

International Phenomenological Society

Emotion and Moral Judgment

Author(s): Linda Zagzebski

Source: *Philosophy and Phenomenological Research*, Vol. 66, No. 1 (Jan., 2003), pp. 104-124

Published by: [International Phenomenological Society](#)

Stable URL: <http://www.jstor.org/stable/20140516>

Accessed: 21/03/2014 14:58

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at <http://www.jstor.org/page/info/about/policies/terms.jsp>

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.



International Phenomenological Society is collaborating with JSTOR to digitize, preserve and extend access to *Philosophy and Phenomenological Research*.

<http://www.jstor.org>

Emotion and Moral Judgment

LINDA ZAGZEBSKI

University of Oklahoma

This paper argues that an emotion is a state of affectively perceiving its intentional object as falling under a “thick affective concept” A, a concept that combines cognitive and affective aspects in a way that cannot be pulled apart. For example, in a state of pity an object is seen *as pitiful*, where to see something as pitiful is to be in a state that is both cognitive and affective. One way of expressing an emotion is to assert that the intentional object of the emotion falls under the thick affective concept distinctive of the emotion. I argue that the most basic kind of moral judgment is in this category. It has the form “That is A” (pitiful, contemptible, rude, etc.). Such judgments combine the features of cognitivism and motivational judgment internalism, an advantage that explains why we find moral weakness problematic in spite of its ubiquity. I then outline a process I call “thinning” the judgment, which explains how moral strength, weakness, and apathy arise. I argue that this process is necessary for moral reasoning and communication, in spite of its disadvantage in disengaging the agent’s motivating emotion from the judgment.

I. The problem

It seems impossible for a moral judgment to have all of the features we think it should have. On the one hand, common sense and the structure of moral language lead us to think that a moral judgment is propositional in form, with a truth value,¹ and that to make a moral judgment is to assert that proposition. It follows that when we make a moral judgment we are in a cognitive state. But that seems to conflict with another important feature of a moral judgment: it is intrinsically motivating. This feature is implied by our practices of moral persuasion. If we want to convince someone to act a certain way for moral reasons, we direct our efforts towards convincing her of the truth of a particular moral judgment. As long as we can get her to make the judgment unreservedly herself, we normally think that she will thereby be motivated to act on it. Of course, we know that she may not be *sufficiently* motivated to act on it because she may also have contrary motives, but the point is that we think that all we need do to get her to feel a motive to act on

¹ And it is not always false. This intuition is therefore incompatible with an error theory like that of J.L. Mackie in *Ethics: Inventing Right and Wrong* (Harmondsworth: Penguin, 1977).

a moral judgment is to get her to make the judgment. I will call the thesis that moral judgments are intrinsically motivating motivational judgment internalism (MJI).²

So we expect a moral judgment to be both cognitive and motivating. But it cannot have both features on the standard Humean view of human psychology according to which no state of taking the world to be a certain way can move us to act, or even incline us to act. Some other psychic state must be conjoined with a cognitive state to do so.³ The standard psychology, then, maintains that cognitive or representational states are necessarily distinct from motivating states. When this view is combined with cognitivism, the thesis that moral judgments are cognitive, it follows that moral judgments are not intrinsically motivating. Cognitivists therefore have the problem of explaining the intuitive attraction to motivational judgment internalism. Non-cognitivists can explain why moral judgments are motivating, but their problem is to explain why they appear to be cognitive. Neither position is altogether satisfying, but we seem to be forced to choose between them.

My position is that we should give up neither the position that moral judgments are cognitive, nor the position that moral judgments in a central group of cases are motivating. There is an important set of problems that highlights our expectation that a moral judgment is both cognitive and motivating, together with the ways in which the two aspects can come apart. This is the problem of moral weakness and the related problems of moral strength and moral apathy. In this set of problems moral weakness, or *akrasia*, is by far the most frequently discussed. What is important about moral weakness is that it seems as if it should not exist, and yet we know that it does. It often happens that a person makes a moral judgment while understanding the judgment perfectly well, and yet lacks sufficient motivation to act on it. Notice that if the making of a moral judgment is a purely cognitive state this is no mystery on the standard psychology since on that view motivation does not come from the judgment anyway. But that does not solve the problem of moral weakness; it simply refuses to see it as a problem. Ever since Aristotle

² This term is used by Russ Shafer-Landau in "Moral Judgment and Moral Motivation," *Philosophical Quarterly* 48, #192 (July 1998), 353-358. David Brink calls this position "appraiser internalism" in *Moral Realism and the Foundations of Ethics* (Cambridge: Cambridge University Press, 1989), p. 40. Some writers prefer to simply call it "motivation internalism," for example, Terence Cuneo in "Are Moral Qualities Response Dependent?," *Noûs* XXXV, 4 (December 2001), 569-591.

³ The older rationalist theories maintained that reason itself provides motivation. See Korsgaard, "Skepticism About Practical Reason" (*Journal of Philosophy* 1986, 5-25) for a discussion of this position. Korsgaard says that the rationalists gave no argument to support their assertion. For a recent defense of the rationalist position on motivation, see Jonathan Dancy, *Moral Reasons* (Blackwell, 1993). For the position that all beliefs require assistance, but certain beliefs entail a motivational disposition, see Philip Clark, "What Goes Without Saying in Metaethics," *Philosophy and Phenomenological Research* LX, 2 (March 2000), 375-379.

the existence of *akrasia* has been treated as something that needs explaining because in some sense we think Socrates in the *Protagoras* must have been right: If a person *really* understands what she is doing when she makes a moral judgment, she would automatically feel motivated to act on it and, in fact, she *would* act on it and would not have to struggle to do so. The Socratic position is even stronger than MJI, but the fact that we nonetheless feel its force shows that the connection we expect between moral judgment and motivation is very strong indeed. This indicates that even though we should admit that our worry about the existence of moral weakness shows we are confused about moral motivation, our worry is not misplaced. That is, we are right to admit that akratic behavior exists, but we are also right to be surprised.

Moral strength is a problem for the same reason, although this is not always recognized.⁴ The real difficulty is not that people often act *against* their moral judgment, but that they sometimes have to struggle before acting. Sometimes the judgment wins and sometimes it loses, but neither the winning nor the losing is the problem. The struggle itself is what strikes us as puzzling. But struggle is not puzzling at all on the cognitivist position.

