

Chapter 7

The Availability of Motives

Two well-known passages in Kant center on a problem that is sometimes called the ‘availability’ of motives. One concerns the naturally sympathetic man whose mind becomes “overclouded by sorrows of his own which extinguish all sympathy with the fate of others”. Kant argues that even in this state, when he has no “inclination” to help others, he can do so, since he can act “from duty alone”.¹ The other passage states that the commandment to love our neighbor cannot mean that we must act from “pathological love”, that is, an emotion or feeling of love. Feeling is not under our control. We can, however, perform acts of “practical love”, that is, acts of assistance from the motive of duty.² One way some philosophers have combined the implications of these two passages is to say that for Kant the motive of duty is always ‘available’—by which they may mean, ‘available when an agent is obligated to act’—but motives like sympathy and “pathological love” are not. Such a claim about the availability of the sense of duty may well be one reason that this motive plays such a central role in Kant’s moral philosophy. It suggests that rational agents have a reservoir of motivation that they can always draw from, even when they believe they have no other normative reason to do what they are obligated to do. These passages also form part of the basis for the assertion, commonly found in Kantianism, that human agents have two qualitatively different sources of motivation, desire (or “inclination”) and reason.³

W.D. Ross takes Kant’s second argument in a surprising direction. He seems to accept it as being completely general.⁴ That is, he seems to claim that no motive, including the sense of duty, is always available (when an agent has a moral obligation).

Ross' focus is on obligation-making motives. He takes his view about the availability of motives, along with the principle that 'ought' implies 'can', to have implications about the deontic relevance of motives. He believes the two claims entail PRC.

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

He writes, "it is our duty to do certain things, and not to do them from certain motives."⁵ In Chapter 1 I mentioned a weakness that has been noted in Ross' reasoning. Any motive any agent now 'has', in some sense, is one that she can act from. So even if we accept that 'ought' implies 'can', the most that Ross is entitled to conclude is that we are never obligated to act from a motive that we do not now have. We began to examine the issue of obligation-making motives in Chapter 3. But we postponed consideration of them, and have since concentrated on wrong-making motives.

Ross' claim about the availability of motives has implications for the case of wrong-making motives—an issue he never considered. In Chapter 3 I presented a case where S refuses to shake R's hand because of R's race. Let us suppose that S believes that a refusal to shake R's hand for this reason is wrong. It might however be true that refusing to shake R's hand from self-interest would be permissible. But if Ross's claim about the availability of motives is correct, then S may not have the ability to refuse to shake R's hand from self-interest. When S does not have this ability, and we assume with Ross that 'ought' implies 'can', then Ross' claim has a further moral implication: S is morally obligated to shake R's hand. However, if S is able to refuse to shake R's hand from an acceptable motive then she acts permissibly in doing so. Note how in this case it

is the availability of a middling motive, not necessarily the sense of duty, which is critical with regard to permissibility.

The issue that Ross raises about obligation-making motives, and the issue about wrong-making motives, both turn on the question of our ability to generate motives in ourselves that we do not ‘have’. If an agent can generate the requisite motive, then Ross’s objection to PRC collapses; likewise, if S can generate an acceptable motive then she is not barred from performing the action her racism favored. In this chapter we will largely concentrate on the issue of wrong-making motives, but our discussion has implications for Ross’ issue. We return to it in the next chapter.

Aristotle and Kant asserted that a type of action can be performed from more than one motive. But one question we will be asking is whether one agent at a given time can choose to perform a given type of action from different motives?

We will now explore these fascinating issues. This chapter is divided into two parts. In Part I we examine the issue of the availability of motives. This part belongs to moral psychology: it is meant to be neutral with regard to substantive moral theory. That is, it is meant to state truths about human motives that any moral theorist can accept. In line with what we found in Chapter 2 I will take the issue of the availability of motives to be largely (but not entirely) coextensive with the issue of the availability of normative reasons. So I will sometimes speak of the availability of motives, and sometimes of the availability of reasons. I propose to address the issues of moral psychology by dividing them. Part A examines the availability of the motive (or ‘sense’) of duty; Part B examines the availability of two other representative motives, self-interest and sympathy. Part A focuses on some of Kant’s claims about the availability of the sense of duty. This

discussion has some interest in its own right, given the centrality of this motive in his moral theory.

In Part II of the chapter we turn to three substantive issues in the casuistry of motives. I have tried to show that consequentialism has the most convincing things to say about the deontic relevance of motives. It can give us a largely objective conception of deontic status, and yet still grant that MM is true.

