

Chapter 8

Three Remaining Questions

At the end of Chapter 1 I posed seven questions that I intended to answer in this book. Three of these questions have yet to be answered. They are the following.

1. If MM is true is that because motives are sometimes wrong-making but never obligation-making? That is, is PR true as well?
2. If motives ever are obligation-making do we then have cases of hypothetical moral obligations?
3. Are hypothetical moral obligations self-contradictory or otherwise objectionable?

Before we find answers to them I propose to examine Thomas Hurka's intrinsic consequentialism. I will give three reasons for concluding that it is not correct to assert, as he does, that motives are sometimes intrinsically valuable. The most plausible theory with regard to the deontic relevance of motives is therefore extrinsic consequentialism. We will use that theory to answer the three questions. It leads us to say that PR is false: there are obligation-making motives. There is nothing conceptually amiss in this assertion. On the other hand, as we will see, it is doubtful whether in practice this idea plays much of a role in our everyday moral thinking.

I

Intrinsic Consequentialism

We can begin with the evaluation of intrinsic consequentialism. Intrinsic consequentialism affirms that motives can be deontically relevant in virtue of their intrinsic value, as well as in virtue of their effects.

Hurka's claim about the intrinsic value of certain attitudes is part of what he calls a recursive theory of value. I will summarize it for our purposes. The 'base level' of the theory incorporates claims about the most basic states of affairs that have intrinsic value. Hurka believes that pleasure, knowledge and achievement are the states of affairs at the base level that are intrinsically good; pain, false belief and failure are the states of affairs at the base level that are intrinsically bad. We will not examine these claims. Our interest is in the further claims he makes about the value of psychological attitudes directed at these states of affairs. Described abstractly these claims are as follows: a love of something intrinsically good for itself is intrinsically good; a love of something intrinsically bad for itself is intrinsically bad; a hatred of something intrinsically bad for itself is intrinsically good; and a hatred of something intrinsically good for itself is intrinsically bad. (11-19)¹ These claims pertain to the second level of his theory. They could in fact be employed in a different recursive theory that made different claims about base-level goods. Hurka also makes claims about third- and higher level intrinsic goods. At the third level of the recursive theory we evaluate attitudes towards the second-level attitudes. He claims (roughly), for example, that a love of a love of something intrinsically good is intrinsically good.² In theory the hierarchy of psychological states

can be extended upward, but it is doubtful that human beings have attitudes above the third level. I will generally focus the discussion on the second level of his theory, but our conclusions will apply to higher levels.

It is important to realize the breadth of Hurka's claims about the attitudes that have intrinsic value. They go far beyond the topic of motives. In fact, the main topic of Hurka's book is morally significant character traits, that is, virtues and vices. However, he thinks of these as manifesting themselves in various psychological attitudes, including desires having the sort of structure just noted. And he seems to be committed to the idea that any episodic desire, whether 'in character' or not, having the relevant structure is intrinsically valuable. (cp. 48-49) Hurka considers the value of a number of other psychological states like wishes that are neither motives nor dispositions to act, but that have the relevant structure. He claims that a wish that something intrinsically good happen 'for its own sake' has intrinsic value. Virtuous people have such wishes.

Hurka does not give us a general characterization of the relevant attitudes. Consider the general 'positive' attitude that is sometimes intrinsically good, and sometimes intrinsically bad. Hurka calls it 'loving'. Some desires and motives will fall under this rubric, as will other propositional attitudes like wishing and hoping for. We saw that desires can be expressed by means of propositions; hopes and wishes can be as well. But both positive and negative attitudes can be taken toward one proposition. I can want it to be true that Smith is President, or fear that it is true that Smith is President. To explain fully the nature of the attitudes that have intrinsic value we would need to say what all the 'favoring' attitudes have in common, and what all the 'negative' ones have in

common. But let us proceed without this. We can begin by examining Hurka's basic claims about the second level of his theory in their full abstractness.

1. **The Abstract Argument and Its Flaw.** He says that the central idea in his recursive theory is captured by "a simple idea", namely, "it is intrinsically good to be oriented positively toward good and negatively toward evil, and intrinsically evil to be oriented negatively toward good and positively toward evil." (17) Elsewhere he writes that "the core idea" of his theory is that "the value of an attitude...depends on the value of its object." (24) Finally, he gives a summary of what his theory says "makes good attitudes good": "if an attitude is one of loving, it is good because its object is good; if it is one of hating, it is good because its object is evil." (189) He says these contentions are "immediately persuasive". (190)

