosophy
2000, Vol. 25, No. 3, pp. 285–307

0360-5310/00/2503-0285$15.00
© Swets & Zeitlinger

Specifying, Balancing, and Interpreting Bioethical Principles

Henry S. Richardson
Georgetown University

ABSTRACT

The notion that it is useful to specify norms progressively in order to resolve doubts about what to do, which I developed initially in a 1990 article, has been only partly assimilated by the bioethics literature. The thought is not just that it is helpful to work with relatively specific norms. It is more than that: specification can replace deductive subsumption and balancing. Here I argue against two versions of reliance on balancing that are prominent in recent bioethical discussions. Without meaning to address the substance or the overall merits of either view I criticize, I attack Gert, Culver and Clouser's implicit reliance on some overall dimension of balancing as a basis of resolving conflicts among norms and Beauchamp and Childress's residual acceptance of 'justified balancing'. The former authors' description of resolving conflicts depends upon a type of value commensurability that (as they otherwise seem to admit) does not obtain, while the latter authors' role for justified balancing would be better served by continued specification.

Key words: bioethics, casuistry, principlism, specification

I. INTRODUCTION[1]

A good number of theorists writing on bioethics have come to see the usefulness of the idea that norms may be specified progressively, an idea that I articulated, drawing on Aristotle and Aquinas, in an article published a decade ago (Richardson, 1990). At the risk of seeming ungrateful, I seek now to argue that the lessons of that article, and the potential of the idea of specifying norms, have not been fully assimilated even by the writers on bioethics who have taken it up.[2] In particular, I want to argue against the continued reliance of the bioethics literature on the metaphor of balancing, a metaphor that specification can and should displace.

Correspondence: Henry S. Richardson, Ph.D., Department of Philosophy, Georgetown University, 37th and O Streets NW, Washington, D.C. 20057, U.S.A. E-mail: RICHARDH@ gunet.georgetown.edu.

Downloaded from https://academic.oup.com/jmp/article/25/3/285/903717 by Johns Hopkins University user on 12 August 2024

The premises of my argument are simple and, I hope, relatively noncontroversial. First, I suppose that justification must meet a weak requirement of publicity. Justifications must be offered in terms of reasons that may be publicly stated. While my argument depends only on *potential* public expressibility of reasons, one might well believe in a stronger requirement that the reasons *actually* be publicly expressed or provided, at least in the context of developing normative constraints on public policy, which bioethics does. Second, I suppose that bioethical theory is meant to be action-guiding. While it may also have a role in articulating understandings and crystallizing awareness, it intends to guide actions in the future. Third, I suppose that developing an adequate set of action-guiding principles – at least in a fast-changing context such as bioethics, if not in human life in general – requires the progressive collaboration of many practitioners and theorists, each building on the work of others. These basic premises will be the basis of my argument that the residual place that the idea of balancing holds in bioethical theory should be ceded to the alternative idea of specifying or interpreting norms.

In the next section, I will explain the relation I see between specification and other modes of interpretation. Before I come to that, however, let me distinguish the different residual roles in which balancing is cast even by those bioethical theorists who also see a role for specification. Balancing may enter in either (1) as a feature or implication of the content of a theory's principles or (2) as part of what the theory says about how conflicts among its principles are to be dealt with. On the first possibility, entering as a feature or implication of a theory's principles, balancing may either be (a) piecemeal and contextual or it may be (b) more global or overall. Piecemeal and contextual balancing, as dictated by the content of some principle, is relatively innocuous and unobjectionable, and is not what I am attacking here. Balancing as a mode of conflict resolution and global or overall balancing, by contrast, are far more troubling. In this paper, I take as my first example of a bioethical theory one that, while open to the importance of specifying and otherwise interpreting norms, nonetheless adheres to an objectionable extent to the idea of global or overall balancing. This is the account offered by Bernard Gert, Charles Culver, and K. Danner Clouser in *Bioethics: A Return to Fundamentals* (1997; to be cited as GCC). A theory that, despite a generous endorsement of the importance of progressively specifying principles, persists in relying importantly on balancing in prescribing how to resolve conflicts among principles is that put forward in the fourth edition of Beauchamp and Childress's *Principles of Biomedical Ethics* (1994; to be cited as B&C). I have no rival bioethical theory to offer in place of these. I am not a bioethicist, and I defer

Downloaded from https://academic.oup.com/jmp/article/25/3/285/903717 by Johns Hopkins University user on 12 August 2024

(on principled grounds; see Richardson, 1999b) to those who are knowledgeable about these topics. My target is not one of the overall theories of either set of authors but their use of the idea of balancing ethical principles against each other; it is to the idea of balancing that specification offers a fruitful alternative for those developing bioethical theory.

I proceed as follows: I begin by clarifying the idea of specification and, in particular, its relation to other modes of interpretation. After next comparing the views of Gert, Culver, and Clouser with those of Beauchamp and Childress, highlighting their different treatments of balancing, I take each of them up in turn. I close with a discussion of the acceptable place for piecemeal, contextual balancing as a feature or implication of the content of norms.

II. SPECIFICATION AND OTHER MODES OF INTERPRETATION

The aim of my article on specifying norms was to displace two leading models of how to bring norms to bear on cases: the model of application (or deductive subsumption) and the model of balancing. Far from being intended as a complete moral theory, the model of specification presupposed that one had a theory, or at least an articulated set of norms, already in hand, and asked a question that then arises. This question was how to bring these norms to bear in guiding action, especially, but not only, when those norms conflict. One way is by deductively subsuming a case under a rule; but this demands, unrealistically, that norms be universal generalizations in their logical form.[3] Another way is by situational or perceptive intuition (as in one version of *phronêsis*), but this leaves the reasons for decision unarticulated. Note that what I have just described are ways of bringing norms to bear on concrete situations, not specifically ways of dealing with conflicts among norms. Specification provides an alternative that can work with norms that are not universal generalizations and can help us articulate our reasons in a potentially public way.