Emotivist theories can explain why moral strength and weakness are problematic since they see moral judgment as intrinsically motivating, a motivation that nonetheless can be overridden by other motives. Non-cognitivism at its best, then, has an advantage with respect to the strength and weakness puzzle. However, non-cognitivists face a related problem, the problem of moral apathy.⁵ The morally apathetic person can make a moral judgment while completely lacking any motive to act on it. Such behavior is no doubt less common than akratic behavior, but it probably does exist. Moral apathy is a problem for cognitivism in the same way moral struggle is a problem, i.e., cognitivists who retain the standard psychology cannot explain why it is puzzling. But non-cognitivists cannot explain how it can occur at all. There should be no such thing as apathy according to non-cognitivism in so far as it takes the motivational force of a moral judgment to be an essential feature of each such judgment. So while cognitivist theories make it too easy to understand how moral apathy can occur, non-cognitivists make it too hard.

John McDowell's way of handling *akrasia* goes part of the way towards rejecting the standard view on the separation of cognitive and motivating states. He argues that the virtuous person has special sensitivity to the moral requirements of a situation in a way that is motivating, but which counts as a form of knowledge and so is cognitive. But McDowell realizes that if he

⁴ Donald Davidson recognized it in his classic paper, "How is Weakness of the will Possible?" (1970), reprinted in *Essays on Actions and Events* (Oxford: Clarendon Press, 1980).

⁵ Alfred Mele calls this problem "moral listlessness" in "Internalist Moral Cognitivism and Listlessness," *Ethics* 106 (July 1996), 727-753. Michael Stocker discusses examples of moral apathy in "Desiring the Bad," *Journal of Philosophy* 76 (1979), 738-753.

wants to hang on to the idea that virtue is knowledge, he has a problem explaining *akrasia* since the akratic person appears to know the same thing as the virtuous.⁶ His solution is to accept Aristotle's position that whereas the virtuous and the akratic know the same thing in a sense, the latter does not appreciate the morally salient features of the situation because his judgment is clouded by desire. The virtuous person's perception of saliences silences considerations that are not silenced for the continent or akratic person.⁷

But there is a problem here. The perception of a salience is presumably a cognitive state. I find the notion of salience rather vague, but I assume that the salient features of a situation are those that stand out or are especially vivid. If so, what makes the virtuous person's judgment motivating and the akratic person's similar judgment not motivating is that the former's cognitive state is especially vivid. Cognitive vivacity explains the difference. But those philosophers who find it hard to believe that a cognitive state can motivate are not going to find the view that an especially vivid cognitive state can motivate any more convincing. If cognitivity *per se* does not motivate, it will not become motivating by becoming more intense. What McDowell needs, I think, is the view that perception of a salience, while counting as knowledge, is not a purely cognitive state. He sometimes says the virtuous person "appreciates" what the non-virtuous does not even when they make the same judgment. And presumably, the converse holds as well. The akratic or otherwise non-virtuous person appreciates aspects of a situation that the virtuous person "silences". So the compassionate person appreciates the fact that a co-worker's child is gravely ill in a way that the discompassionate does not. The envious person appreciates the fact that a co-worker's salary is higher than his own in a way that the non-envious does not. This suggests that to appreciate is not just to understand, but to *feel the force* of that which is appreciated. But this way of putting the view makes it hard to resist the conclusion that the virtuous and the non-virtuous simply respond differently to the same cognitions. And that takes us back to the standard division between cognitive and affective states, adding that there is an important causal relation between them that differs with the degree of virtue of the agent. In that case the problem we are addressing in this paper remains.

I conclude that what is needed in an account of moral judgment is an explanation of how central examples of moral judgment are both cognitive

⁶ McDowell, "Virtue and Reason," *The Monist* 62 (1979), 331-350; reprinted in *Virtue Ethics*, edited by Roger Crisp and Michael Slote (New York: Oxford University Press, 1997), 141-162.

⁷ McDowell makes essentially the same point in "Are Moral Requirements Hypothetical Imperatives?" *Proceedings of the Aristotelian Society, Sup. vol. 52* (1978), p. 28: "The way out is to attenuate the degree to which the continent or incontinent person's conception of a situation matches that of a virtuous person. Their inclinations are aroused, as the virtuous person's are not, by their awareness of competing attractions: a lively desire clouds or blurs the focus of their attention on 'the noble'."

and intrinsically motivating, but which also includes an account of how the motivational aspect of the judgment can be detached from it. Given the difference between cases of moral weakness, in which motivation is present but is not sufficient to counter other motives, and cases of moral apathy, in which motivation is reduced to none at all, it is preferable if the account explains how motivation is detached in stages. Socrates' insight suggests that there is an important and central class of moral judgments in which strong motivation is intrinsic to the judgment. But in many other cases the judgment is accompanied by a weaker and often insufficient motivation, and in a few cases the motivation disappears entirely.

In what follows I will begin with an account of what I call a ground level moral judgment, a judgment that is both motivational and cognitive. Specifically, I will propose that a ground level judgment has the following features: (1) It is an expression of emotion, and for that reason is intrinsically motivating. (2) It is propositional in form, with a truth value (and it is not always false), and the agent asserts that proposition when making the judgment. The account I will give is compatible with the view that a moral judgment is primarily about some aspect of the world outside the mind of the judge, so the objectivity of moral judgment has not been precluded, but I will not enter the debate over moral objectivism or realism in this paper.⁸ My account requires rejecting the standard view that cognitive and affective states are never necessarily connected, and I will propose that the structure of emotion requires that they be so connected. I will then propose that there is a process I call "thinning" the judgment, from which arises the possibility of making such judgments with very weak motivation or even in the absence of motivation. I will argue that this process is a normal one since it is necessary for making moral judgments that are not ground level, but it also makes moral strength, weakness, and apathy possible.

II. *The structure of ground level moral judgments*

A. *Preliminaries*

The class of moral judgments I want to focus on are what I call ground level moral judgments. These judgments can be identified by two important features. First, they are "here and now" judgments. That is, they are made in a context in which the person making the judgment is confronted with the object of the judgment. Second, they utilize what I call "thick affective concepts." These concepts mostly coincide with what Bernard Williams calls "thick evaluative concepts." Examples of ground level judgments include, "She is pitiful," "He is contemptible," "That remark is rude," "That's a lie."

⁸ Theories that are not willing to say that what a moral judgment is about is something in the world outside the mind of the judge, but is instead about an affective state of some kind in the mind of the judge are cognitivist, but not realist about moral phenomena.