There are, in fact, two meanings that the word 'object' has here. One way of understanding what Hurka finds persuasive is the following. We first find 'base level' goods like knowledge and pleasure. These goods, or instances of them, can be the objects of certain favoring attitudes. For example, suppose S is pleased that she knows about the theory of evolution. Here 'pleased' is a 'factive' attitude that implies that S has the relevant knowledge. But there is another way to interpret what the 'object' of an attitude is: here the object is defined by the agent's beliefs. This sort of 'object' may not exist, or even be capable of existing. The object of S's desire can be finding the Fountain of Youth. With this meaning in mind Hurka speaks of the "intentional and therefore internal relation" that an attitude has toward its 'object'. (32) It is important to realize that it is this second meaning that Hurka has in mind in the second level of his theory. He makes this clear when he says that a false belief about an object's existence does not by itself

affect the value of a favoring attitude towards it. If a scientist like Aristotle desires knowledge as an end, the fact that his methods of investigation will only yield erroneous conclusions does not detract from the intrinsic value of his desire. (34; Cp.178; 183)

Hurka's basic argument thus actually involves an equivocation. It admittedly has some force when it utilizes the first, or factive, understanding of 'object'. If we have picked out an existing piece of knowledge it is plausible to say that a love of it is valuable, at least in some way. If S is pleased that she knows the theory of evolution her attitude is in a clear sense a response to an existing intrinsic good. Similarly, if we assume that a desire will lead to the production of some knowledge then a desire for it as an end is valuable, in some way. (I will say more about this shortly.) But his conclusion about the value of attitudes employs the intentional meaning of 'object'. Sometimes what an agent believes is intrinsically valuable really is so. But sometimes it is not.³ In the latter case it is not clear why having a favorable attitude towards this sort of 'object' is intrinsically valuable. Suppose that S is pleased that she 'knows' astrology for its own sake. (She is not pleased with this knowledge because of what it can do for her.) Hurka writes, "The value of an attitude...depends on the value of its object." If we take 'object' in the factive sense, this statement actually seems to argue against saying that S's attitude in this case has intrinsic value, rather than in favor. Astrology is not knowledge, so S does not have an attitude that responds to something having intrinsic value. In the assertions constituting the second level of his theory Hurka takes 'object' to have the 'intentional' meaning. But the base level of his theory does not assert that all of the objects that agents believe are intrinsically valuable really are intrinsically valuable.

Another way to put the problem is this. If S knows about the world, knows what has intrinsic value, and has the appropriate attitudes to them, then these attitudes might well be thought to inherit the intrinsic value of their (factive) objects. But if S's attitudes are based on false beliefs it is not obvious that they have intrinsic value, even if they reflect her convictions about what has intrinsic value. By utilizing the intentional meaning of 'object' with regard to second level attitudes Hurka loses the ability to claim that they have intrinsic value simply in virtue of being 'about' the states of affairs at the first level that have intrinsic value.

Now, Hurka can support his views about the intrinsic value of some attitudes in some other way. He can appeal, for example, to 'reflective equilibrium'. But I do not think he is correct in saying that there is an "immediately persuasive" argument from the intrinsic value of base level goods to the intrinsic value of all of the attitudes that he claims have it.

2. The Priority of Producing Good. If there is no strong argument for the claim that a large class of attitudes has intrinsic value we can ask whether there are any reasons to claim that many of them, anyway, have extrinsic value. There are. This is obvious in the case of the desires that are motives, since they lead people to do things that produce states of affairs that are intrinsically valuable. And we saw in chapter 3 that motives have other effects that can make them relevant deontically. People care about why other agents are performing the actions that they do, and this can bring about significant psychological effects. But I want to continue looking at Hurka's claims in their full generality and to think more about the whole range of psychological attitudes that he is interested in. He does not merely present an abstract argument for the intrinsic value of certain attitudes.

He later discusses a number of specific attitudes. The critical issue in considering these attitudes is not whether they can be valuable: obviously they can be. The issue is the kind of value they have: is it intrinsic or extrinsic?⁴ To answer this question we need to examine the basis of the value judgments that Hurka makes. Many of these judgments actually turn out to be most plausibly construed as involving extrinsic value.

We can grant Hurka that there are certain 'base level' values. We do not even need to decide what they are, so long as we understand them to be 'agent-neutral'. That is, certain objects, properties or states of affairs are characterized as intrinsically good, whatever their relation is to a given agent. So, if knowledge is intrinsically good then anyone's knowledge is intrinsically good. And if some kinds of knowledge are better than others, then anyone's having that kind of knowledge is better.