MM: There is an action X such that if X were performed from one motive it would fall into one deontic category, and if X were performed from another motive it would fall into a second deontic category in virtue of this difference in motives.

But unlike the Kantianism embodied in the Formula of Humanity and in Slote's virtue ethics, consequentialism does not assert that there are any strongly wrong-making motives. That is, consequentialism does not assert that there is any motive such that whenever someone acts on it she acts wrongly. Also, consequentialism can give us a plausible account of why motives are deontically relevant in the cases that they are. I will utilize consequentialism in addressing the casuistical issues. For ease of exposition I will employ extrinsic consequentialism.

In Part I of the chapter I will show that all motives share three structural features. Seeing this tends to weaken the sense that there is anything distinctive about the availability of the sense of duty. On the other hand, I will grant that there are some differences between the sense of duty and an emotional motive like sympathy. In Part II of the chapter we will address these three questions: the deontic relevance of feelings; the importance of unconscious motives (and operatively covert desires); and whether it is ever possible at one time for an agent to act permissibly if she has one reason for

performing a certain action and to act wrongly if she has another reason for performing that action. My answer to this last question—which is surprising—is that this is sometimes possible.

Part I

The Availability of Motives

A

The Availability of the Sense of Duty

I believe it was Lawrence Blum who, in alluding to the passages in Kant just mentioned, first used the term ‘available’ to attribute to Kant the view that agents may always act from the sense of duty. He writes that Kant believes that “the moral motive must be something available to any moral agent at any time.”⁶ But in his earlier, seminal discussion of Kant, Bernard Williams alluded to Kant’s well-known doctrine that only actions performed from the sense of duty have moral worth.⁷ Williams portrayed Kant as believing that it would be “hideously unfair” if “the supreme value achievable by human beings” (that is, moral worth) were not available to everyone.⁸ The implication is that Kant did not believe that nature was unfair in this way. We are not directly interested in Kant’s doctrine of moral worth, but we are interested in the availability of motives.

It is important to conceive of the general issue of the availability of motives properly. The question is not whether a motive is ‘available’ simply as an inert psychological state such as a wish, or even as an evaluation to the effect that an action would be desirable. It is whether an agent can act from the motive. Furthermore, we have to picture the agent’s situation in a certain way. First, we need to stipulate that she is

deciding within a relatively limited span of time. This is to keep our discussion within the boundaries that are governing Kant's discussion of the sorrowful man who no longer feels sympathy. Kant and Ross are aware of the fact that we can transform our motives over time.⁹ Kant's point seems to be that we can suddenly become aware of someone who needs our help now. (Ross speaks of producing a motive in one's self "at a moment's notice".¹⁰) So the question is whether motives are as readily available as actions themselves seem to be. But, second, we need to stipulate that the situation as the agent conceives of it stays constant. Suppose that S is about to walk across the street to buy a newspaper. She then sees a friend in the same direction whom she is eager to talk to. This obviously could make a new motive available to her for walking across the street. With regard to the helping case we must assume that the issue that Kant is raising concerns an agent who has already become aware of the needy person. Third, we will limit the discussion to what we might call direct rational production of motives. We will not consider indirect methods of producing motives like taking a drug.¹¹ Fourth, we will grant that occurrent desires are available as motives. One implication of this is that if S already is aware that she has two normative reasons to X, then she can often choose to act from either one. (This may not be true in some cases, since she may regard each of the reasons as necessary normatively, so that one of the two is not 'reason enough' to act.) The issue we are investigating is whether any other motive besides these two is available.

A word is also necessary about the sense of duty as a motive. We will be thinking of the sense of duty in a generic sense. Different moral theories represent different specific forms of thought that determine, among other things, which actions are right. We will focus our discussion on Kant's statements about the sense of duty, but none of his

remarks or our conclusions pertains only to the specific forms of thought and motivation that constitute the sense of duty as that is conceived of in Kantian moral theory.

Let us now consider the availability of the sense of duty in the generic sense. As Blum states Kant's idea, it is obviously false. That is because there are epistemological preconditions for acting from the sense of duty. A rational agent could not normally perform an arbitrarily-specified action like turning on the television from the sense of duty. That is because she does not believe that doing this is her duty. For the same reason she could not necessarily find some action or other to do now from duty, because she might not believe that there is anything she is now obligated to do. A necessary condition of a rational agent's being able to act from the sense of duty is what we can call the 'epistemic availability' of the motive: in order for S to be able to X from the sense of duty S must be able to know, or believe, that she is morally obligated to X.¹²

To say that the epistemic availability of the motive of duty is a necessary condition of being able to act on it does not mean, of course, that an agent can only act from duty when she really is obligated to perform a certain action. If S falsely believes that X is obligatory she may be able to X from the sense of duty. Since we are interested in moral psychology, we can simply ignore the question of whether S does have an obligation. But I need to say a little more about epistemic availability.