Williams claims that thick evaluative concepts combine descriptive and evaluative aspects in a way that cannot be pulled apart, but as far as I know, he does not demonstrate that they have this feature.⁹ I want to make the parallel claim, that concepts like *pitiful*, *contemptible*, *rude*, and *lie* have both cognitive and affective aspects that cannot be pulled apart, and will defend that claim in what follows.

I call the type of moral judgment I want to discuss ground level because I think it enjoys a certain primacy both conceptually and in the learning of moral behavior, but I cannot demonstrate that without giving a full theory of moral judgment, of which this is only a preliminary sketch. My conjecture is that the type of moral judgment I call ground level ends the regress of moral justification and can serve as the foundation of a viable foundationalist ethical theory.

One more preliminary point: The class of judgments I am discussing are evaluative, but they are not restricted to the moral since I do not think that the structure of a judgment such as “That’s a lie” differs from a judgment like “That’s rude,” or at least, I do not think that the differences affect the problem of this paper.

B. *Emotion and thick concepts*

I claim that an emotion is a unitary psychic state that is both cognitive and affective, where the cognitive and affective aspects are not separable states.¹⁰ There are many theories of emotion according to which an emotion has a cognitive aspect as well as an affective aspect, but what I wish to stress here is the inseparability of the two aspects. “Emotion” is not simply a word for an amalgam of two more basic states. What makes the distinctness thesis

⁹ The idea of thick ethical concepts is best known from its use by Bernard Williams in *Ethics and the Limits of Philosophy* (London, 1985). He says the idea came from Clifford Geertz in *The Interpretation of Cultures* (Basic Books, 1973), who in turn credits Gilbert Ryle in the essays “Thinking and Reflecting,” and “The Thinking of Thoughts,” *Collected Papers*, vol 2 (London: Hutchinson, N.Y.: Barnes and Noble, 1971). But my reading of Ryle does not reveal that he was using a concept whose descriptive and evaluative components cannot be pulled apart, the most important feature of the concept as used by Williams, and I’m not even sure about its use by Geertz. Subsequent commentators on the distinction sometimes misrepresent the thesis as the claim that all concepts that combine the descriptive and the evaluative are such that the two aspects cannot be pulled apart. For example, Simon Blackburn objects to the view that there are thick concepts by discussing “lewd” and derogatory words for “fat.” Blackburn claims that the descriptive and evaluative aspects of these terms can be pulled apart, and he might be right about these terms. But that does not show that there are no thick concepts in the sense I mean. At best it shows that these terms are not among them. See “Through Thick and Thin,” *Proceedings of the Aristotelian Society*, Sup. vol. 66 (1992), 285-299.

¹⁰ To my ears the term “state” already suggests separability, whereas “aspects” clearly does not. Claiming that a moral judgment combines cognitive and affective states loads the issue in favor of separability. But not everything cognitive need be a state, and even more obviously, not everything affective need be a state.

tempting even to many of those theorists who recognize both cognitive and affective components in emotion is that a cognitive state similar to the cognitive aspect of an emotion can exist apart from the affective aspect of the emotion, and an affective state similar to the affective aspect of an emotion can exist apart from the cognitive aspect.¹¹

The thesis that cognitive and affective states are necessarily distinct is generally accompanied by the view that affective states are responses to cognitive or perceptual states. That is, we have feelings in response to a prior state of representing the world a certain way. And it cannot be denied that some feeling states are nothing but responses to a prior representational state. But not all are, and this is important. To see the difference, let me begin by comparing two concepts, *nauseating* and *rude*.

Nauseating does not have any descriptive content other than *whatever causes a feeling of nausea*. Hence, there is no limit in principle on the kinds of thing that can be nauseating. Whatever a person finds nauseating *is* nauseating. Of course, if someone systematically applies the term “nauseating” to very different things than other persons do, we may doubt that she understands the word. But the point is that the mere fact that someone occasionally calls something “nauseating” that most other persons do not neither shows that she misunderstands the meaning of the word nor that she is speaking falsely. If it causes her a feeling of nausea, it *is* nauseating—to her, no matter what else is true of it descriptively.

Rude is different. As Philippa Foot has argued,¹² “rude” only applies when certain descriptive facts apply, facts that are not limited to causal facts. These descriptions do include a causal property—that the behavior gives offense, but

¹¹ Even the friends of the relation between emotion and moral judgment generally accept the standard psychology, usually maintaining that an emotion is an appropriate response to something perceived or cognized. David Wiggins has a view of this kind, arguing that something in an object *merits* a certain feeling. As he puts it, a property and an attitude are “made for each other.” (“A Sensible Subjectivism,” in *Needs, Values, and Truth*, Oxford: Blackwell, 1987, p. 199). Likewise, Charles Taylor contrasts the feeling of nausea with the feeling of respect for persons by saying that in the latter case the object merits our reaction, whereas in the former, the connection between the two is a brute fact. (*Sources of the Self*, Harvard University Press, 1989, p. 6). In earlier continental philosophy versions of the appropriate response view of emotion were defended by Brentano, Scheler, and von Hildebrand, but ironically, a similar view is implied by Hume’s position that there is a natural fit between the emotion and the object: “It must be allowed that there are certain qualities in objects which are fitted by nature to produce particular... feelings.” (*Of the Standard of Taste*, para. 16). For the continental philosophers mentioned see Franz Brentano, *On the Origin of Our Knowledge of Right and Wrong* (N.Y.: Humanities Press, 1969); Max Scheler, *Formalism in Ethics and Non-formal Ethics of Values*, tr. by Manfred S. Frings and Roger L. Funk (Evanston: Northwestern University Press, 1973); and Dietrich von Hildebrand, *Christian Ethics* (N.Y.: David McKay Co, 1953).

¹² “Moral Arguments,” in *Virtues and Vices and Other Essays in Moral Philosophy*, Berkeley and Los Angeles: University of California Press, 1978, p. 102.

Foot specifies that what makes rude behavior rude is that it gives offense *because* it indicates lack of respect. Now Foot may be mistaken about the precise descriptions that apply in cases of rudeness. In particular, it seems to me that actually giving offense is neither necessary nor sufficient for being rude, but my point is not about the way in which Foot is wrong, but about the way in which she is right because clearly, rudeness has something to do with both giving offense and expressing a certain attitude—either disrespect or something similar. It does not matter what the precise attitude is. The point I want to make is that “rude” applies to only a certain range of behavior.¹³ A person can get it wrong. She may take offense at something even though it is not rude, and I would add, although Foot would not, that it may be rude even though she does not take offense. In contrast, if she feels nauseated by something, it is nauseating, and if she does not feel nauseated by it, it is not nauseating. This difference is expressed in a difference in the grammar of “rude” and “nauseating.” We say that something is or is not nauseating *to her*, but we do not add the words “to her” when asserting that something is rude.