I will now mention six points that Hurka makes. They represent some of the remarkable elaborations of the recursive theory of value that he makes. Some of these points can be seen as qualifications to the agent neutrality of the base level goods. **1.** The value of an attitude depends in part on its object's relation to actuality. (118f; 125; 168-9) This means, for example, that the compassion directed at an imaginary character like King Lear is less valuable than that directed at a real human being.⁵ Attitudes directed at remote possibilities have less value or badness than when those same attitudes are directed at real possibilities. **2.** The attitudes themselves have different values, even when their objects have the same proximity to the actual world. Desires to realize goods have more value than mere wishes to bring those same goods into existence; desires to bring about bad states of affairs are worse than mere wishes for them to happen. (123) **3.** Pain that an agent feels at harm she herself has caused is better than pain that she feels for a

harm of the same magnitude caused by someone else. Similarly, it is worse to enjoy a fantasy of torturing someone than it is to enjoy to the same degree a fantasy of witnessing someone else torture a person. (204-5) This is one of the modifications that Hurka makes to the agent neutrality of his theory. Others include the following. **4.** Desiring a certain amount of my own pleasure as an end just as much as I desire the same amount of others' pleasure as an end, and desiring that amount of pleasure for others more than I desire that amount of my own, are equally valuable, under certain conditions. That is to say, under these conditions an agent's having an altruistic bias towards the pleasure of others is just as valuable as her desiring her own pleasure equally with others'. (208-212) **5.** Being pained by one's own vice—a third-order attitude—is more valuable than being pained by someone else's vice of the same magnitude. (206) **6.** In contrast, being pleased by one's own virtue is not more valuable than being pleased by someone else's virtue of the same magnitude. Hurka plausibly describes a disproportionate pleasure in one's own virtue to be “an objectionable form of moral pride”. (206)

I consider these claims to be morally sensitive, and largely in line with the evaluations of the relevant attitudes that are made by reflective people. But we need to understand their basis. Consider Hurka's second point about the value of wishes as opposed to desires. To wish for a certain good is to have a favoring attitude towards it, but so is a desire to produce it. Why, then, is a desire to produce it more valuable? Hurka's approach suggests that reflection on the relationship between a wish for a good and that good, as compared to a desire for it and that good enables us to answer this question. I think, though, that the only plausible answer is that these judgments reflect a focus on producing good results. A desire to produce a certain good is more likely to

produce that good than a wish for it is. If this is the explanation for our thinking, then the evaluations that we are making concern the extrinsic value of these two attitudes, not their intrinsic value.

A similar idea underlies Hurka's first point. Our normative thought does not deny that wishes and hopes, as well as attitudes directed at fictional beings, have value. And we are indeed inclined to agree that, say, compassion directed at a remotely possible being is less valuable than compassion directed at an actual being. The issue is why we think this. And, again, the obvious explanation is that we think that compassion directed at actual beings, or at beings whose misfortunes are likely to occur, will tend to lead agents toward the real amelioration of their suffering, or some other being's suffering. To conceive of the evaluations in this way is again to see our thinking as involving judgments of extrinsic value.

Now consider Hurka's fifth and six points. These concern certain attitudes towards one's own virtue and vice, as compared to attitudes towards the virtues and vices of others. If we were looking at these evaluations abstractly it seems that we would say that a pleasure in my own middling honesty is just as valuable as the same amount of my pleasure in your middling honesty. Similarly, we would say that being somewhat pained by my own moderate callousness is just as valuable as my being somewhat pained by your moderate callousness. Hurka correctly says that we do not tend to equate the value of the relevant attitudes. The most plausible explanation of these judgments, I think, appeals to certain deep-seated, but well-known, propensities of human nature. We generally are rather too impressed with our own virtue, and therefore insufficiently motivated to improve ourselves. So our judgment about the value of an agent's pleasure

in her own virtue is somewhat discounted. We likewise think that humans tend to ignore their own vices as compared to the vices of others. Furthermore, an agent's pain at her own vice is more directly tied to reforming action than is her pain at another person's vice, since that second pain may not even be known by the other agent. If these thoughts are underlying our evaluations we are again making judgments about the likely effects of the relevant attitudes, given human nature.

I submit that in fact all six of the modifications that Hurka makes to his recursive theory of value can plausibly be seen as reflecting a concern with producing good results. The departures from strict agent-neutrality, for instance, seem to respond to the influence of background beliefs about human nature and its propensities. They make morally enlightened corrections for the prevailing motivational winds, as it were.⁶ We thus can see why certain attitudes are more or less valuable than we would have supposed if we considered them merely as abstract propositional structures.

I do not assert that no psychological attitudes are intrinsically valuable. It is possible, for example, that a factive pleasure in knowledge that one actually possesses is intrinsically valuable. We do not need to resolve this question. It is sufficient for our purposes if we can see that many of the attitudes that Hurka regards as intrinsically valuable are more plausibly regarded as only extrinsically valuable.