In Chapter 2 I distinguished between normative beliefs about what is desirable and motives, claiming that the latter are certain desires. In what I called the standard case a rational agent desires to do what she believes is desirable and seeks to satisfy the desire. Beliefs about which actions are obligatory are thus normative beliefs about a distinctive kind of desirability. We can make one further assumption here about these beliefs. If S

suspects that she is obligated to X she may not understand why she is obligated to X. And it might then be doubted that she really can X from the sense of duty. Ross asserts, “Rightness is always a resultant attribute, an attribute that an act has because it has another attribute.”¹³ And he suggests that it is only by thinking or knowing that an act has another attribute that we think or know that it is right.¹⁴ So it is plausible to assume that if an agent rationally believes that she is obligated to X then she understands why she is obligated to X, that is, what attribute of it makes it obligatory. There are many other serious issues about the epistemic availability to an agent of beliefs and knowledge about what she herself is obligated to do. But we will not pursue them.

Let us now formulate one claim about the availability of the sense of duty, where it is understood that the abilities mentioned exist within the time frame. AD-N states a necessary condition for acting from duty. It captures the point that there are epistemic preconditions for acting from duty.

AD-N: When S is rational practically and S can X, if S can X because it is obligatory (i.e., from the sense of duty) then S is able to form the rational belief that she is obligated to X.

I have inserted two other clauses in AD-N that allow us to focus precisely on the issue of the availability of the motive of duty. The clause about S’s being able to X is inserted for the following reason. If S mistakenly believes that she is obligated to X then she may be unable to X from the sense of duty because she is unable to X at all. We need to insure that S can perform a certain action before we can discuss whether she can do it from duty.

The clause about S being rational practically is inserted to make explicit that it is the rational generation of motives that we are considering. It also bypasses issues pertaining to weakness of will. Let us grant that stipulating that S is rational practically means that she will do X if she believes she is obligated to do it.

AD-N seems to be true. (We will see shortly, though, that it contains an ambiguity.) And it certainly makes an important clarification about how available the motive of duty could possibly be.

A second claim, AD-S, states a sufficient condition for acting from duty. It raises some important further issues about this motive.

AD-S: When S is rational practically and S can X, if S is able to form the rational belief that she is obligated to X then S can X because it is obligatory (i.e., from the sense of duty).

AD-S also seems to be trivial. But, surprisingly, Kant himself in effect expressed doubt about it.¹⁵ In another well-known passage he asserts that experience makes it doubtful that any action has moral worth. Even when it seems as though only the sense of duty has moved a person to act, we can not be sure that some “secret impulsion” was not operating.¹⁶ The passage explicitly concerns the operation of a motive that the agent is unaware of. But the same issue arises when an agent is aware of a consideration. I described this situation as involving ‘operatively covert desires’. Suppose that S is aware that self-interest favors Xing. If S decides to disregard that and X only because it is her duty, Kant’s point implies that she cannot be sure that self-interest has nothing to do with her Xing.¹⁷

This passage (and some others¹⁸) suggests that Kant may not have accepted AD-S. We will use the following supposition to structure our investigation of this issue:

S chooses to X only from duty.¹⁹

We can imagine that S is aware of other considerations like self-interest that favor Xing, but chooses to disregard them, or that S has unconscious desires that favor Xing, or both. We saw that, in general, to have two conscious ends is “to ride two horses at once”, and therefore sometimes to reduce the chances that one will achieve both. So S might choose to X only from duty because that will actually make it more likely that she will achieve both aims. But this sort of case also has moral significance because S’s choosing to act only from duty is often thought to bespeak an admirable or virtuous character. Finally, the case is important in the Kantian doctrine about the moral worth of actions.²⁰

When S chooses to X only from duty S three possibilities exist. Roughly stated, they are:

P1. S actually Xes only from the sense of duty.

P2. S actually Xes from some motive or motives other than the sense of duty;

P3. S actually Xes from the sense of duty and from some other motive or motives.