The principal rivals of the idea of specification, then, are the ideas of application (or deductive subsumption) and of balancing. Each of these ideas floats relatively free from the content of moral theory. In judging their relative merits, one must take them to be combined with some set of norms, to be put to work in filling out some ethical theory or other.[4] The considerations I will be advancing here, however, are largely independent of the content of moral theory. Instead, I am concerned with different ideas about how one might bring moral theory to bear in settling concrete issues.[5]

Downloaded from https://academic.oup.com/jmp/article/25/3/285/903717 by Johns Hopkins University user on 12 August 2024

If we shift the focus to situations in which two principles or other norms conflict, specification again offers an overlooked alternative. Intuition here becomes intuitive balancing, an intuitive sense of which set of considerations "weighs" more. Corresponding to deductive subsumption is an approach that would seek a fixed way lexically or otherwise to rank principles. Both of these approaches assume that the set of principles remains fixed, and either (a) must already contain a priority rule that resolves the question at hand or (b) is supplemented by intuitive balancing to determine which ones override. Central to the role of specification, however, is that it is defined as a relation between two norms: an initial one and a more specific one that is brought to bear on practice. Specification, then, can sometimes resolve conflicts by filling out – and thereby changing, at least by addition – the set of norms. In a recent article, Robert Veatch appears to deny this third possibility. He writes that "the only other possibility for resolving conflicts among principles [aside from balancing them] is to attempt to rank-order them" (1995, p. 210). When he later comes to a discussion of specification, he assimilates it to these two other possibilities via an assumption that specification will be possible only when the two principles in conflict may be lexically ranked (p. 216). This ignores, however, that a specification may be quite context-specific, and so not even speak to broad aspects of the initial principle's domain, let alone take a position on its relative ranking within those areas.[6]

Specification is not the only mode of modifying principles. In the remainder of this section, I seek to distinguish it from some of the others. Let us assume, together both with Gert, Culver and Clouser and with Beauchamp and Childress, that the main work that needs to be done in order to achieve progress in bioethics at the present time is in the direction of greater concreteness. The work of *abstraction* – of crystallizing a few broad principles from out of a mass of messier materials – has already been done by other philosophers and bioethicists.[7] Even in the direction of greater concreteness, however, not all moves count as specifications, in my sense. It will, I hope, be useful and interesting to distinguish some of the possibilities.

To begin with, we should distinguish between interpreting norms and deriving subordinate norms. Subordinate norms may be derived from initial norms either by deductive subsumption or by less formal causal reasoning. An example of the former kind of derivation would be deriving "do not lie" from "do not do anything with the intent to deceive" plus a definition of lying as stating a falsehood with the intent to deceive. An example of a more informal derivation in Gert, Culver, and Clouser is their derivation of "do not drink and drive" from "do not kill" (GCC, p. 53).

Downloaded from https://academic.oup.com/jmp/article/25/3/285/903717 by Johns Hopkins University user on 12 August 2024

Given causal facts about automobiles, alcohol, human bodies, and the general practice of sharing the roads, drinking and driving poses too high a likelihood of killing to be acceptable. While the prohibition of drunken driving is an important and firm one, there is no need to regard it as *interpreting* the rule against killing. If the relevant causal facts changed sufficiently, we would change our views about drinking and driving. The moral prohibition against drunken driving is thus *derivative* from a prohibition on killing (or on imposing undue risk of injury or death). The initial prohibition itself remains unmodified by the derivation, which merely links it to a conclusion by causal (or conceptual) facts. These links supplement the initial norm without changing it. By contrast, an interpretation, as I'll be using the term, *modifies* the content of a norm.[8]

Specifying norms is but one of at least four modes of interpreting them.[9] To be able to set these possibilities out, I need to mention some further details of my analysis of the specification relation. When I initially defined specification, I was concerned in the first instance, as I have mentioned, with the question of how an initial norm may be brought to bear upon a situation. Accordingly, it was important to set out in a relatively precise way in what relation the more specific norm stood to the initial norm, such that we might be licensed in saying that the initial norm is brought to bear when we use the more specific norm. In order to capture this feature, I invoked the semantic condition of *extensional narrowing*: everything that satisfies the specified norm must also satisfy the initial norm (or, if the norms are not logically absolute, everything that satisfied the absolute counterpart of the one must satisfy the absolute counterpart of the other).[10] Yet to fit the intuitive notion of specification, not every norm that happens to be narrower can count. Accordingly, I added a syntactic condition of *glossing the determinables*. That is, I required that a specification narrow a norm by adding clauses spelling out where, when, why, how, by what means, to whom, or by whom the action is to be done or avoided. You can see from this explanation that there are two kinds of interpretation that do not count as specification but are interesting and important nonetheless, namely moves that simply narrow without glossing, and moves that gloss without narrowing. An example of *narrowing* that does not gloss is the move to "do not torture" from "do not harm." The former is indeed extensionally narrower, but since "torture" is a well-understood notion on its own, there is no need to generate this more specific norm by adding clauses to the initial norm prohibiting harming. An example of *glossing* that does not narrow would be any gloss that purports to replace an initial formulation by definition rather than supplementing it. For instance, "do not have sex in the office" could be glossed as "do not

Downloaded from https://academic.oup.com/jmp/article/25/3/285/903717 by Johns Hopkins University user on 12 August 2024

have sex in the office, by which we mean: do not engage in any act involving contact (of a certain kind) with the genitals of another." This formulation adds words but purports, at least, not to narrow: it is simply spelling out what "having sex" already meant.

Or do we understand sex so definitely? So far, I have been speaking as if the condition of extensional narrowing were non-controversially applicable – as if we both knew in general what it was for norms of various types to be satisfied and when it is that they are. There are, however, both general problems about the former (which I will not go into here) and frequent cases in which our norms are so vague that it is indeterminate whether or not a given type of action would satisfy them. I owe to my colleague Mark Murphy the thought that *sharpening* a vague norm cannot count as a specification because the specification relation depends upon the extensional narrowing condition and whether the narrowing condition is met will be in principle indeterminate if the initial norm is sufficiently vague (as I had been assuming, in the last paragraph, that the prohibition on having sex was not). For example, "do not drink more than twelve beers at a single sitting" cannot count as a specification of "do not drink inordinately" – not only because the former drops the word "inordinately," but also because it is indeterminate whether every action that satisfies the latter will satisfy the former. Perhaps drinking eleven beers at a single sitting is also inordinate. The idea of "drinking inordinately" is too vague to determine whether the proposed interpretation counts as a specification of it or not.

Often we are in this position. Beauchamp and Childress write that "a typical example of a rule that specifies the principle of respect for autonomy by giving it more content is, 'Follow a patient's advance directive whenever it is clear and relevant'" (B&C, 1994, p. 39). Strictly speaking, in the terminology that I am now suggesting, this is a sharpening that may or may not narrow, but it is certainly not a specification. It does not proceed simply by adding clauses to the principle of respect for autonomy. Further, whether all cases of following such advance directives are cases of respecting autonomy may well be, because of the vagueness of the latter notion, too indeterminate to settle.[11] (No more indeterminate, perhaps, than a similar judgment involving the notion of freedom, but indeterminate nonetheless.) The connection that Beauchamp and Childress make out between respecting autonomy and following advance directives is a theoretical achievement that takes a relatively complex argument to set out. To be sure, *once* that connection has been made, it will be plausible to state the resulting principle as "Respect the autonomy of patients by following their advance directives whenever they are clear or relevant." This

Downloaded from https://academic.oup.com/jmp/article/25/3/285/903717 by Johns Hopkins University user on 12 August 2024

does count as a specification of the initial norm. All I am insisting on here is that, by definition, a specification wears its connection back to the initial norm on its face, for my requirement that specifications proceed by adding clauses implies (or was intended to imply) that the initial clauses will remain.[12] In this way, "respect the autonomy of patients" bears a transparent relation back to "respect autonomy." In the absence of such a transparent link, what Beauchamp and Childress put forward is technically an interpretation of autonomy that sharpens it in a certain way for a certain context rather than a specification of the principle of autonomy.