3. The Desire to Produce New Knowledge as an End. I conclude the discussion by returning to our topic, the psychological attitude of desire. If we are discussing the value of a kind of desire we are not quite focused on motives, since all motives are desires, but not all desires are motives. This is because an agent may have a desire in either a dispositional or occurrent sense, but not act on it. It is conceivable that Hurka

believes that those desires that are acted on have more intrinsic value than desires that are only had occurrently or dispositionally. (This might be regarded as an extension of the second of the six points above.) But I propose to examine one of the general claims he makes about a certain desire, whether it is acted on or not. This is the claim that a desire to produce new knowledge as an end is intrinsically valuable. This can serve as a test case for his views about the intrinsic value of certain motives.

Hurka conceives of what he is evaluating as a desire and its intentional object. If we focus our minds on them properly, he contends, we will affirm that the desire to produce new knowledge as an end has intrinsic value even if an agent who acts on it never succeeds in producing new knowledge. Furthermore, Hurka seems to think that his view is the only one that can explain why we esteem the efforts of a scientist like Aristotle, in spite of the fact that he succeeded in producing very little new knowledge. (34; 183)

I disagree with Hurka's explanation of why such a desire is valuable. We can think of the evaluation of this desire in a different way. When we focus our minds on another question we can see that the desire to produce new knowledge as an end is indeed valuable, but only extrinsically valuable. In other words, we can affirm the following statement: in general, the desire to produce new knowledge as an end is extrinsically valuable, that is, valuable only because of its likely effects.

Let us make what I believe is a false assumption, and say that Aristotle's desire to advance human knowledge of biology did not in fact produce any new knowledge. Is it still possible to say that his desire to produce this knowledge as an end was extrinsically valuable? To speak of the value of this desire 'in general' is to adopt a long-term social

perspective.⁷ Different relative evaluations of this desire from this perspective will seek to bring about different degrees of prevalence of this desire in a group of people, different degrees of strength for it on average, and so on. Furthermore, we need to keep in mind that these desires will often provide imperfect human agents with the ultimate normative premises of their practical reasoning. These agents will not know before they engage in research whether they will succeed in producing new knowledge. Given this crucial fact, an extrinsic consequentialist, like Hurka, will have to focus her evaluation on ‘the desire to produce new knowledge as an end’ in the intentional sense. That is, she will be evaluating a desire that will move people to do things that they believe will produce new knowledge. Let us suppose that it is clear from this social perspective that many people who act on this desire will fail to produce new knowledge. (This is probably too pessimistic an assumption, but weakening it does not affect the point I am making.) It may still be true that the net result of the efforts of all these researchers is that much new knowledge is produced. This establishes that the desire is indeed extrinsically valuable, in spite of the low success rate of researchers. Hurka’s error in thinking about the value of this kind of desire stems from his focusing his thinking too narrowly on the results of what he takes to be one failed researcher’s efforts.

We need to combine some of the conclusions we have reached. Hurka does not have a convincing simple argument to show that a wide range of second-order psychological attitudes are intrinsically valuable. In particular, he does not have such an argument to show that the desire to produce new knowledge as an end—taking this in the ‘intentional’ interpretation—is intrinsically valuable. Furthermore, an extrinsic consequentialist can explain what Hurka seems to think she cannot, namely, how the

desire of a researcher who fails to produce new knowledge can be seen as extrinsically valuable. These conclusions support my assertion that this desire is only extrinsically valuable.

There is some further support that can be given to this assertion. Curiously enough, Hurka's view gets no support from what we might call the internal perspective on motivation. This was discussed in Chapter 2. An agent who desires to bring about a certain state of affairs as an end need not regard this state of affairs as intrinsically valuable. If S desires to produce new knowledge as an end she need not believe that this knowledge is intrinsically valuable; she may regard it as an all-purpose means, and have no particular use for it in mind for now. But it also possible that S desires to produce this knowledge as an end because she believes it is intrinsically valuable. She practically affirms that certain new knowledge is desirable as an end because she believes that it is intrinsically valuable. In so acting S is not affirming that her desire is intrinsically valuable; she is affirming that the knowledge she seeks to produce is intrinsically valuable. In saying these things we are formulating part of S's practical syllogism, and also articulating S's beliefs about knowledge and its value. Some of these beliefs may be mistaken. But the fact remains that if S herself accepts Hurka's claim about the intrinsic value of knowledge, and seeks (however fecklessly) to produce some, she will not be guiding her activity by an evaluation of this desire itself.

I take all of this to support the more far-reaching conclusion that extrinsic consequentialism is more plausible than intrinsic consequentialism in its treatment of the value of motives.