It is, of course, an empirical question which of these three possibilities is possible or likely for human agents.²¹ I think we have reason to agree with Kant to this extent: P3 seems to be fairly common. This is confirmed, for one thing, by observing how people who sincerely profess to be performing some action ‘only because it is right’ actually do perform it. If S claims to be giving T the correct change only because it is right, but is

observed checking on how often T patronizes her store, we can reasonably suspect that a “secret impulsion” is part of the explanation for her action. S herself can carry out such an observation. In fact, the psychologist Timothy Wilson suggests that “self-outsight” is the better route epistemologically to our unconscious motives in many cases.²²

Let us then consider P2. What would it be like for a person to consciously choose to X only from the sense of duty, but in fact X only from, say, self-interest? The possibility of this being true seems more remote. Indeed, the general idea of S consciously choosing to X only from one motive but Xing only from another motive is strange. There are reports by psychologists of cases of ‘post-hypnotic suggestion’ that we might interpret in this way. The subjects’ behavior is bizarre.²³ In one case a subject in a trance was told to remove a flower pot from a window, wrap it in a cloth, put it on a sofa, and bow to it three times. After he awakened he did these things, and when asked why he had done so, said that he thought that the flower looked cold, that it needed to be warmed, and so on. It is natural to conclude that this person was completely mistaken about why he was doing what he did.²⁴ It is another question, of course, whether such a radical disconnection between an agent’s conscious end and her actual motive can occur with a person in a normal state of awareness. I do not think we can rule this possibility out a priori, so we should grant that P2 sometimes occurs.

The three possibilities thus reveal an ambiguity in AD-S. (The ambiguity also exists in AD-N, but we did not notice it.) There are two purely conative sorts of availability that concern motives as such. One is manifested when an agent consciously acts for a certain reason, or consciously pursues an end. We can call the availability of these motives ‘affirmative availability’. (This label is drawing on the idea that, in the

standard case, if S consciously decides to X for a certain reason it is as if she affirms that that reason is a good reason for her to X now.) If S believes that X is obligatory then the affirmative availability of the sense of duty for her means that she can choose to act for that reason. Furthermore, if S believes that self interest and duty both speak in favor of X the affirmative availability of the sense of duty for her means that she can choose to X only from duty, or from both.

The other sort of availability of motives can be called ‘operative availability’. This is the ability of an agent to make her conscious ends actually guide and produce her action. We can call the motives that actually guide and produce action, operative motives. We can also say, drawing on a distinction made earlier, that if M is an operative motive for an action then M explains that action, at least in part. Kant’s point in the passage about “secret impulses” can now be stated in two ways, corresponding to P2 and P3.

P2’: If S consciously chooses to X only because it is obligatory then her operative motive for Xing is some other desire, or her operative motives for Xing are other desires.

P3’: If S consciously chooses to X only because it is obligatory then her operative motives for Xing include the desire to do her duty, but also include other desires.

In either case, there turns out to be a distinction between the affirmative and the operative availability of the sense of duty.

We can see this distinction if we return to AD-S.

AD-S: When S is rational practically and S can X, if S is able to form the rational belief that she is obligated to X then S can X because it is obligatory (i.e., from the sense of duty).

Let us modify this statement slightly to clarify the issues. Consider AD-S1:

AD-S1: When S is rational practically and S can X, if S is able to form the rational belief that she is obligated to X then S can X **only** because it is obligatory (i.e., only from the sense of duty).

With respect to affirmative availability AD-S1 is still true, perhaps even trivially true. So long as S is able to arrive at the rational belief that she is obligated to X she can choose to X because it is obligatory. And she can choose to X only because it is obligatory—even if she believes that other normative reasons like self-interest favor Xing. We can therefore say the sense of duty as a motive is completely available affirmatively to a rational agent (in the time frame) if it is epistemically available to her (in the time frame).

What about the truth of AD-S1 with regard to operative availability? The case of S choosing to X only from the sense of duty shows that this is a different matter. The idea of there being “secret impulses” involves the two possibilities, P2' and P3'.

When P3' holds and S chooses to X only because it is obligatory, then S does make, and therefore can make, the sense of duty one of her operative motives for Xing. So the motive of duty is operatively available to S in the sense that she can make it one of her operative motives. But when P3' holds S does not make the sense of duty her only motive for Xing. The crucial question then is this: can S make it her only operative motive for Xing? Kant himself implies that the answer is, ‘No’. In one of his later writings he asserts that we have a duty to perfect ourselves morally. He takes this to mean that we must develop and exercise the ability do our duty only from the sense of duty. But he says that developing this capacity is difficult:

It is man's duty to strive for this perfection, but not to reach it (in this life), and his compliance with this duty can, accordingly, consist only in continual progress.²⁵

So Kant believed that our ability to make the sense of duty our only operative motive for doing our duty (in the time frame) is incomplete.