I have distinguished four modes of interpretation that move in the opposite direction of abstraction: specification itself; extensional narrowing and glossing; the components of specification; and sharpening, a prerequisite of specification. Having these in mind will help us as we turn to the ways in which our two sets of bioethics authors take up the idea of specification and, more generally, acknowledge the need for interpreting principles.

III. COMPARING THE TWO VIEWS

From the point of view of this discussion, the respective approaches of Beauchamp and Childress and of Gert, Culver, and Clouser differ saliently in only one major respect. It is a difference that emerges from a background of considerable similarity and complementarity – a background that is important to set out so that we may have the outlines of their views before us.[13] To be sure, my focus is on the abstract structure of their views, and hardly at all on the specific content, let alone on content specific to bioethics. Let me start by mentioning five important structural similarities.

First, both sets of authors put forward *a relatively small number of central principles or rules* that they draw from "our common morality" (B&C, 1994, pp. 6, 37, 101; GCC, 1997, pp. 16-17, 33–34) and that they use to help generate more concrete guides to action. Beauchamp and Childress offer four: autonomy, justice, beneficence, and non-maleficence. Gert, Culver, and Clouser have ten. They factor non-maleficence into five rules, each directed against a particular kind of harm, and add prohibitions on deceiving, breaking promises, cheating, breaking the law, and not doing one's duty (GCC, p. 34). From my point of view I am less interested in the specific lists than in the fact that each group of authors lists a small finite number of independent moral principles or rules.

Second, both sets of authors characterize their rules or principles as representing what W.D. Ross called "*prima facie* duties." That is, each of

Downloaded from https://academic.oup.com/jmp/article/25/3/285/903717 by Johns Hopkins University user on 12 August 2024

the duties they list is subject to being overridden by other moral considerations. Further, both sets of authors describe this process of norms being overridden in a way that invokes the metaphor of weighing or balancing. Beauchamp and Childress align themselves more explicitly and directly with Ross in this respect, putting the point in the quantitative language of balancing: each of the principles holds "unless it conflicts on particular occasion with an equal or stronger obligation" (B&C, pp. 33–36). Gert, Culver, and Clouser do not directly state that it takes a moral consideration to override a moral rule and do not speak of objectively "stronger" obligations; instead, they elaborate a hypothetical standard for when rule violations may be justified. According to them, a moral rule may be violated only if "an impartial rational person [could] advocate that violating it be publicly allowed" (GCC, p. 37).[14] Rationality here requires that the person advocating the violation have an "adequate reason" in mind, as judged by "some significant group" (GCC, pp 26, 28). Impartiality requires not being influenced by which persons are affected by the violations (GCC, p. 31; cf. Gert, 1998, chap. 6). While this combination of rationality and impartiality does not entail that only moral reasons (or what some significant group takes to be such) may override a moral rule, they do tend in that direction.

Third, both sets of authors take an important step beyond Ross's model of balancing conflicting *prima facie* duties by insisting on the importance of *interpreting* the moral principles, in the sense developed in the last section. Both stress that these interpretations must take account of existing social practices without kowtowing to them. Both indicate that it will be necessary to adapt the general principles from which they start so as to fit the special requirements of the biomedical context.[15]

Fourth, both sets of authors indicate that there will be a *need to specify* the principles or norms in the process of interpreting them. In general, their calls for interpretation reflect an awareness of the highly general character of the principles and rules that they invoke and of the gap that therefore arises between them and the guidance of action in particular circumstances. The interpretations aim to make the principles more specific in ways that take account of concrete aspects of the biomedical context. In Beauchamp and Childress, this call for specification is pervasive (B&C, pp. 28–32; and see index *s.v.* "specification"). Gert, Culver, and Clouser stress instead other modes of interpretation, but do take comfort in their rivals' adoption of specification, for they think that specifying the four principles well will mark at least a "way station" on the journey to truth – that is, towards their own alternative view of the moral rules governing bioethics (GCC, p. 91). They distinguish between what they call

Downloaded from https://academic.oup.com/jmp/article/25/3/285/903717 by Johns Hopkins University user on 12 August 2024

'general' moral rules, which all rational persons have reason to follow, and 'particular' moral rules, which are contextual specifications and interpretations of the general rules (GCC, pp. 52–3). General moral rules, they note, are so abstract that they need "culturally sensitive specification" to be usable at all (GCC, p. 55). To be culturally sensitive, this specification will necessarily have to be different in different communities.

Fifth, both sets of authors nonetheless claim that there will be an important and necessary place in the theory for the idea of *balancing* competing considerations. Beauchamp and Childress describe balancing as being of coordinate importance with specification (B&C, pp. 32–4). For Gert, Culver, and Clouser, as we will see, balancing is even more central, on account of their attempt to unify morality around the idea of minimizing harm.

In addition to these schematic similarities, there are ways in which these two views are complementary (as noted in Beauchamp, 1995). This is sometimes obscured by polemical rhetoric in which one of these rivals tasks the other for lacking a certain account. By calling the features I am about to mention 'complementarities', though, I am suggesting that I see no reason that these useful aspects of each account could not be adopted by both – or, indeed, by anyone. I have in mind two pairs of such helpful features.

The first pair of complementary contributions arises from the authors' respective attempts to constrain the reasoning whereby *prima facie* norms get overridden. Gert, Culver, and Clouser offer a helpful checklist of questions to ask in determining whether it is justifiable to violate one of the moral rules (GCC, p. 38). Some of these are quite obvious – e.g., What are the harms and benefits? Others are less so. Some, for instance, are culled from other moral theories: what relationships among people are involved? How does the difference between harm merely foreseen and harm intended as an end or means enter in? Although they cast this list of questions as a list of "the morally relevant features" (GCC, p. 37), this is to oversell it. After all, they themselves go on to utilize a much wider set of morally relevant features, ranging from social expectations of privacy to the proper understanding of death. Still, their list of questions is useful to keep in mind when dealing with conflicts and helpful in explicating what they mean in requiring that all moral agents must be able to understand the rationale of any rule violation. Similarly useful is Beauchamp and Childress's partially overlapping list of "conditions that restrict balancing" (B&C, p. 34). In particular, they supplement Gert, Culver, and Clouser's question about whether there are morally preferable alternatives with demands that the infringement be minimized and its negative effects be

Downloaded from https://academic.oup.com/jmp/article/25/3/285/903717 by Johns Hopkins University user on 12 August 2024

mitigated, also adding that the agent must check the realism of his or her projections of expected harms and benefits.