If this is correct then we need to recognize a further point: the treatment of concealed motives that intrinsic consequentialism allowed for is not available. In Chapter 3 I suggested that one advantage that intrinsic consequentialism seems to have over extrinsic consequentialism is that it can recognize that even perfectly concealed motives might be relevant deontically in virtue of their intrinsic goodness or badness. But motives do not have intrinsic goodness or badness. This means that motives that are concealed from others are relevant deontically only because of their effects. These are not limited, of course, to the direct effects noted in Chapter 3. Acting from a given motive can have effects on an agent's character, and these can have significant further effects on others, as well as the agent herself. For some agents and some motives these indirect effects may be significant enough to be relevant deontically. What we cannot grant is that motives have any deontic relevance aside from these various effects.

II

Taking Stock

Before we consider two last problems, it will be useful to take stock of what we have established thus far. We set out to see which of the major moral theories could give an acceptable account of the deontic relevance of motives. I argued that the two central ideas in Kantianism cannot do this, and that the version of virtue ethics defended by Michael Slote cannot do so either. Consequentialism can give an acceptable account. We have just seen that Hurka's claim that some motives are intrinsically valuable is not

convincing. Extrinsic consequentialism seems to be the only moral theory left standing. I want to explain in what sense this is true.

The main points that have been established thus far are the following. Motives can be deontically relevant because of their effects. The two most important types of effects I pointed to are effects on the feelings of other agents, and effects on how an action is performed. These effects are most likely to be deontically relevant in cases of personal interaction. They are unlikely to be relevant, for example, in cases of large-scale philanthropic activity that saves peoples' lives. The way in which motives are deontically relevant entails that there are in theory no strongly wrong-making or obligation-making motives. Motives sometimes have a large enough effect to change the deontic status of the actions that they lead to, but they are not the only relevant factor. They contribute a finite amount to the effects of actions, but the contribution of other factors will often outweigh them. It is possible that a consequentialist will treat certain motives as in practice strongly-wrong-making

I will now mention some issues that I have not resolved. I have not addressed the question of what states of affairs have intrinsic value, which is a central topic in any consequentialist theory. I have assumed that whatever is intrinsically valuable is 'agent neutral'. I have also assumed that the determinants of deontic status are 'largely objective'. In the case of consequentialism this means that the states of affairs that have intrinsic value and that determine whether an action is right, wrong or merely permissible largely exist independently of an agent's beliefs and desires. Of course, in another sense consequentialism could allow that these determinants are 'subjective' if it takes a hedonistic form. For a theory like utilitarianism the determinants of deontic status are the

net amount of pleasure produced by the various options open to an agent, and pleasure is a subjective state. Still, it is an objective matter whether a given option does produce the greatest net pleasure or not.

Furthermore, I have said nothing about whether the best form of consequentialism is maximizing in structure or not. For all I have said, some ‘satisficing’ version of it, or a form having ‘an agent-centered prerogative’ is theoretically preferable.⁸ Nor have I really addressed any questions about ‘act’ versus ‘rule’ forms of consequentialism.⁹ The examples I have used are probably most easily understood as supporting an ‘act’ version of the theory, but ‘rule’ versions of the theory can take account of all the points made. If, for example, there are motives that are in practice strongly wrong-making then a rule consequentialist can accept a rule that prohibits acting from that motive.

It might be more accurate to say that our investigation has led to what we could call a ‘causalist’ approach to the deontic relevance of motives. That is to say, our investigation has found motives to be deontically relevant only because of their effects. It is conceivable that some non-consequentialist theory could accept causalism.¹⁰ But it seems to me that our conclusions are most at home in some form of consequentialism. So I will continue to say that our conclusions have supported the admittedly vague idea of extrinsic consequentialism.

III

Two Arguments Supporting PR

We can now return to some questions that were left unresolved, limiting our attention to extrinsic consequentialism, and how it will address them. It can show, in a plausible case, why MM is true: a motive like racism sometimes makes an otherwise permissible action wrong. We must address a question first posed in Chapter 1: if MM is true is that because motives are sometimes wrong-making but never obligation-making? Another way to pose this question is to ask whether Prichard and Ross' Claim is true.

PR: All actions that are morally obligatory are specified without mentioning their motives.