This passage raises interesting questions. Korsgaard glosses Kant as maintaining that “valuing humanity...is an internal labor with which we are never simply done.”²⁶ But, what sort of “internal labor” does Kant envision us undertaking? If I cannot choose all at once to make the motive of duty my only operative motive, what can I do over time to insure that it is? We cannot rule out that, say, psychologists may discover that there are rational techniques that an agent could use, but it seems fair to say that, whatever they might be, it is not now true that a human being can insure in the time frame that her only operative motive is the sense of duty. To this extent, at least, the sense of duty is incompletely available operatively, even when an agent rationally believes that she is obligated to perform a certain action.

I will dispense with discussing the implications with respect to operative availability of P2' holding when S chooses to X only because it is obligatory. That is a rather bizarre possibility, and we already have reached the critical conclusion. That is, we have concluded that the sense of duty is incompletely available operatively for rational human agents when they have well-founded beliefs about what they are obligated to do.

Before we consider the availability of other motives, it is worth noting how Kant's remarks about “secret impulses” undercut the contrast he tries to draw between emotional motives and the sense of duty. Kant's remarks about “secret impulses” are certainly plausible insofar as he is saying that P2' is often the case when an agent chooses

to act only from duty. And we cannot rule out that P3' is sometimes true. The description that Kant gives of the sympathetic man whose mind becomes “overclouded by sorrows of his own” asserts that this man, who has no “inclination” to help others, can still help them, since he can act “from duty alone”.²⁷ We are to picture someone who seems to lack any non-moral desire to help another person. Yet, Kant implies, he has access to a reservoir of motivation that still makes it possible for him to help. But if P2' is true in this case, his operative motive will not be the sense of duty. If P3' is true in this case, his operative motive will not only be the sense of duty. In these cases, the reservoir does not provide any motivation—if it even exists; or it is not the only source of motivation.

B

The Availability of Other Motives

In order to understand the general issues here we will now consider the availability of two other types of motive. One was discussed by Kant himself, *viz.*, sympathy; the other is self-interest. We will then have considered an emotional and a ‘cool’ non-moral motive. There are three sorts of availability to examine in each case: epistemic, affirmative and operative. We can proceed by asking whether each motive is available in each of the three ways. I again stipulate the conditions mentioned in the last section: we are dealing with a relatively short span of time; the factual situation of the agent as she conceives of it remains constant, etc. We will see that the availability of these other motives largely parallels the availability of the sense of duty.

There are, however, a few differences among the three motives that require some comment. I will mention one now. Sympathy probably differs from the sense of duty in that the latter motive requires more sophisticated conceptual abilities. A child who does not have many conceptual or linguistic abilities could act from sympathy, but she could not do something ‘because it is right’ in the richer sense that we have discussed. Still, the most appropriate comparisons to the sense of duty would involve agents who have developed conceptual and cognitive abilities. It is with such agents that epistemological issues can explicitly arise, and they can still act from sympathy. Similarly, there is a sense in which an animal or an infant could act from self-interest when, for example, it seeks food or shelter. But, again, we should focus on cases where an agent with developed conceptual abilities chooses to do something because she believes it furthers her own interests.

Epistemic Availability. We begin with the epistemic availability of sympathy. We again see the relevance of epistemic considerations. If S is rational she could not, say, walk to T’s house from sympathy unless she believed certain things about T’s situation and her own activity. Abstractly stated, the general ‘desirability characterization’ of sympathy seems to be this: if S acts from sympathy with respect to T then S desires to relieve T’s suffering as an end.²⁸ Let us assume with sympathy what we implicitly assumed with the sense of duty, namely, that the agent accepts the general desirability characterization when she enters the situation under discussion, or, at least, that she is disposed to relieve people’s suffering as an end in some circumstances. Even if S is in general disposed to relieve people’s suffering as an end in some circumstances, she cannot X from sympathy unless she believes that Xing will relieve someone’s suffering.

Obviously the point relates to the structure of the practical syllogism, and strictly parallels the sense of duty.

However, the epistemic availability of sympathy may be more complex than that of the sense of duty. The possible complexity is due to the ‘weighing’ of normative reasons. The desirability of a course of action may need to be weighed against its undesirability, and against the desirability of other actions. It is possible that this part of the process of reaching a rational normative belief is radically simplified with regard to the sense of duty, since some philosophers hold that a moral obligation always provides an agent with the strongest normative reason for action that she could have. If so, the weighing issue is always straightforward with regard to the sense of duty, unless moral obligations can conflict. If sympathy does not always provide an agent with the strongest normative reasons then the issue of epistemic availability includes the question of whether the agent has the intellectual resources to do the normative weighing rationally.