The second pair of complementary contributions I want to highlight arises in the authors' respective accounts of moral justification. Beauchamp and Childress give a helpful exposition of Rawls's notion of reflective equilibrium (B&C, pp. 20-28), in which justificatory priority is not given in principle to abstract considerations over particular ones or vice versa. As a result, deductive arguments are not privileged over inductive ones. Instead, justification is a matter of making arguments in both directions. Although Gert, Culver, and Clouser give no such exposition of coherence, their methodology is not the "deductivist" one that some have charged it with being.[16] Instead, as I have mentioned, they too view the norms they articulate as having been drawn (inductively) from the common culture. For their part, Beauchamp and Childress lack the stress on public justification that is central in Gert, Culver, and Clouser and central to my argument, below. As the latter authors put it, "To justify morality is to show that morality is the kind of public system that all rational persons would favor as a guide for everyone to follow" (GCC, p. 26). It follows from this, as they indicate, that to be justified in overriding one of the moral rules, one must have grounds that a rational person can impartially and publicly advocate (GCC, p. 37). Putting these two aspects together – reflective equilibrium and publicity – we arrive at Rawls's idea of public justification, in which each rational (and reasonable) individual supports a set of norms in reflective equilibrium, and each publicly accepts as reasonable the basis on which each other person supports them (Rawls, 1996, p. 387). My sense is that this complementarity of views about justification is not merely a conceptual possibility, and that each set of authors would, in fact, endorse their combination. Certainly – and this is important for my later argument – there is no indication that Beauchamp and Childress would dissent from the requirement of publicity, stressed by Gert, Culver, and Clouser.

Finally, I come to what is, from my point of view, the salient difference between the two positions. Despite the plurality of rules they put forward, Gert, Culver, and Clouser claim that their account is unified around the idea of preventing harm and pervasively use the idea of minimizing harm in indicating how they would interpret their principles. "The goal of morality," they write, is "to lessen the overall evil or harm in the world" (GCC, p. 62).[17] Each of their ten moral rules, they claim, is either a prohibition on harming people in a certain way or else is justified instrumentally on the basis of its importance in preventing harms. Thus, they cast the list of 'morally relevant features' as providing instructions for how to

Downloaded from https://academic.oup.com/jmp/article/25/3/285/903717 by Johns Hopkins University user on 12 August 2024

balance harms and benefits (GCC, p. 86). Accordingly, while the 'relevant features' give shape to their discussions of rule violations, entailing that their view does not coincide with any simple consequentialism, their bottom-line question is nonetheless often put in simple, global balancing terms: Is the harm involved in acting against the rule greater than the benefit to be attained by doing so?[18] Beauchamp and Childress, by contrast, place no harm-minimization principle at the core, instead resting content with the degree of unification their four principles provide.

IV. "LEAST HARM" AS A MERELY INTUITIVE STANDARD

In using the idea of harm-minimization to clothe their view with a semblance of unity, Gert, Culver, and Clouser make global balancing a feature of their theory, one that, they claim, is implied, or at least suggested, by the various moral rules.[19] Now the question is whether it is plausible to claim that this balancing might proceed in a way that satisfies the requirement of publicity. I will argue that it cannot.

It matters greatly, in this respect, whether values are commensurable, i.e., whether the reasons against action represented by harms may all be adequately arrayed, for the purposes of deliberation, on a single dimension. On some past theories, which were monistic about value, it would be easy to make sense of the idea of minimizing harm. If harms and benefits were to fall on one dimension, all one would need to do would be to compute where on that dimension an alternative fell. Such theories purport to proceed on a publicly assessable basis. Their problem is not with publicity, but with the facts. Gert, Culver, and Clouser, to their credit, do not succumb to such false claims of commensurability. Instead, they insist that there are five qualitatively distinct dimensions of disvalue, five irreducibly different types of bad: death, pain, disability, loss of freedom, and loss of pleasure. That these are qualitatively distinct and not all measurable on one scale seems correct. If this is accepted, though, then their claim that the idea of minimizing harms unifies morality more strongly than do Beauchamp and Childress will be meaningless unless there is a systematic and publicly explicable way to balance these incommensurable harms.

In fact, however, there is none. Gert, Culver, and Clouser certainly offer none. Instead, they fall back on intuitions in particular cases. Here is one of their examples: Although "wearing an orange necktie with a fuchsia shirt" may displease some people, and one will violate the moral rule against depriving people of pleasure if one wears such a combination with

Downloaded from https://academic.oup.com/jmp/article/25/3/285/903717 by Johns Hopkins University user on 12 August 2024

this intention, such sartorial choices are not, in fact, morally prohibited, because the harm involved in depriving people of the freedom to make such choices is "greater" than the harm involved in the displeasure they cause (GCC, p. 56). Now, we must ask, how do Gert, Culver, and Clouser purport to know this? Have they done research about sartorial freedom and color-clash wincing? No; and, since harm is not one kind of thing, there is not any obvious research that could possibly settle this question. Instead, it is simply supposed to be intuitively obvious which harm is greater. Perhaps we could at least test the case in thought by varying the hypothetical. What if a speaker wore a brightly-colored sixties tie with a relatively plain nineties shirt? What if the speaker wore nothing at all? Would that cause the audience so much displeasure that it would be morally justified to deprive him or her of this freedom? Would we make moral exceptions for people who look like Cindy Crawford or Mel Gibson, the sight of whose nakedness, even on a public stage, might result in a net *increase* in pleasure? I submit that the idea of avoiding what results in the greater harm does not name or suggest a usable or publicly explicable way to settle these sorts of issues.

In fact, as these rhetorical questions begin to suggest, the features that affect the balance of resulting pleasure are by no means the most salient among those that matter to the moral appropriateness of one's attire. Given the social conventions that prevail in most of the world, a proper respect for others dictates wearing *some* clothing, irrespective of the possible pleasure produced by going without. Duties of consideration for others interact complexly with evolving traditions to yield more nuanced and specific norms that forbid, say, wearing beach clothes to a funeral. Quantitative balancing of pleasure and displeasure will not track the distinctions that matter in such cases: the pleasure that someone gets from wearing beach clothes to a funeral is rightly discounted. A further fact confirming that the amount of displeasure is not dispositive in these kinds of case is that the considered judgment with which Gert, Culver and Clouser begin, which favors individual liberty in the choice of wardrobe within some range, is quite insensitive to the size of the audience. If what mattered were the *amount* of displeasure caused by someone wearing an orange tie with a fuchsia shirt, then we ought to be *adding up* the displeasure of each audience member; and if the audience gets big enough, then it would turn out to be a moral violation to wear such a combination. But that, on due reflection, is neither what we think nor how we think.[20]

And in any case, there is no hint of a method, here, for determining which harms are greater than which other ones. This leaves me with the suspicion that what is actually happening is that Gert, Culver, and Clouser

Downloaded from https://academic.oup.com/jmp/article/25/3/285/903717 by Johns Hopkins University user on 12 August 2024

are noticing which moral conclusions we actually come to and then *reading off* from those judgments the assertion that we judge one harm to be greater than another. If this is the case – and this is certainly the direction in which economists working with revealed preference theory proceed[21] – then judgments of relative harm can never provide a *way* of determining which norm overrides. Rather, they provide a way of restating the conclusion that one overrides the other.