So *nauseating* and *rude* are concepts that are related in different ways to the descriptive situations in which they apply and to the subjective states they bring about. As an aside, let me say that I find *offensive* an in between case. Is there any limit, in principle, to what can be offensive? Is *offensive* like *nauseating* or is it like *rude*? *Giving offense* seems to be like *nauseating*, but I’m not sure about *offensive*, which, like *rude*, might apply even when no actual offense is given. So *offensive* may be midway between *nauseating* and *rude*. Or perhaps it is not midway, but is ambiguous between a sense like *nauseating* and a sense like *rude*.¹⁴

Let us look a little more carefully at the contrast between *nauseating* and *rude*. I have said that *nauseating* applies to anything that causes a certain subjective state, the feeling of nausea, and I have claimed that there is in principle no limitation on the kinds of things that can cause nausea. It follows that if I call something nauseating, I am not calling attention to any descriptive features of it other than the fact that I feel nauseated by it. *Rude* is not like that. There is in principle few, if any, limitations on what can cause me to feel offended, just as there is in principle few limitations on what can cause me to feel nauseated. But not just anything is rude, even if it offends

¹³ Michael Stocker gives objections to Foot’s claim about the relation between rudeness and giving offense in *Valuing Emotions* (Cambridge University Press, 1996), 74–77.

¹⁴ A special way of being offensive is to be disgusting. Martha Nussbaum has examined the feeling of disgust for the implications it ought to have in the law. From the way in which Nussbaum distinguishes the feeling of disgust and the feeling of indignation, I suspect that she thinks that *disgusting* is much closer to *nauseating* than to *rude*, whereas what one feels indignant at is closer to the rude. See “Secret Sewers of Vice: Disgust and the Law,” in *Emotions and the Law*, edited by Susan Bandes (New York University Press, 2000).

me. Causing nausea is sufficient for something to be nauseating, but causing offense is not sufficient for something to be rude. Only things with certain descriptive properties other than the property of causing offense are rude, and when I call something rude I am, in part, calling attention to those properties, or, to put it more carefully, I am in part calling attention to properties that would not pertain were it not for those descriptive properties.

Someone who accepts these observations might reply that this does not yet show there is anything special about *rude*. If it is not like *nauseating*, then it is just an ordinary descriptive concept like *solid* or *square*, a concept whose meaning has no connection to subjective states at all. But this cannot be right because feeling offended *is* intrinsically connected to the meaning of “rude,” just as feeling nauseated is intrinsically connected to the meaning of “nauseating.” I have said that it is not necessary that each and every rude act produces a feeling of offense. In this respect rude is like red. It is not necessary that each and every red object produces a sensation of red either. Still, the fact that beings with our emotional dispositions normally have a certain feeling in the presence of rude behavior is part of what it means to be rude just as the fact that beings with our perceptual capacities normally have a sensation of red in the presence of red objects is part of what it means to be red. The fact that affect is built into the possession of thick concepts in my sense has the consequence that there are concepts the possession of which is not a purely cognitive state. This distinguishes thick concepts in my sense from the notion of thick evaluative concepts as used by Williams. Williams and others following him make the distinction in terms of the content of the concept, not in terms of the state of the possessors of the concept. I am not denying that the content of a thick concept is more than descriptive, but I think it is more helpful to focus on what it is like to have such concepts and to use them in typical instances. To grasp a thick concept in the sense I am after is to be in (or to have been in) a state that is more than cognitive.¹⁵ And I also suggest that the fact that there are such concepts should not be so surprising since the grasp of color concepts also requires that one is in or has been in a state that is more than cognitive. The possession of the concept *rude* is like that.¹⁶

¹⁵ This raises the issue of the difference between a concept and the state of having a concept. I am skeptical about this alleged difference because it seems to me that a concept is not analogous to an object that can be passed around from person to person, much less is it an object that can exist apart from its possession by any person. But my argument in this paper does not depend upon the way this issue is settled. The question of whether moral judgments are motivating is not the question of whether a proposition or the concepts that are its constituents are motivating. The issue is whether the person making a moral judgment is in a motivational state. In what follows I will argue that that state is one that is intrinsically and inseparably both cognitive and affective.

¹⁶ A number of philosophers have pursued the idea that value concepts have parallels with color concepts. See, for example, John McDowell, “Values and Secondary Qualities,” in

Someone might propose instead that *rude* is simply a conjunctive concept, conjoining *having descriptive property R* with *causing a feeling of offense*. For example, Foot proposed *displaying a lack of respect* as her candidate for descriptive property R.¹⁷ So the proposal might be that something is rude just in case it displays descriptive property R and causes, or typically causes, a feeling of offense. But this cannot be right either because it does not explain why there is anything special about *this* descriptive property that causes offense and other descriptive properties that also cause offense but are not rude. If the relation between descriptive property R and feeling offended is merely causal, it is no different than the relation between other descriptive properties and the offense *they* cause. This observation forces us to look closely at the relation between feeling offended and those descriptive properties, and what we see, I propose, is that the feeling of offense in instances of rudeness is not merely caused by certain features of the situation, but is a matter of feeling offended *at* certain features, features that are construed a distinctive way when they are the object of the feeling of offense. These features of the situation are the intentional object of the feeling, not just the cause of it. Feeling nauseated, in contrast, has no intentional object.

Here, then, is the important difference between *rude* and purely descriptive concepts, on the one hand, and a concept like *nauseating* on the other. When I see something as rude, I feel offended *at* the offensive features of the situation, and those features cannot be fully described independently of their quality as intentional objects of the feeling of offense. That is, there is no purely descriptive feature of the object of offense that adequately captures that which is offensive. This is not to deny that the cognitive grasp of purely descriptive features of some situation, together with my emotional dispositions, *causes* my emotional response; in fact, I believe that that is probably the case. Nor is it to deny that offensiveness supervenes on those descriptive features, although I am not going to address the supervenience thesis here. It is simply to say that to see something as rude is not just to see it as having those descriptive properties, nor is it to see that something has caused me to feel offended, nor is it to see the two together.¹⁸ It is to see it as the intentional

Morality and Objectivity, edited by T. Honderich (London: Routledge & Kegan Paul, 1985), 110-129, and Barry Stroud's Tanner Lectures, 1989, vol. 10. I prefer not to tie my view of the ontological status of moral value to the outcome of the debate on secondary qualities. My point here does not depend upon there being a close analogy between *rude* and *red*. The similarity to which I am alluding is merely that in both cases the state of grasping the concept is more than cognitive.