If we assume that sympathy does not always provide an agent with the strongest normative reasons we can formulate the following necessary condition for the availability of sympathy. We assume that X is some action that S believes relieves T’s suffering.

AS-N When S is rational practically and S can X, if S can X because it is desirable as an end to relieve T’s suffering (i.e., X from sympathy) then (i) S is able to form the rational belief that Xing is desirable as an end; (ii) S is able to form the rational belief that the correct weighing of all the relevant normative reasons determines that S has most reason to X; (iii) S is able to form the rational belief that the desirability of relieving T’s suffering as an end constitutes S’s strongest reason to X.

We will find again an ambiguity in this statement but, setting that aside, AS-N is a plausible way to capture the fact that for a rational agent the motive of sympathy has

epistemic preconditions largely parallel to those of the sense of duty. (If moral obligations are not always an agent's strongest normative reasons for acting then AD-N needs to be revised to create a strict parallel to AS-N.)

The issues of the application of general desirability characterizations and weighing also arise with respect to self-interest in its conceptualized form. If S is rational, then she can only, say, make a certain investment from self-interest if she can rationally believe that this investment furthers it. And again, if S believes that making the investment is desirable because it furthers her self-interest, she might also believe that she has a stronger normative reason to perform an act which is not in her self-interest, or which furthers it more effectively. I could present a formulation exactly parallel to AS-N for self-interest, but I will dispense with that.

The same two issues of epistemic availability arise for a cognitively sophisticated agent with regard to any motive whatsoever. As far as the application of general desirability characterizations go, we can see that one critically important reason why a specific motive is not epistemically available to an agent is that her objective situation—in the comprehensive sense that includes her possible futures—will not allow for it. It is possible for S to rationally believe that she will realize some value by a given action when this is not true. But insofar as agents are rational they will seek to realize values that their situation calls for, or at least allows.

Affirmative Availability. We granted above that if 'acting only from duty' in AD-S1 is interpreted to mean 'consciously choosing to act only from duty' then it is true: if S is able to form the rational belief that she is obligated to X then S can consciously choose to X only because it is obligatory. The parallel statement about sympathy (AS-S1)

is also true in this sense. We again assume that X is some action that S believes relieves T's suffering.

AS-S1 When S is rational practically and S can X, if (i) S is able to form the rational belief that Xing is desirable as an end; (ii) S is able to form the rational belief that the correct weighing of all the relevant normative reasons determines that S has most reason to X; (iii) S is able to form the rational belief that the desirability of relieving T's suffering as an end constitutes S's strongest reason to X then S can X only because it is desirable as an end to relieve T's suffering (i.e., X from sympathy).

These points are also completely general, and pertain to any motive whatsoever.

A similar formulation could be stated for self-interest.

Operative Availability. The sense of duty is incompletely available operatively when it is available affirmatively. This means that when 'acting only from duty' in AD-S1 is interpreted to refer to the agent's operative motive it is not always true. It is plausible to claim that sympathy is also incompletely available. To see this we can consider the case where S chooses to X only because it is desirable as an end to relieve T's suffering and the three conditions mentioned in the protasis of AS-S1 are true. When these conditions exist we can again distinguish two possibilities concerning S's operative motives.

P4: S's operative motive for Xing is some other desire, or her operative motives for Xing are other desires.

P5: S's operative motives for Xing include the desire to relieve T's suffering as an end, but also include other desires.

If the parallel to the sense of duty is correct, we can conclude that, in the conditions mentioned, P5 is often true, and P4 may sometimes be true. When 'acting only from

sympathy' is interpreted to refer to the agent's operative motives, and either possibility obtains, AS-S1 is false. S cannot X only from sympathy.

Self-interest might be thought to operate a little differently. Mary Mothersill once noted that Kant never suggests that the sense of duty could be a "secret impulsion". He apparently believed that choosing to act only from a higher motive might be accompanied (or even replaced) by the operation of lower ones, but he did not believe that choosing to act only from a lower motive might be accompanied or replaced by the operation of a higher one.²⁹ If this view is correct it suggests that when the three conditions parallel to the conditions in AS-S1 hold with regard to self-interest, and an agent chooses to act only from self-interest, that self interest actually is completely available operatively. But even if we accept that higher motives never operate as "secret impulsions" we believe that choosing to act from only one lower motive, like the pursuit of money, might be accompanied by the operation of another lower one, like the pursuit of power. In this sense self-interested motives are also incompletely available operatively.