We saw that writers like Austin pretty explicitly state that utilitarianism entails PR. I called this the substantive argument for PR. There are two other arguments sometimes used to support MM, emanating ultimately from Kant, focused exclusively on the deontic category of the obligatory. They thus really concern PR. They rest on alternative interpretations of the idea of an obligation-making motive. Ross nicely captured these alternatives. On the one hand, the idea could be, "I ought first to have the motive, and in consequence to act under its influence."¹¹ On the other hand, the idea could be, "it is your duty to act from a certain motive if you have got it."¹² One argument for PR (and against the idea of an obligation-making motive) asserts that the idea of an obligation-making motive in the first sense violates the dictum that 'ought' implies 'can'. This is because on some occasions an agent will be obligated to act from a motive that is not available to her. The other argument for PR asserts that an obligation-making motive in the second

sense would constitute a hypothetical moral imperative. This, it contends, is a contradiction. Extrinsic consequentialism could simply sidestep these arguments if it said that MM is true and that PR is also true. To make these claims is to say that motives are sometimes wrong-making, but never obligation-making.¹³ However, this combination of assertions faces a significant obstacle. The reasoning I offered to show why consequentialism supports MM is symmetrical: if some motives are extrinsically bad then other motives are extrinsically good. That means that sometimes MM would be instantiated in cases where an act that would otherwise be merely permissible is obligatory because of its good motive. That motive is obligation-making, which means that PR is false. If this is correct then we cannot ignore the arguments supporting PR.

The first order of business is to address the dilemma that is in effect created by the two arguments. This requires us to choose one of the two interpretations of the nature of the obligation in question. We might accept Ross' first interpretation, and say that S would be obligated to act from M1 whether it is available to her or not. Or we might accept his second interpretation and say (roughly) that S is obligated to act from M2 if she now 'has' it.

It does seem that the first interpretation is problematic, and that consequentialists should accept the dictum 'ought' implies 'can' as it applies to the relevance of motives to obligations. In fact, insofar as the issues with respect to obligation-making motives are conceptual any other moral theory is free to accept the claims that consequentialism makes about them. Various forms of deontology also accept the dictum about 'ought' and 'can'—indeed, it is said to originate with Kant.¹⁴ Their proponents can thus agree that if there are any obligation-making motives then the obligations that they give rise to should

be thought of as conforming to Ross's second interpretation. So let us proceed on the assumption that when motives are obligation-making it is because something similar in form to Ross' second interpretation is true.

It will be helpful to think a little more about the interpretation of the idea of an obligation-making motive. Following Ross, we might put the idea this way:

OMM1: If S now 'has' the desire M1, which favors doing X, then she is obligated to X from M1.

Stocker in effect pointed out that a motive is available when a desire is 'had' or experienced.¹⁵ But our discussion of the availability of motives showed that it is also available in other cases. Keeping this point in mind will allow us to modify OMM1 somewhat. If S is rational and a certain reason is epistemically available to her then she can realize its application to her situation. Thus, S might be able to form a certain desire if she thinks about her situation properly. If this is true then this motive is available to her, and she can practically affirm it. The availability of this motive does not involve some mysterious use of will power to directly conjure up desires. It involves in the first instance rational thought. The desires will follow, in a rational agent. Therefore, if we want to say that only motives that are available to an agent can be obligation-making we should accept this second formulation of the idea of an obligation-making motive:

OMM2: If S would do X from M1 then S is obligated to X from M1.

If we accept this interpretation of an obligation-making motive we can also accept that S is obligated to X from M1 only when S can X from M1. This obligation does not violate

the dictum about ‘ought’ and ‘can’. However, if this is correct then S may come to have an obligation to X from M1 at a time when she does not ‘have’ the desire M1. It is the possibility that she will be moved by this desire to X that creates the obligation to do it, not her occurrent desire.

We now need to consider the other argument that relies on a claim about the concept of moral obligation. Kant asserted that moral obligations are categorical, and fundamentally different from hypothetical imperatives. So he can be taken to hold that a hypothetical moral imperative (or obligation) is impossible. His thinking might be put as follows. Moral obligations are things that we must do. That means that they must be done even if we strongly want not to do them. And it means that they must be done even if we have no desire to do them. But if S has a hypothetical moral obligation to X that means that S is morally obligated to X if she wants to X. This contradicts the idea that moral obligations are the practical necessities we take them to be.

The first point to make here is that it actually is not clear that Kant himself completely rejects the idea that an agent’s desires (or aims) can create a moral obligation for her. That is because in the Kantian system there are duties of ‘imperfect obligation’ such as beneficence and gratitude. With such duties “inclination is allowed to play a role in determining exactly what and how much we will do to carry them out.”¹⁶ Some Kantians seem to be uncomfortable with this idea, in part because of its suggestion that these duties are not (completely) categorical.¹⁷ But Thomas Hill, for one, is prepared to say that one sort of “latitude” allowed by such duties is the “freedom to choose to do X or not on a given occasion, as one pleases.”¹⁸

Furthermore, the Kantian objection seems to have in mind trivial or hedonistic desires. It is thought to be absurd to say that S is obligated to X because she likes doing X or finds it enjoyable to X. But the motives that would occasionally make an act obligatory could be gratitude or patriotism or even the sense of duty. Still, it might be said, the idea that any desire could generate an obligation is absurd, since desires come and go in a way that does not track our thinking about what agents must do.