If we put this together with Gert, Culver, and Clouser's denial that harm is one kind of thing, what we see is that their claim to have unified morality more than have their rivals is spurious. Like their rivals, they identify four or five independent norms. Like their rivals, they see that each of these can sometimes override the others.[22] Although Gert, Culver, and Clouser's push for unification ends up laying more stress on the language of balancing than do Beauchamp and Childress, the former group has no more to offer by way of a weighing procedure.

This is ironic, given Gert, Culver, and Clouser's healthy emphasis on the publicity of moral justification. The principal difficulty with the metaphor of balancing, when used to cover a judgment made without any quantitative basis, is that it tends to mask the real reasons at work. At the very least, it fails to encourage the articulation of the real reasons, of the kind that a more frankly qualitative account, such as the one I began to sketch regarding nakedness, might bring out. Public justification cannot be built without each person articulating what his or her reasons for judgment are. The metaphor of balancing here provides an excuse for laziness in this regard, and does little else for us. Given the falsity of value commensurability, relying on a principle whose content implies or features global balancing will inevitably clash with the requirement of publicity in these ways: by depending upon intuitive quantitative balancings whose basis cannot be publicly expressed because there is no actual quantitative dimension backing them up and by failing to encourage the public articulation of the actual, qualitative bases of such judgments.

V. MUST WE EVER RESORT TO BALANCING TO RESOLVE CONFLICTS AMONG PRINCIPLES?

The other troubling use of the balancing metaphor is as a suggestion about how to resolve conflicts among principles. Balancing that resolves conflicts either rests upon articulable reasons or it does not. I will use the term 'intuitive balancing' for instances in which the deliberator is unable to articulate his or her reasons for weighing matters differentially. The re-

Downloaded from https://academic.oup.com/jmp/article/25/3/285/903717 by Johns Hopkins University user on 12 August 2024

maining acts of balancing are based on an underlying reason; these are what Beauchamp and Childress call 'justified acts of balancing'. These latter acts, by definition, "entail that good reasons be provided for one's judgment," or at least underlie the judgment implicitly (B&C, p. 33). Accordingly, justified acts of balancing do not directly violate the publicity requirement in the way that intuitive balancing does. Indeed, as defined, justified acts of balancing meet the strong requirement of publicity, which demands that the reasons be actually expressed, not merely publicly expressible. While Gert, Culver, and Clouser are distinctive in the extent to which they rely on (what turns out to be) intuitive balancing, both sets of authors invoke justified balancing of plural considerations as a recommended way of bringing their plural principles to bear on practice. This recommendation I dispute. I will argue that to rely on balancing, rather than specification and other modes of interpretation in dealing with conflicts among principles, is to go against the requirements of the cooperative development of action-guiding theory. Instead of balancing norms, we should specify or otherwise interpret them.

You will underestimate the potential for the various modes of interpretation if you think in terms of a dichotomy between 'interpreted' and 'uninterpreted' norms. As the history of Talmudic interpretation shows, it is possible to keep adding further layers of interpretation. More generally, it is important that all four types of interpretation distinguished above were defined in terms of a *relation* between two norms: a norm is a specification, or a narrowing, or a gloss, or a sharpening *of another norm*. The ideas of specification and its kin are all *relative to an initial norm*. A norm that is a specification of an initial norm may in its turn be specified. Hence, when Gert, Culver, and Clouser complain of what they term my "failure to formulate any procedure for dealing with conflicts between *specified* norms" (GCC, p. 89, emphasis added), they obscure the point that just because the norms that are in conflict are specifications of some other norms, this does not mean that the conflicting norms cannot be *further* specified so as to relieve the conflict. Of course, if by "a procedure for dealing with conflicts" they mean a decision procedure that will guarantee more automatic results than can the process of specification, which rests on deliberative rationality, then I do not believe in any such procedure; nor have they begun to offer one. What allows the idea of specification to offer a third way of reflectively coping with conflicts among principles is the fact that it offers a change in the set of norms. It will be important to keep this possibility of continuous or progressive specification in mind as we turn to the question of whether justified balancing is needed as a distinct mode of addressing conflicts among principles.

Downloaded from https://academic.oup.com/jmp/article/25/3/285/903717 by Johns Hopkins University user on 12 August 2024

To be sure, not all conflicts among moral principles are resolvable. This is something I explicitly noted in my original article on specification. Contrary to what Gert, Culver, and Clouser suggest (GCC, p. 89), I did not there "fail to realize that some moral disagreements are not resolvable" – a failure, they think, which explains the supposed failure (dissolved in the last paragraph) to provide any procedure for dealing with conflicts among specified norms. Setting aside their conflation of conflicts and disagreements, I would note that I explicitly allowed for the possibility of genuine, unresolvable moral dilemmas, i.e., for unresolvable conflicts of the strongest sort (Richardson, 1990, text at n. 48).

As I have noted, Beauchamp and Childress are quite supportive of the idea of progressive specification. They insist, however, that justified balancing ought to be retained as a distinct and complementary mode of dealing with conflicts among principles. Against the suggestion of simply submerging the idea of balancing into that of specifying, they respond as follows:

> Balancing often eventuates in specification, but it need not; and specification often involves balancing, but it might only add details. Accordingly, we do not propose to merge the two methods. The point is that balancing does not compete with specification, and they both coherently augment the model of coherence. We therefore propose that balancing and specification be seamlessly united with a general model of coherence that requires us to defend the reasons we give for actions and norms… [B]alancing is particularly useful for case analysis, and specification for policy development (B&C, p. 34).

In the body of their book, they follow through on this stance, calling in many places for "further specifying and balancing" of principles (B&C, pp. 37, 101–102, 107, 328, 331, 334, 412, 433, 471). The narrow question on which I part company with them is whether balancing is ever to be recommended as a distinct mode of resolving a conflict among principles. In effect, my argument will build on the observation that, in a context such as bioethics, at least, case analysis cannot be separated from policy development.