Affective Availability. At this point an objection will naturally arise about what we have found with regard to sympathy. It will be said that acting from sympathy does not consist merely in choosing to relieve someone's suffering as an end. That is too simple a picture. It neglects the element of feeling or affectivity. If S is sympathetic to T then S is pained by T's situation, or 'feels bad' about it. In terms of the philosophy of emotion we could say that the picture of sympathy utilized thus far is defectively 'cognitivist', since it interprets acting from an emotion as simply choosing to act for a certain reason. Kant's argument about pathological love seems to rest on this very point. He seems to be assuming that feelings are an essential component of the emotion of love

(and presumably sympathy). But the having of these feelings cannot be commanded, as the performing of an action can. We can put his point into our terminology by saying that we do not have ‘affective availability’ with regard to the feelings that are involved in emotional motives.

The question of the role of feeling in emotion is one that reaches well beyond the topic of motives, since emotions are often felt in contexts where there is no opportunity for action. (I mentioned grief as an example.) Now, it is clear that, say, the anger of a rational agent is not merely a distinctive feeling. If it were, it would be surprising that it is generally followed by a desire to hurt another person. But the desire ‘makes sense’, given the emotion. This strongly suggests that the emotion of a rational agent contains a thought, for example, that another person treated the agent contemptuously. However, it is a further and interesting question whether emotions necessarily involve feelings. I think that emotions do incorporate feelings, although there is some debate about this in the literature.³⁰ A paper by Michael Stocker gives a good argument to show that emotions involve feelings.³¹ He asks us to imagine that S is sitting in an airplane in flight. She might entertain the thought ‘this plane could crash’ completely coolly and unemotionally. But she might entertain the same thought with anxious feelings, and hence, emotionally. The difference between entertaining this thought unemotionally and emotionally consists in having certain feelings.

If having an emotion does involve having certain feelings, then an agent’s acting from that emotion also involves her having those feelings as she acts. Kant is therefore correct in claiming that the feeling element in emotional motives is not available to a

rational agent in the time frame we are considering. The availability of emotional motives differs from that of the sense of duty, and indeed, from self-interest.

While the affective element in emotions presents an extra issue of availability, it is important to note that otherwise sympathy works exactly like self-interest and the sense of duty. All three motives incorporate desirability characterizations. Such characterizations have epistemic preconditions: in order to believe that it would be desirable to relieve T's suffering as an end S must believe that T is suffering, for example. When these preconditions are satisfied it is possible for an agent to choose to act because the action is desirable in the specified way. But in choosing to act for this reason the agent cannot guarantee that it is her only operative motive.

Here, again, Kant's own arguments have an equivocal force. He says that "practical love" is available when "pathological love" is not. I think his point is related to the one I just made. But Kant errs in suggesting that the issue concerns moral obligations. It really concerns rational action. There is 'practical fear' and even 'practical lust'.³² These are desires that occur without the feelings normally associated with them. A rational agent could choose to act on these desires.

Conclusions. These, then, are our general conclusions about the availability of motives. I have simplified them somewhat for purposes of summarization.

- 1) All of the motives of rational agents have epistemic preconditions.
- 2) For every non-emotional motive, if it is epistemically available to a rational agent then it is affirmatively available to her.
- 3) For every non-emotional motive, if it is epistemically available to a rational agent then it is incompletely available operatively.

- 4) Emotional motives have an affective component.
- 5) The affective component of emotional motives is not available in a small time frame.
- 6) For every emotional motive, if the desirability characterization of the motive is epistemically available to a rational agent then it is affirmatively available to her.
- 7) For every emotional motive if the desirability characterization of the motive is epistemically available to a rational agent then it is incompletely available operatively to her.

Part II

Three Casuistical Questions

A

The Importance of Feelings

The conclusions of the previous section are propositions in moral psychology. They are not part of some moral theory, and they should be acceptable to proponents of all moral theories. We now need to consider three moral questions that are raised by our conclusions. In answering the second and the third question we will utilize extrinsic consequentialism. The first question I want to consider is the possible deontic relevance of the feeling component of emotional motives.

We saw that writers since Sidgwick have suggested that recipients of certain helpful actions welcome them more if they stem from affection rather than a sense of duty. This suggests that the recipients of helpful actions might care specifically about the

distinctive feelings experienced by an agent who helps them; they might welcome the actions more when they believe that the agent ‘feels bad’ about their plight as she acts.³³ (This does not mean that they have any interest in whether, say, ‘sympathy’ is the correct term for the motive moving the agent.) Given what we have found about the relevance of motives in consequentialism this attitude could in theory make a difference in the deontic status of actions.