¹⁷ *Displaying a lack of respect* is probably not a purely descriptive concept, in which case the proposal would be that there is something still more basic that the agent perceives in cases of rude behavior and which is purely descriptive.

¹⁸ In *Descartes' Error* (Putnam Pub, 1994), Antonio Damasio describes cases in which persons with a particular kind of brain damage have no trouble recognizing the descriptive features of a situation and understanding that they are thought to be rude without, as I put it, seeing it as rude. This work implies both that it is possible to perceive the descrip-

object of the feeling of offense. I do not just feel offended *by* the rude behavior, I feel offended *at* it. *Rude* is, then, a concept that combines descriptive and affective aspects in a way that cannot be pulled apart. The conclusion is that *rude* is a thick affective concept.

There are other thick affective concepts. I have already mentioned *pitiful*, *contemptible*, and *lie*. Most virtue and vice concepts are probably in this category as well—concepts such as *compassionate*, *kind*, *honest*, *merciful*, *dogmatic*, *close-minded*, and *brutal*. Besides *rude*, non-moral thick affective concepts probably include *petty* and *tacky*. *Honorable* is a thick concept that is ethical, but is moral only in a broad sense of moral.¹⁹ There are also concepts that are both descriptive and affective, but where the two components can be pulled apart, such as *fat*, which adds a separable affective component to the descriptive concept *overweight*, and so they are not thick in the sense I mean. Of course, many descriptive terms have some affective coloring, although in many cases it is not clearly part of the meaning of the word, for example, “meticulous”, “novel”, or “assertive”. There are also concepts that are connected with feeling, but that mean nothing more than causing a feeling of a certain kind. This was my point about *nauseating*, and it probably also applies to *exciting* and *frustrating*, but there are probably cases in between such as *funny* and *boring*.²⁰

I have said that in an emotional state we see the intentional object of the emotion a certain way that differs from the way the object is seen by a person who is not in the emotional state, and which cannot be designated in purely descriptive terms. I will reserve expressions of the form “I see *x* as *A*,” where *A* is a thick affective concept, for those cases in which the thing I see *as A* is the intentional object of an emotional state. So when I see something as *rude* I am in a distinctive emotional state. An emotion is therefore a unitary state that has both a cognitive aspect and an affective aspect that are necessarily connected. An emotion is a state of feeling a characteristic way about something seen *as* rude, as pitiful, as contemptible, and so on. So the intentional object of an emotion is seen as falling under a thick concept. When I feel the

tive features of a situation without seeing it as falling under a thick concept and that it is normal not to.

¹⁹ I doubt that there is any clear distinction between the moral and the ethical, although the latter is sometimes treated as broader. In the Western world the moral is associated with guilt and innocence, whereas the honorable is usually associated with cultures oriented around shame rather than guilt.

²⁰ We do say that something is funny or boring *to her*, but we also think that it is possible to make a mistake about what is funny or boring. Most philosophy professors judge that if a student finds Plato boring somebody is making a mistake, either the professor for not demonstrating how interesting Plato really is, or the student for not finding him so. This group of concepts suggests that not all feelings relate to their objects in the same way as the emotions I am considering in this paper, and yet there may be no reason to deny that they are emotions. I am not offering an account of the complete range of emotions here, but only those emotions that can be expressed in ethical judgments.

kind of offense characteristic of response to rude behavior, I see something *as* rude; when I feel pity, I see something *as* pitiful; when I feel contempt, I see something *as* contemptible; and so on. Probably not all feelings whose intentional objects are seen as falling under thick concepts have names. For example, I do not know the name of the feeling that involves seeing something as tacky, but I believe there is one. Similarly, probably not all thick concepts under which the intentional objects of emotions fall have names, for example, the intentional object of the feeling of indignation.²¹

The view I am proposing here differs in an important way from the view that emotions are responses to value. The emotion of pity is not a response to something seen as pitiful; the emotion of contempt is not a response to something seen as contemptible, and so on. Rather, seeing something as pitiful is a feature of the state of being in the emotion of pity; seeing someone as contemptible is a feature of the state of being in the emotion of contempt. Emotion is a kind of value perception that feels a characteristic way.²²

It is time to make a clarification about my use of the term “feeling”, which I have been using loosely to refer to any affective state or affective aspect of a psychic state. My position is that the feelings that are aspects of emotions have a different “feel” than those that are similar but are not components of emotion. For example, I have suggested that it is possible to feel offended when the offense is not directed at an intentional object. In that case the feeling of offense is a “mere” feeling. The feeling of offense that is directed at an object seen as rude is a distinctive way of feeling offended. I propose that no one can feel *that* way without seeing something as rude, and no one can see something *as* rude without feeling that way. For each emotion there is a thick affective concept correlated with it.

Emotions are potentially motivating because they combine affectivity and intentionality. Unlike non-intentional affective states such as sensations and moods, in an emotional state there is something specific in the world around us towards which affect is directed and which leads us to respond in ways characteristic of the emotion. So in a state of pity someone is seen as pitiful

²¹ There may be a word for it in Greek. Bernard Williams says that *nemesis* usually means “worthy of indignation” in *Shame and Necessity* (University of California Press, 1993, p. 80), yet it also seems to be the appropriate reaction to injustice or moral wrongdoing of other kinds.