One way to restate this question about the conflict-resolving role of justified balancing – intuitive balancing having been disposed of by my argument in the last section – is as follows: given that one has a reason for resolving a conflict one way rather than another, what compelling reason might one have for refusing to incorporate that reason into a further specification of one or the other of the competing principles?[23] This, in effect,

Downloaded from https://academic.oup.com/jmp/article/25/3/285/903717 by Johns Hopkins University user on 12 August 2024

is what we saw happen in the sartorial cases of the last section. There, we ended up with something like "one's freedom to wear what one wants in public, despite offense one may give to others, is to be respected so long as one at least wears something." While Beauchamp and Childress are correct to note that justified balancing does not always eventuate in specification, it seems that it always *can*. Consider the following simple example, which is one of the rare instances in which Beauchamp and Childress mention conflict-resolving balancing without linking it to specification: the case of a conflict between the virtues of generosity and tolerance and a duty (perhaps of justice) which calls for outrage or punishment (B&C, p. 67). Now, suppose that the reason or consideration that justifies tipping the balance in favor of punishment is as follows: "While the virtues of generosity and toleration call for us to be generous to all *persons*, some *behaviors* are beyond the pale and demand that we express our outrage and punish them severely." Plainly, however, this reason embeds a distinction that will be usefully incorporated into our interpretation of these virtues. It can generate the following specifying move: from "be generous and tolerant" to "be generous and tolerant towards all persons even when they have transgressed, but towards their behavior only when that behavior is within the pale." This specification, in turn, could well guide a wiser policy of treating transgressors as not irredeemably bad.

Let me give another example where specification is a more fruitful and explicit way of resolving a concrete issue than is balancing. Again, I am not a bioethicist, and do not purport to judge the soundness of the reasoning I will describe in this paragraph. Instead, I offer as an example a speculative reconstruction of some actual reasoning that seems reasonable, even if its results are not correct. What matters, here, are the kinds of moves that specification makes available. Historically, Beauchamp and Childress's principles arose in tandem with the Belmont Report's principles governing research ethics. One issue that has continued to be hotly debated is how to treat research that is both carried out on and intended to benefit children. Children cannot meaningfully give consent; yet sometimes the potential benefits of a proposed research protocol are so great that it seems crazy to block it solely on that basis. Because the Belmont principles are so highly entrenched in this area, it sometimes seems to the commentators writing on these issues that the only question is whether the principle of autonomy is here to be 'balanced' against the principle of beneficence (or perhaps justice, if the relevant group of beneficiary children is not well off) or whether, instead, the principles must be 'ranked' against each other in some lexical fashion that prohibits these trade-offs within certain ranges.[24] To approach these questions using this dichotomy,

Downloaded from https://academic.oup.com/jmp/article/25/3/285/903717 by Johns Hopkins University user on 12 August 2024

however, is to fail to make explicit the normative stance that pervades this whole context, namely, a presumption against the permissibility of research on human subjects.

To make this normative stance explicit, we should recognize that in addition to the general principles of autonomy, beneficence/non-maleficence, and justice, we also (if we accept the content of the Belmont report or something approximating it) start out with a principle that unites them and specifies how they are to be brought to bear on the research context. It has the form, "It is impermissible to engage in research on human subjects unless the principles of autonomy, beneficence, and justice are adequately satisfied." Call this the 'protean research-limiting principle'.

With reference to this protean principle, we may fruitfully recast the debate about whether the general principles of autonomy, beneficence, and justice are to be 'ranked' or 'balanced'. The idea of specification can help us articulate the different interpretive alternatives, which show up in two possible ways of specifying the protean principle's vague notion of 'adequate satisfaction'. The less restrictive specification is: "It is impermissible to engage in research on human subjects unless the principles of autonomy, beneficence, and justice are *satisfied on balance*." The more restrictive specification is the following: "It is impermissible to engage in research on human subjects unless *we do so in a way that respects their autonomy, proceeds justly, does no (intentional?) harm, and produces (significant) benefits*." Call this "the restrictive research-limiting principle."

Once the research-limiting principle is spelled out, however, we can see that the debate about whether to allow trade-offs among autonomy and beneficence in the case of research on children can be recast as a debate about whether to qualify or specify this one principle in a way that is relevant to the differences between children and adults. The commission apparently did make a distinction between children and adults, a distinction that showed up in different ways as they specified the restrictive research-limiting principle for the two cases. In neither case did they take the principle to be absolute, or to apply without qualification. One qualification had to do with degree of risk. In the case of adults, this specificatory qualification, in effect, read as follows: "It is impermissible to engage in research *posing more than minimal risk* to human subjects unless…." Given the inability of children to give meaningful consent, however, the restrictive research-limiting principle, so specified, still conflicted sharply with the aim of benefiting children. Accordingly, the National Commission put forward a tentative compromise that in effect attempted to resolve this residual conflict by specifying the restrictive research-limiting princi-

Downloaded from https://academic.oup.com/jmp/article/25/3/285/903717 by Johns Hopkins University user on 12 August 2024

ple differently for the case of children. In effect, what they suggested was a principle beginning as follows: "It is impermissible to engage in research posing *more than a minor increase over minimal risks* to human subjects *who are children*, unless…" The clause pertaining to expected benefit was compensatingly beefed up.

As is plain from this example, the bare idea of specification does not indicate how one ought to specify principles so as to resolve a concrete problem, nor is it meant to. Multiple alternative, incompatible resolutions might be reached by specifying (as by intuitively balancing). Being discursively explicit, however, specifications can be defended on the basis of reflective equilibrium: by making arguments that show how they may be supported by their fit with what we continue to believe on due reflection. Having a substantive ethical theory that one is working out will be practically indispensable in making strong connections of this kind, and will obviously help narrow down among possible specifications. Conversely, making explicit connections by specifying norms is indispensable to making progress in ethical theory. The kind of chain of specification exemplified by my reconstructed National Commission example enables one to connect conclusions to initial principles while at the same time developing more nuanced and definite guidance that responds to what are taken to be morally relevant distinctions among different concrete situations.

Can we imagine any reason that justifies a balancing outcome that *could not* be incorporated into a specification in this way? One apparent possibility is that some third norm come in as a contingent tie-breaker. Suppose that we are undecided as to whether it is permissible for the pauper to steal medicine from the pharmacist in order to save his wife's life. If so, then the fact that the pharmacist happens to have promised his entire supply of that medicine to someone else may tip the balance. Further, since the connection between stealing and promise-breaking is, here, entirely contingent, it seems hasty to try to build it into some further specification of our norms. But this only appears to be a case of justified balancing. It obtains that appearance from the speciously precise suggestion that one began with some sort of quantitative tie. Once that false suggestion is subtracted, it is no longer apparent that two principles here win out over one. To settle that question, one is thrown back on intuitive balancing after all. To be sure, if the compelling aspect of the additional consideration is not the promise, *per se*, but the fact that someone else, apparently, needs this life-saving medication, then a possible specification is, after all, in the offing: "One may not steal medicine to save the life of someone one loves when doing so deprives someone else, legally entitled to them, of the life-

Downloaded from https://academic.oup.com/jmp/article/25/3/285/903717 by Johns Hopkins University user on 12 August 2024

saving properties of that medicine." I, for one, cannot imagine a case of justified balancing that could not generate a specification.