Now, we have just seen that the best way to capture the idea of an obligation-making motive is OMM2. OMM2 does not make the occurrence of a desire the fact that is relevant for the existence of an agent's obligation. So we need to ask why there is something objectionable in asserting that sometimes agents have moral obligations because statements of the form of OMM2 are true of them.

Consequentialism can answer this question. It asserts that motives have deontic significance (when they do) because of their effects. If X from M1 is obligatory then the effects of X from M1 are better than the effects of not-X. But then shouldn't X from M1 be obligatory? Obligations in general are determined by the value of the (broad) effects of actions, and the motive of X in this case produces a difference in that value. In other words, the way that motives relate to obligations for the consequentialist is simply an application of how actions in general become obligatory. So unless there is something objectionable about a principle requiring people to maximize the amount of value in the world, there is no separate objection to idea that a motive can make an action obligatory. This means that the sense of 'must' in a statement of the form, 'S must do X from M1'—given OMM2 and consequentialism—is the same as it is in a statement of the form, 'S

must do X'. Consequentialism should therefore accept MM and reject PR. There can be both wrong-making and obligation-making motives.

Non-consequentialists can also accept the response just sketched to the second Kantian argument supporting PR. They can grant that if a motive does generate a moral obligation then the agent 'must' perform the relevant action in the same way that she 'must' perform an action like telling the truth. The non-consequentialist would need to explain what overarching moral principle makes it true that the motive generates a moral obligation. The consequentialist appeals to a principle requiring the maximizing of value. But it is certainly possible that there are other plausible principles that could be appealed to instead.

IV

A Final Puzzle

There remains one puzzle. Consequentialism should deny PR, and assert that there are obligation-making motives. But the fact is, we have little or no familiarity with the idea of an obligation-making motive, while we are familiar with the idea of a wrong-making motive. Our investigation began by noting how we understand statements to the effect that it is wrong to do some things for certain reasons, or from certain motives. But we do not say (or think), 'You must drive her to the airport from gratitude', or 'She must enlist in the army from patriotism'. We do say (and think) things like, 'we must be grateful to those who have helped us', and 'I must be kind to people'. But it is possible to

interpret these statements as concerning the development of the relevant motives, rather than acting from them on a given occasion. So there is something puzzling about the dearth of evidence for obligation-making motives.

We can begin to address this problem by noting an important difference between obligation-making and wrong-making motives. We have assumed that the dictum about ‘ought’ and ‘can’ applies to our concept of moral obligation. If we believe that it is wrong for S to X from M1 then we are likely to suppose that S is obligated to not-X. (In theory it is possible that S is permitted to X from M2, but this possibility probably does not occur to many people in many cases.) If S is obligated to not-X then the issues about her motivational capacities pertain to this action, not the action of Xing. Furthermore, we know that our deontic concepts are largely objective, so that it is likely that S’s obligation to not-X will seem objective, too. If so, it will seem that PR applies to it, and that S can not-X from any motive. Therefore, wrong-making motives do not narrow down the types of motive that allow an agent to fulfill her obligations.

On the other hand, an obligation-making motive does narrow down the types of motive that allow an agent to fulfill her obligations. She must act from one motive. The situation with regard to common sense moral thinking about this possibility seems to be the following. We all realize that many obligations (such as those regarding not-harming, promise-keeping and veracity) are objective. This makes PR plausible. We also accept that there are some moral obligations pertaining to gratitude and beneficence where motives seem somehow involved. Kant and Ross, among others, have tried to articulate the vague yet real requirements that are thought to exist here. But, finally, we accept the dictum about ‘ought’ and ‘can’. All of this easily leads us to the conclusion that PR is

true without qualification. Ross, indeed, thought that it represents the consensus of moral theorists throughout history, and presumably also the verdict of moral common sense.¹⁹

In contrast, a moral theory like act consequentialism entails that PR is false; it must allow that obligation-making motives exist in theory. Obviously act consequentialism does not completely correspond to moral common sense. But even on its account of things obligation-making motives would be hard to notice. If act consequentialism is true three conditions would have to hold in order for someone to reasonably assert a statement of the form, S is obligated to X from M1. To reasonably assert this she would have to believe that:

- 1) S's Xing from M1 has better consequences than her not-Xing;
- 2) It is at least partly in virtue of being from M1 that this is so;
- 3) M1 is available to S.

It is rare that someone (including S herself) knows that these conditions hold. And this may well mean that in practice people ignore this possibility altogether, and think more simply in terms of obligations to act. If so, the difference in practice between the consequentialist position on obligation-making motives and moral common sense may not be great.