²² There is also the theory that a value property just *is* the disposition of the bearer of the property to give rise to certain affective states in suitable observers. The color analogy allegedly supports this theory. (See, in particular, John McDowell, “Values and Secondary Qualities,” *op. cit.*). Variations of the response dependent view have become increasingly popular. I will not discuss the advantages of my theory over this class of theories here, although I believe it can be shown that the theory of this paper shares the advantages of response dependent theories without being vulnerable to its objections. For an insightful discussion of some problems with response dependent theories, see Cuneo, *op. cit.*

and given the opportunity to respond, the feeling of pity motivates us to take action to stop the suffering towards which the pity is directed. In a state of love someone is seen as lovable and the feeling of love motivates us to treat the beloved in characteristically loving ways. In a state of contempt someone is seen as contemptible and the feeling of contempt motivates us to express our contempt in some way. Of course, there are no doubt a number of other conditions necessary for actually acting on an emotion, but those conditions are necessary on any account of motivation, including the Humean desire account. The point of this paper is to argue that making a ground level moral judgment is in the right category to be motivating.²³

Of course, I am not suggesting that being in an emotional state is sufficient to act in the right way. What I am proposing is that emotions are potentially motivating because of their intrinsic features. Their affective aspect is “pushy”. And their intentional objects give them something to which the agent is motivated to respond. Emotions generate desires to change the world in ways characteristic of the given emotion; hence, they are act-dispositions.²⁴ For example, the emotion of compassion characteristically generates a desire to alleviate suffering; the emotion of anger characteristically generates a desire to “get even” or something similar. These are vague and fairly unformed desires. Action requires something more specific to direct the agent to act in one way rather than another way that satisfies the same desire. Since the motivating aspect of emotion is relatively unfocused, it needs to be shaped by experience, knowledge, and an understanding of the entire context. In short, it needs to be shaped by practical reason. But in this paper I am not focusing on what it takes to act in the right way, but on the psychology of making a moral judgment.

One way an emotion can be expressed is by asserting that the intentional object of the emotion falls under the thick concept. So in saying “That remark is rude,” or “He is contemptible,” or “She is pitiful,” I can be expressing the emotions of offense, contempt, or pity, respectively. Not all uses of these sentences express emotions, and we will get to that presently, but I propose that the ground level use of these sentences, the one we use in learning thick concepts, and the one we continue to use as the point of reference for all other uses, is the expression of emotion. Since my judgment expresses an emotion, it expresses an intrinsically motivating state. And since the

²³ Emotions directed towards a character in a novel are not motivating because there is nobody to whom one can respond in action. But I believe the psychology is the same as in cases of emotions directed towards actual persons. In both cases, if a practicable opportunity to respond through action presented itself, the agent would be motivated to act. Perhaps we can permit ourselves a greater latitude in emotional response to characters in novels because we know that we will never be called upon to act upon our emotions.

²⁴ For an interesting account of the emotions as act-tendencies, see Nico Frijda, *The Emotions* (Cambridge University Press, 1986), esp. section 2.8.

judgment also asserts that some person, object, or state of affairs falls under the thick concept that applies to the intentional object of that emotion, it is propositional in form, is about the intentional object of the emotion, and I am in the cognitive state of taking the intentional object to fall under the thick concept. If it does fall under that concept, the judgment is true; if not, it is false. Hence, judgments like “That is rude” or “She is pitiful” are both cognitive and intrinsically motivating when these judgments are expressions of emotion.²⁵

If the foregoing account of emotion is possibly true, then it is possible that there are psychic states that are both cognitive and affective and which are such that the cognitive aspect of the state cannot exist apart from the affective aspect, and the affective cannot exist apart from the cognitive. It follows that the Humean doctrine that cognitive and affective states are necessarily “distinct existences” is false, but several provisos are in order. First, the Humean doctrine is usually posed in terms of an essential distinction between belief states and states of desire. But it might be true that beliefs and desires are necessarily distinct even though it is false that every cognitive state is necessarily distinct from every affective state. On the account I have given, an emotion is not a combination of a belief and a desire. Second, the Humean position is usually posed as a claim that any given belief (or cognitive) *state* can exist apart from any given desiderative (or affective) *state*. The problem is therefore posed as one of comparing an independently identifiable state that is wholly cognitive with an independently identifiable state that is wholly affective. This way of presenting the problem arguably rules out of consideration in advance possible states that are not independently characterized as wholly cognitive or wholly affective.²⁶ The Humean doctrine might therefore be true of those states Humeans consider. My point is that there are states they have not considered, in particular, emotions. Third, since the Humean doctrine denies the *possibility* of essentially united cognitive and affective states, there remains the issue of what kind of possibility is denied and how its presence or absence can be determined. The usual procedure is introspection. The reader is invited to imagine herself in the one state without the other and for each cognitive and feeling state to which the Humean points, he says he can always imagine them existing apart. I have proposed states in this paper that I find fail the test, but I doubt that the test is conclusive anyway. It is unlikely that many of us have a clear idea of what constitutes being in a state

²⁵ Some theorists on the emotions analyze an emotion as a type of judgment, e.g., Martha Nussbaum in *Upheavals of Thought* (Cambridge University Press, 2001). I find it implausible that the cognitive component of an emotion is as structurally complex as a judgment. I have argued that an emotion has a cognitive component and the emotion can be expressed in a judgment, but the conditions for making a judgment are more complex than the conditions for being in the emotion that is expressed.

²⁶ See note 10.

of an emotion such as pity, or even how to individuate our emotional states in advance of a plausible theory of emotion. If we had a clear idea in advance, we would not be faced with so many wildly differing accounts of emotion. Perhaps states of emotion as I have characterized them here are not possible, but if so, that has not been demonstrated by anything Hume says.

C. Summary

The general picture of ground level moral judgment I have been describing goes as follows: A situation has certain descriptive features D. My perceptual and cognitive awareness of D, together with my emotional dispositions, causes me to be a certain emotional state E. E is an affective state whose intentional object is the D situation seen as falling under a thick concept A. I cannot see situation D *as* A without being in emotion E. I can express emotion E by making the judgment, "That is A." For example, a situation has the descriptive feature of being a remark that expresses disdain for me and is uncalled for. My awareness of this feature, together with my emotional dispositions, leads me to be in a distinctive emotional state consisting of feeling offended at the rudeness of the remark. In this emotional state I see the remark as rude. I may express my emotion by simply saying "That is rude." This judgment expresses an intrinsically motivating state since it expresses an emotion; it is also propositional in form with a truth value, and I am asserting that proposition when I say "That is rude." To take a moral example, a situation has the descriptive feature of being one in which a person with whom I do not identify is suffering. (We might also need to add that I find the sufferer beneath me in status.²⁷) My awareness of this descriptive feature, together with my emotional dispositions, leads me to be in the distinctive emotional state of feeling pity for the sufferer whom I see as pitiful. I may express my emotion by simply saying, "She is pitiful." This judgment expresses an intrinsically motivating state since it expresses the emotion of pity; it is propositional in form, with a truth value, and I am asserting that proposition when I say "She is pitiful."