If we can always specify in cases in which we have justifying reasons for resolving a conflict among bioethical principles, then we should – with only one caveat, to be mentioned shortly. To explain why we should, I revert to the second and third assumptions I stated at the outset of the paper: that bioethics aims to guide action and that its doing so requires the progressive collaboration of many practitioners and theorists. Now, like both sets of authors I have been discussing here, I view the progress of ethical thinking, even when it goes well, as halting and tentative. It is not a matter of monotonically filling out a sketch. There will be erasures. Whole patches will have to be painted over and begun afresh. An important part of the progress, however, will be improvement in sorting out conflicts in advance. To the extent that this is not done, the theory is not guiding action clearly. Indispensable to the progress of the enterprise in sorting out conflicts is the public airing of hypotheses, as we might put it: of candidate specifications and interpretations that are put forward for consideration as ways of resolving like conflicts in the future. In order to contribute to this broader enterprise, then, theorists and deliberators who have the opportunity to articulate the resolution of a conflict in the form of a specified principle ought to do so. Only by doing so will they project their resolution forward to future cases. Therefore, specification, which seems always to be possible in cases in which balancing is justified, should for these reasons always supersede balancing.

The one caveat I want to enter here is this: sometimes we may feel so tentative about the resolution that we have reached in a concrete case that we feel we are not in a position to project it into the future. Sometimes, we will do better just to admit that a problem needs further work. Notice, though, that in such instances, resting with balancing is being commended only weakly, as a second-best outcome.

VI. IS THERE ANY PLACE FOR BALANCING IN BIOETHICS?

Do I, then, recognize no appropriate role for balancing, apart from the caveat just entered? The one type of balancing distinguished at the outset that I have not yet addressed is 'contextual' or 'piecemeal' balancing that enters as a feature or implication of the content of an ethical principle. This limited sort of balancing must be allowed.

A good example of this limited place for balancing arises within Beauchamp and Childress's discussion of the principle of nonmaleficence. Es-

Downloaded from https://academic.oup.com/jmp/article/25/3/285/903717 by Johns Hopkins University user on 12 August 2024

sential to medical contexts is the need to interpret this principle in light of the idea of net harm or benefit or in terms of the balance of harms and benefits. Now consider what obligations not to treat might arise from this idea. In the case of competent patients, Beauchamp and Childress suggest, there is no need for the relevant principle to call for us, the moral assessors, to balance the harms and benefits. Instead, emphasis falls on patient consent, the patient being deemed well able to balance the harms and benefits as he or she sees fit. In the case of incompetent patients, however, there seems to be no good alternative to having the moral assessors do their own balancing. By this route, we arrive at a conclusion such that "the burdens [of treatment] can so outweigh the benefits to the incompetent patient that the treatment is wrong rather than optional" (B&C, p. 212). Here we have a principle that does not put forward a global balancing principle; rather, it confines balancing to a special case within the adaptation of the principle of nonmaleficence to the context of medical treatment. Local, context-specific balancing of this kind seems a sound feature of our principles, and one that is difficult to eliminate at this level of abstraction. (It may turn out, however, that *further* specification of this principle will move us away from the appearance of seeking a quantitative "netting out" of burdens and benefits and towards a far more context-sensitive and precedent-informed set of judgments about which sorts of burdens so outweigh prospective benefits that treatment ought not to be pursued.)

I conclude, then, not that there is no appropriate place for balancing in bioethics, but that its place is limited to contexts that are both relatively narrow and shaped by surrounding principles. Global balancing, attractive to some as a way of seeming to bring unity to ethical theory, is ruled out by the principle of publicity, which is especially important in a domain of public-policy ethics such as bioethics, in conjunction with the fact of value incommensurability. Conflict-resolving balancing, whether intuitive or "justified," is ruled out by the superiority of specifying norms (and otherwise interpreting them) in such contexts. This superiority, as I have argued, lies in the greater contribution of specifying and interpreting norms to the overall enterprise of progressively developing action-guiding principles. Contributing to this enterprise, I assumed at the outset, is part of the purpose of work in bioethics.

NOTES

1. An earlier version of this paper was presented at a conference commemorating the twentieth anniversary of the Belmont Report, held at the University of Virginia in

Downloaded from https://academic.oup.com/jmp/article/25/3/285/903717 by Johns Hopkins University user on 12 August 2024

April, 1999. I am grateful to James F. Childress and Harold Shapiro for inviting me to participate. I have benefited from several rounds of useful comments from Tom L. Beauchamp and from the criticisms of James F. Childress, Ezekiel Emanuel, Robert Veatch, and other conference participants.

2. I am, in fact, quite indebted to the writers I will be discussing for having raised many important questions about the idea of specification. Special mention is also due in this connection to DeGrazia (1992). I have benefited from many instructive discussions with Dr. DeGrazia over the years.
3. Unless a norm is universally quantified, in the logicians' sense (beginning "for all…"), deductive subsumption cannot work with it. Contrary to GCC, (p. 89), I did not confuse and am not confusing the absoluteness of logical form, which logicians designate a "universal generalization" (as opposed to an existentially quantified statement), with universality or universalizability across the domain of persons or moral agents.
4. Accordingly, in defending the fruitfulness of specifying, it is not part of my task to develop substantive criteria of moral decision. Gert, Culver, and Clouser write that I "cannot" develop such criteria (GCC, p. 89), and I certainly do not claim that the idea of specification can generate such criteria from thin air. Instead, my question is whether, in bringing to bear morally relevant considerations on practical problems, specification or balancing offers a more productive model of how to proceed.
5. In Richardson (1995), I argue that the idea of specification suggests how to steer between the rocks of consequentialism and deontology; but I do not presume the correctness of this argument here.
6. It is possible that, rather than ignoring the possibility of specifications that distinguish between different contexts, Veatch is implicitly asserting that we do not need to make distinctions among contexts in stating correct moral principles. To the contrary, I believe that we do. For instance, I believe that Rawls (1971) made real moral progress in grappling with the conflict between the vague values of freedom and equality by developing a specification of them that reconciled their demands in a way that was tailored to a specific context, namely the appropriate configuration of the 'basic structure of society'. Whatever one thinks of Rawls's substantive conclusions, one ought to be impressed with the force of some of his arguments, such as that in §48 appealing to the distinction between 'legitimate expectations' which flow from fair institutions and free-standing claims of moral desert, which respond to features specific to the context of issues surrounding the basic structure. He is right to claim that determining when basic socio-economic institutions are fair is a moral question to which fundamentally different considerations apply than the question of what distributive shares are deserved by persons who have cooperated within fair socio-economic institutions.
7. Cf. the characterization of the National Commission's work in GCC (p. 73). In reality, of course, specificatory and abstracting work need to be interweaved.
8. This divergence between interpretation and derivation is related to the logical form of the initial norm. If it is a universal generalization, all instances may be derived from it by deductive subsumption. If it is not, however, interpretive glosses will be needed to indicate which of its potential implications are to be *counted as* following through on the norm. In this way, interpretation modifies norms that are logically loose enough to permit this sort of supplementation. As I argued in Richardson (1990), most of our norms are loose enough for this.
9. Gert, Culver, and Clouser claim that I do "not recognize that in order to make norms culturally sensitive, it is necessary to allow for some degree of interpretation of the

Downloaded from https://academic.oup.com/jmp/article/25/3/285/903717 by Johns Hopkins University user on 12 August 2024

norms" (GCC, p. 89). Now, specification being a mode of interpretation, it is plain that I allow for some degree of interpretation. While my original article did not discuss cultural variations or how to account for them, Nussbaum (1990, p. 235) has developed the idea of 'local specification' to do just that.