¹ References to pages in VVV will be included in the text.

² More precisely: a love for itself of a love of something intrinsically good for itself is intrinsically good.

³ The structures here are similar to the practical syllogism, and once again there are problems that arise from false minor premises.

⁴ There is some debate about exactly which features of an object count as determining its intrinsic value. Hurka argues for relaxing Moore's strictness about this. VVV 6-7.

Hurka's view "allows a state's intrinsic goodness to be affected by its relational

properties.” VVV 6. I have no quarrel with speaking of the value of an attitude and its intentional object as ‘intrinsic’. Hurka clearly thinks that if the causal properties of an attitude are the only basis of its value then this is not intrinsic value. This is the question that I am addressing in the text. For more on the concept of intrinsic value see Shelly Kagan, “Rethinking Intrinsic Value,” The Journal of Ethics 2 (1998), pp. 277-97; Ben Bradley, “Is Intrinsic Value Conditional?” Philosophical Studies 107 (2002), pp. 23-44.

⁵ Moore had held this. Principia, op. cit., pp. 219-21.

⁶ This is perhaps least clear with regard to the fourth point Hurka makes. I see him as here trying to reconcile the agent-neutrality at the base level of his theory with the obvious fact that we tend to regard altruistic desires as better than self-interested ones. Sidgwick noted that the value of self-interested desires has been a controversial question historically, but even consequentialists who recognize the good results that self-interested desires often produce have tended to say that these desires are only of neutral value. ME 366; Bentham Introduction, op. cit., p. 115. There is a strain of thought stretching back to Butler that attempts to resolve this paradox. See Joseph Butler, Dissertation of the Nature of Virtue, paragraph 9; ME 428; Moore, Principia, op. cit., p. 172; Rawls, A Study, op. cit., pp. 207-10. These authors are insisting that human beings commonly act from self-interested desires, so that the amount of good they produce establishes a sort of baseline for our evaluation of other motives. The motives we regard as good are the ones that produce more good than that, ‘in general’. (I will say more about this important concept shortly.) This point is fully consistent with an extrinsic consequentialist approach to the value of motives.

⁷ This is the perspective on motives that Bentham adopts in An Introduction, op. cit., p. 115f. We also utilized this idea briefly in Chapter 6 in thinking about revenge and malice.

⁸ For the first idea, see Michael Slote, Common-sense Morality and Consequentialism (London: Routledge and Kegan Paul, 1985), esp. ch. 3; for the second, Scheffler, Rejection of Consequentialism, op. cit.

⁹ The locus classicus of act consequentialism is Smart in Utilitarianism, op. cit. Rule utilitarianism is defended in Richard Brandt, A Theory of the Good and the Right (Oxford: Oxford University Press, 1979); Brad Hooker, Ideal Code, Real World (Oxford: Oxford University Press, 2000).

¹⁰ Scanlon, Moral Dimensions, op. cit., possibly accepts this idea. But Scanlon does not address all of the issues that I have in this book, and so it is hard to be certain about this.

¹¹ Ross, Foundations, op. cit., p. 115.

¹² Ibid., p. 116.

¹³ In “Motive and Rightness,” op. cit., I suggested that this is the best way to support MM.

¹⁴ On this, see Sullivan, Kant’s Moral Theory, op. cit., p. 320, n. 6.

¹⁵ Stocker, “Intentions,” op. cit., pp. 589-90.

¹⁶ CKE 20. Kant’s division of duties is presented briefly in G 89 (422) and more elaborately in his later work, MetM.

¹⁷ Barbara Herman finds Kant’s implication in the Groundwork that there are obligations that have an “exception for the sake of inclination” (G 89) “very un-Kantian”. PMJ 63, n. 27. Marcia Baron finds imperfect duties as such to be “obligatory only in a highly attenuated sense.” That is because one is not required to help others, for example, “at

every opportunity or as much as possible.” Kantian, op. cit., p.162. Baron’s concern is not with the role of inclination in fixing a person’s obligations, but with the fact that a particular action is not required or necessary. This leads her to the surprising conclusion that an act of beneficence cannot have moral worth, despite what Kant says at G 66 (398). Ibid., p.163. Baron’s thought is that actions have moral worth only if they are done simply because they have to be done. But, she continues, imperfect duties generally don’t have to be done at a given time.

¹⁸ DPR 155. Hill says that there are two other sorts of “latitude” that are involved in imperfect duties. Ibid. Kant uses the term “latitude” at MetM 194 (390); 211 (411); 229 (433 n); 240 (446).

¹⁹ Ross, Foundations, op. cit., pp. 138-40.