It follows that one way of expressing an emotion is to assert a certain kind of proposition. So when I say "That is rude" or "That is pitiful" or a host of other statements of that form using thick concepts I may be making a judgment that is cognitive and is intrinsically motivating. Further, the judg-

²⁷ It is interesting to compare the differences in the descriptive features of situations that elicit compassion and those that elicit pity. In both cases the situation involves a person who is suffering or harmed or whose dignity has been damaged in some way, but there is a difference in the relation between the agent and the sufferer in the two cases. Pity seems to arise when the agent does not identify with the suffering person or perceives him as lower in status than herself. Compassion arises in situations in which the agent does identify with the sufferer or considers him an equal. This is probably the reason why compassion has always been considered a morally good state, whereas pity is morally ambiguous.

ment is not precluded from being primarily about a reality outside my mind, although I will not address that issue here.²⁸

III. *The thinning of moral judgment*

I cannot see something as rude without feeling offended in the characteristic way that goes with rudeness, but I can see or judge *that* something is rude without feeling offended. I can do that in a way that roughly parallels the way in which I judge colors. I cannot see something as red without seeing red, i.e., without having a sensation of red, but I can see *that* something is red without a red sensation. I can do that if I see signs of its redness. Perhaps I am looking at it in the dark and I see that it looks the way red looks in the dark. Or my sensory clues may be more indirect, as when I see a traffic sign that has the characteristic shape of a red stop sign. In each case I may be able to judge that the object is red without having the sensation of red, that is, without, as I put it, "seeing it as red." Similarly, I cannot see something as pitiful without feeling pity, but I can see *that* it is pitiful without feeling pity. When I judge that something is pitiful without feeling pity, as when I judge that something is red without having a red sensation, my judgment that the property applies is indirect. When I judge that it is pitiful while feeling pity, as when I judge that it is red while having a red sensation, my judgment that the property applies is direct. In an indirect judgment it is likely that I must still be able to imagine seeing something as pitiful or as red, but even imagining requires a faint copy of the pity feeling and the red sensation in my mind. (Here I think Hume probably had it right). We need to be able to do this to even understand what is said when someone judges that something is red or pitiful. Without a memory trace of the sensation or feeling we cannot understand the judgment at all.

The ability to judge that something is pitiful without feeling pity, like the ability to judge that something is red without having a red sensation, is a critically important ability. One of the most important uses of concepts, in fact, the main reason we need concepts at all, is to have a way to talk about the world when we are not in the state that originally gave rise to the concept, and when the objects to which the concepts apply are not here and now. We need to talk about horses and tables and sunshine and red, and the

²⁸ It is not precisely correct that my position requires that the same state can be both cognitive and motivating. It requires that there are states that are both representational and motivating, but it is not essential to my view that every representational state is cognitive. According to the standard psychology perceptual states, like cognitive states, are representational but non-motivating. Conative states, like affective states, are non-representational but motivating. The differences between perceptual and cognitive states and between conative and affective states are not important for my purpose here. My purpose is to argue that the nature of moral judgment is a problem for any psychology that asserts that representational and motivating states are essentially distinct and never more than causally connected.

pitiful, the rude, and the contemptible. I have said that we can use the concept of pity to apply to a person before us who is pitiful even when we do not feel pity. More frequently we want to talk about a pitiful person we have only heard about, or we may think of one hypothetically, as we are doing now, or we may read about one in a novel, or the pitiful person may have existed in the past. The ability to make a judgment like "He is pitiful" without feeling pity is necessary for much of our discourse even though I have proposed that we learn the concepts from ground level, here and now situations. There is therefore nothing aberrant about such indirect judgments. They are the first step in what I call the "thinning" of a moral judgment, and I call them Level 2 judgments.

It follows that the words "That is pitiful" can be used either to express a ground level judgment that is an expression of the emotion of pity, or it can be used to merely express the proposition *that* the object is pitiful. Similarly, the words "That is red" can be used either to express a judgment that includes the sensation of red, or it can be used merely to express the proposition that an object is red. So the same words can be used to express either a Level 1 or a Level 2 judgment.

There is also a difference between the statements "That is pitiful" and "I pity that." Either one can be used to express the emotion of pity, but there is a difference in intent and ontological commitment in the two cases. When I say "That is pitiful" I am willing to assert that the intentional object of my pity has the property of being pitiful, whereas if I only say "I pity that", I am calling attention to myself and my emotion, and this may indicate that I am unwilling to say that I have construed the intentional object of my emotion correctly. It is even possible to say "I pity her, but she is not pitiful." We cannot make statements like that without higher order reflectiveness on our own emotional states and their intentional objects. A statement like "I pity her, but she is not pitiful" is like "I have a red sensation, but that object is not red." We understand how it is possible to make such judgments because we also understand how they are descendants of the ground level uses of the concepts *red* and *pity*.

I call a judgment that something is pitiful without feeling pity or the judgment that something is rude without feeling offended a level 2 judgment since it is one level removed from ground level judgments. Level 2 judgments also utilize thick concepts, and I have said that the agent cannot understand these concepts without having had the feelings that go with them when they are part of level 1 judgments, and I have conjectured that she probably must have a memory of those feelings, a memory which gives her a faint copy of those feelings, even though she does not have the feeling at the time she makes a level 2 judgment. Level 2 judgments are necessary because it is important to engage in moral discourse about situations that are not here and

now, and we want to be able to examine the logical consequences of our judgments, which is often more difficult to do with Level 1 judgments. But these judgments are motivating only to the degree that a faint copy of the motivating feeling remains. This means that we lose something in the thinning process—emotional practice. Level 2 judgments might be compared to judgments about color or musical sound or the taste of wine carried on without seeing color or hearing the music or tasting the wine. We gain something, but we lose something.