10. Gert, Culver, and Clouser accuse me of failing to recognize that norms that are not logically absolute, and hence allow of exceptions, can nonetheless be universal in the sense that "they apply to all rational persons" (GCC, p. 89). This accusation is a mistake, which derives from their not noticing that in the relevant passage (Richardson, 1990, pp. 292–5), the only sort of universality I mentioned was not universality over the domain of persons but universality over the domain of acts ("for all acts…").
11. We might push for analysis to firm up our norms sufficiently to allow specification to begin. What 'analysis' should mean, in this context, is a deep question. For reasons to think that autonomy cannot be given an analysis in terms of necessary and sufficient conditions, see Dworkin (1988, esp. p. 6).
12. To solidify this implication, I should have spelled out that the only syntactic changes to the norm are those that flow from glossing the determinables.
13. For additional reasons to regard the two views as complementary, see Beauchamp (1995).
14. For more detail, see Gert (1998, ch. 9).
15. As I will suggest below, some of the moves that Beauchamp and Childress describe as specifications I would count as interpretations of a more generic kind. Gert, Culver, and Clouser discuss the interpretation of the rules in GCC, pp. 54-60.
16. Beauchamp and Childress acknowledge this (B&C, p. 20). Gert and Clouser are charged with "deductivism" by, e.g., Lustig (1992).
17. Cf. GCC (p. 80) on moral rules and moral ideals. That their theory is unified around the idea of minimizing harm is the only apparent basis for Gert, Culver, and Clouser's claim (GCC, pp. 20, 88) that they recognize the unity of morality in a way that Beauchamp and Childress do not.
18. Cf. GCC, pp. 56, 58, and 245. On p. 58, Gert, Culver, and Clouser indicate that the calculation of benefits and harms is to be supplemented by a potentially qualitative public consideration of the acceptability of a proposed modification to the rules. In this paper, I am in effect arguing that the latter standard of acceptability on due reflection, or of reflective equilibrium, should be taken to supersede any purported notion of global, quantitative balancing.
19. For present purposes, I take on face value Gert, Culver, and Clouser's claim to have unified their theory in this way. For doubts as to whether they actually succeed in this, see my (1999b).
20. For a nuanced discussion of the limitations on aggregative thinking, see Scanlon (1998, chap. 5). For elaboration of the suggestion that moral theory should draw on modes of deliberation that we accept on reflection as well as on judgments that we accept on reflection, see the outset of Richardson (1995).
21. See my criticism of preference-based models of deliberation (Richardson, 1994, sec. 15).
22. Gert, Culver, and Clouser even say, at one point, that moral ideals can sometimes override the moral rules (GCC, p. 21) – and this despite their heavily criticizing Beauchamp and Childress for failing to take seriously enough the distinction between the two categories of norms (B&C, p. 77).
23. Beauchamp and Childress mention David DeGrazia as having noted the possibility of using the reason that makes the balancing 'justified' to generate a specification in this way (B&C, 1994, p. 34).

Downloaded from https://academic.oup.com/jmp/article/25/3/285/903717 by Johns Hopkins University user on 12 August 2024

24. I am here simplifying a more subtle discussion in Veatch (forthcoming). I also draw for this article the factual basis (such as it is) from my speculative reconstruction of the National Commission's reasoning, below in the text. I am grateful to Professor Veatch for his permission to cite this article.

REFERENCES

Beauchamp, T.L. (1995). 'Principlism and its alleged competitors,' *Kennedy Institute of Ethics Journal* 5, 181–198.

Beauchamp, T.L. and Childress, J.F. (1994). *Principles of Biomedical Ethics*, 4th ed., Oxford University Press, New York.

DeGrazia, D. (1992). 'Moving forward in bioethical theory: Theories, cases, and specified principlism,' *Journal of Medicine and Philosophy* 17, 511–539.

Dworkin, G. (1988). *The Theory and Practice of Autonomy*, Cambridge University Press, Cambridge, England.

Gert, B. (1998). *Morality: Its Nature and Justification*, Oxford University Press, New York.

Gert, B., Culver, C. and Clouser, K.D. (1997). *Bioethics: A Return to Fundamentals,* Oxford University Press, New York.

Lustig, B.A. (1992). 'The method of principlism: A critique of the critique,' *Journal of Medicine and Philosophy* 17, 487–510.

Nussbaum, M.C. (1990). 'Aristotelian social democracy,' in R.B. Douglass, *et al.* (eds.), *Liberalism and the Good*, Routledge, New York, pp. 203–252.

Rawls, J. (1971). *A Theory of Justice*, Harvard University Press, Cambridge.

Rawls, J. (1996). *Political Liberalism*, rev. ed., Columbia University Press, New York.

Richardson, H.S. (1990). 'Specifying norms as a way to resolve concrete ethical problems,' *Philosophy and Public Affairs* 19, 279–310.

Richardson, H.S. (1994). *Practical Reasoning About Final Ends*, Cambridge University Press, Cambridge.

Richardson, H.S. (1995). 'Beyond good and right: Toward a constructive ethical pragmatism,' *Philosophy and Public Affairs* 24, 108–141.

Richardson, H.S. (1999a). 'Bioethics: Root and branch,' *Hastings Center Report*, September-October, 40–42.

Richardson, H.S. (1999b). 'Institutionally divided moral responsibility,' *Social Philosophy and Policy* 16, 218–249.

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

Veatch, R.M. (1995). 'Resolving conflicts among principles: Ranking, balancing, and specifying,' *Kennedy Institute of Ethics Journal* 5, 199–218.

Veatch, R.M. (forthcoming). 'Ranking, balancing, or simultaneity: Resolving conflicts among the Belmont principles,' in a volume edited by J.F. Childress and H. Shapiro.

Downloaded from https://academic.oup.com/jmp/article/25/3/285/903717 by Johns Hopkins University user on 12 August 2024