When we get to what I call Level 3 judgments, thick concepts disappear entirely. These judgments use what Bernard Williams calls “thin evaluative concepts,” concepts such as *right*, *wrong*, *duty*, *should*, *ought*, and *good*. I propose that these judgments are two steps removed from ground level judgments. The process of getting to a level 3 judgment by thinning a ground level judgment could go roughly as follows: First, I see someone as pitiful and make the judgment “She is pitiful,” expressing the emotion of pity. At Level 2 I can judge that she is pitiful, but do not feel pity. Once I am able to judge that she is pitiful without the feeling of pity I can make a similar judgment of other people in a way that is not “here and now.” So at Level 2 I can judge that persons are pitiful whom I have never met, or persons I merely imagine.²⁹

At Level 3 I make a judgment like “I should do something to help her,” or “It is right for me to help her.” Level 3 judgments shift the focus from the person or state of affairs to which I am responding to my response. Level 3 judgments, like Level 2 judgments, serve both theoretical and practical purposes. First, we have a theoretical need to understand the connections among the different thick moral properties objects and situations can have. It is likely that the pitiful, the contemptible, the honorable, the rude, the dishonest, etc. have little in common. But what they do have in common is that they call for responses of various kinds, and the responses can be categorized as those that are called for, those that are uncalled for, and those that are neither the one nor the other. I propose that thin concepts categorize situations according to the categories of our response. This leads to the second need for thin concepts, a practical one. We need linking concepts between thick concepts like *pitiful* and practical response in action. Most thin concepts apply to action responses (*should*, *ought*, *right*, *duty*); they do not even purport to be properties of the situation to which we are responding. An

²⁹ I suggested in note 23 that the psychology of judgments about characters in novels is similar to the psychology of judgments about actual persons. Such judgments can be ground level. But assuming that imagination is generally thinner than perception, level 2 judgments about fictional characters would also be common. The fact that the thinning process permits degrees of variation between levels 1 and 2 can explain the fact that there seem to be cases in between feeling a full emotion and merely judging that the intentional object of the emotion has the requisite thick property.

exception is the pair of concepts *good* and *bad*, but curiously, these are almost never the ones used in debates about whether moral judgments are motivating. I suspect that the logic of good and bad differs from the logic of the other thin concepts, but I will not attempt a discussion of that issue here.

Typically, discussions of the issue of whether moral judgments are motivating take only level 3 judgments as examples. It is no wonder, then, that many theorists feel forced to say that moral judgments are not motivating since level 3 judgments are the farthest removed from the class of judgments that are most clearly motivating. I have argued that level 3 judgments are not basic, so whereas they are not themselves motivating, they are derivative from a class of judgments that are motivating. Moral discourse would be impossible without judgments of at least level 2, which also can be made in the absence of emotion. But even level 3 judgments can be motivating in some persons who have acquired an emotional response to concepts like *duty*, *right*, and *should*, but one can understand these concepts perfectly well without ever being in an emotional state when using them, and I doubt that they are initially acquired through the experience of emotion, as I have suggested happens with the ground level thick concepts.

The thinning of moral judgment makes possible the problem cases of moral motivation: moral strength and weakness, and moral apathy. In situations of moral struggle, a person making a Level 2 or Level 3 judgment is not sufficiently motivated to act on it because the emotion that provides the motivational force is absent. But if I am right about the thinning process that produces these judgments, it would be common for some degree of motivational force to remain because the relation between a higher level judgment and the emotion present in ground level judgments would be part of our past moral experience. Since the original judgments are expressions of emotion, emotion is part of the point of making moral judgments, even those judgments that are derivative. This would explain why we find moral strength and weakness somewhat surprising even while we acknowledge how common it is. Not only does a flicker of emotion commonly remain in Level 2 and Level 3 judgments, but the judgment does not make sense without a tie to a more basic judgment in which emotion is present. That is, even Level 3 judgments are linked to emotion in that they use concepts of response to situations that fall under concepts whose ground level use involves emotion.

The morally strong person does not have sufficient motivation arising from the judgment itself to act on the judgment, and must rely on additional sources of motivation that are also part of moral training. In fact, training in what used to be called "will power" has that function. The morally weak person is in the same situation except that she lacks sufficient additional sources of motivation to act on the judgment. The weak motivational force of a Level 2 or 3 moral judgment can sometimes be reduced to none at all,

which is what happens in the morally apathetic person. The causes of apathy are typically other emotional or physical states such as depression or weariness that are strong enough to completely blot out the motivational force of some moral judgment. Since the apathetic person lacks all emotional connection to the judgment she makes, she even lacks the motive to call on the additional sources of motivation commonly used by the morally weak and strong.³⁰

In short, moral motivation is confusing because we cannot account for the connection between our motivating emotions and our judgments by one single mechanism. Nonetheless, our psychology is economical. There is an important class of cases in which the judgment itself is motivating because these judgments are expressions of emotion. Those moral judgments that are not expressions of emotion are the descendants to those that are, and while they require the development of a backup mechanism of moral motivation, they are not wholly devoid of the motivating force of those judgments from which they derive.

IV. Conclusion

The account of moral judgment I have given here shows how the advantages of moral cognitivism and motivational judgment internalism can be combined in a single type of moral judgment that enjoys primacy conceptually and in the learning of moral behavior, and that is not relative to the judgments of persons who are virtuous or even especially rational. Derivative forms of moral judgment allow for the common experience of a dissociation between the cognitive aspect of moral judgment and motivation, thereby making moral strength, weakness, and apathy possible. This account also shows why thick concepts are so important in moral learning. An interesting consequence is that it may support Bernard Williams's well-known claim that moral reflection takes away moral knowledge. That is because Level 2 and especially Level 3 judgments are the result of reflectiveness, but the reflectiveness needed to generate these judgments detaches the judge from the emotion out of which the ground level judgments were made. If ground level

³⁰ This account may also explain the psychology of amorality and some forms of wickedness. An amoral person never learned level 1 concepts, but may know the appropriate use of level 3 concepts (although one may doubt that he really understands them). The same point applies to the form of wickedness in which a person sees the bad as good. There is another and more interesting phenomenon discussed by Michael Stocker, which he calls desiring the bad (*ibid.*), indeed, desiring the bad under the aspect of bad. The psychology of this is clearly complex, but my conjecture is that it involves seeing "the bad" not merely as bad, but as falling under distinctive thick concepts such as the shocking or the disgusting, where the attraction is not merely to the bad as such. In other forms of desiring the bad, an agent may react to certain negative features of the good, such as those characterized by Nietzsche, so that he is not rejecting the good as such, but the good as included in certain unappealing thick concepts. For example, the concept of a person who turns the other cheek suggests to many people an unacceptable wimpiness.

judgments give us an important kind of moral knowledge, akin to the knowledge we get from perceptual judgments like “That is red,” then the thinning of these judgments would be a serious loss, even though I have argued it is a necessary loss, and it does not come without an important gain.³¹

³¹ An early version of this paper was the Keynote Address at the Central States Philosophical Association meeting in Des Moines, Iowa, Oct. 1998. Other versions have been presented at the University of Arkansas, Little Rock, SUNY Buffalo, Valparaiso University, the University of Notre Dame, the University of Genoa, Tulane University, and the University of Glasgow.