I believe, however, that anyone who accepts all of our conclusions in moral psychology would not have this attitude. If S only chooses to act because doing so will relieve T’s distress, but S does not ‘feel bad’ about his plight, S is not pretending to be concerned about T’s well-being, nor is she acting from self-interest or some other motive that characterizes the relief of T’s suffering as a means. She accepts the same desirability characterization that moves someone acting from sympathy. If T understood all this, I do not think he would care about S’s feelings.

B

The Deontic Relevance of Operative Motives

There is a difference between the motives (that is, reasons) that agents consciously act on (or affirm), and the motives that actually move them to act. The two cases where they differ are unconscious motives and operatively covert motives. I have claimed that consequentialism can give the most satisfactory defense of MM; this involves asserting that sometimes motives affect the deontic status of actions. Now we need to ask this question: which motives affect the deontic status of actions, the affirmed

ones or the operative ones? I will argue that the answer is the operative motives, or at least many of them. But in giving this answer I must emphasize that these are not mutually exclusive categories. Consider the generalization of P3' above:

GP3': If S consciously chooses to X only because of motive M1 then her operative motives for Xing include M1, but also include other motives.

When GP3' is true, affirmed reasons are a proper subset of operative reasons: if S chooses to act for a reason then it is her operative motive, although non-affirmed reasons may be operative as well. We could say that in these cases non-affirmed motives piggy-back on the affirmed ones. When the generalization of P2' occurs, though, the agent's affirmed and operative reasons are mutually exclusive:

GP2': If S consciously chooses to X only because of M1 then her operative motive for Xing is some other motive, or her operative motives for Xing are other motives.

I suggested that such situations are much less likely to occur. But my conclusion below about the relevance of operative motives applies to such cases too.

It is important to note that we are considering unconscious desires and operatively covert desires that directly bring about or modify action. Unconscious desires can also affect what a person believes, by distorting her evidence and her reasoning, both theoretical and practical. In this way unconscious desires can also indirectly have an effect on what a person does. In fact, these more indirect (and hence, insidious) processes may be much more common.³⁴

It will be helpful to have before us an example of an unconscious motive for acting. In keeping with our focus thus far, let us consider a wrong-making motive. I will

modify an example discussed by Timothy Wilson.³⁵ Suppose S chooses to play tennis with her twelve year old nephew T. She chooses to play ‘for fun’. But S has a competitive desire that operates during the match and causes her to play ever more aggressively. This desire leads to a style of play that eventually upsets T, who is not a skilled player. S’s consciously chosen reason was to enjoy a game of tennis, and this was operative. However, she had another operative motive.

Consequentialism can give a good account of the deontic significance of motives. How will the extrinsic consequentialist, in particular, deal with a case like this? Well, even motives that an agent is unaware of can have bad effects, as our example shows. So it seems that it is operativity that this approach will take to be fundamental. It is true that unconscious motives often bring about effects in different ways than affirmed motives do. And people who are affected by what they believe are unconscious motives respond to them differently than they do to what they believe are affirmed motives. An agent is much more fully ‘behind’ a consciously chosen motive than she is an unconscious one. Still, unconscious motives cannot be ruled out as deontically irrelevant. The extrinsic consequentialist can therefore say that, given certain further facts (discussed below), S acts wrongly in playing so competitively.

It may be objected that this emphasis on operative motivation is mistaken, since a deliberating agent can only take into account features of the world (including herself) that she is aware of. Unconscious operative motivation is out of view for purposes of deliberation. Any sort of deliberation recommended by Bentham (or Kant³⁶, for that matter) concludes with reasons that are consciously chosen.

This objection is only partly correct. A deliberating agent can take account of features of the world that she is able to become aware of, not only features of the world that she is aware of. And sometimes an agent can become aware of—or, at least, surmise—her unconscious motivation, or its effects. While playing tennis S may become introspectively aware of clues that her competitive desires are operating. For example, she may become aware that she is starting to think of her nephew in hostile terms. Or S may use “self-outsight” to see that her behavior is changing as she plays, so that her serves are getting faster, whatever exactly her motive for doing this might be.

Obviously there is not the same difficulty in the case of operatively covert motives. If a mature agent is aware that X would further her self-interest, say, but she does not affirm that as a reason to X, she has a sort of warning that her self-interest may be operative when she does X.

This does not mean that for extrinsic consequentialism all operative motives affect the deontic status of an action that they lead to. What it means is that sometimes operative motives that that an agent does not consciously choose to act on are relevant deontically. A rough idea³⁷ of which operative motives are relevant is captured by this principle concerning the deontic relevance of unaffirmed motives.

DRUM: If S does X and an unaffirmed motive of hers, M1, in Xing has a bad effect such that (i) Xing in this way has worse effects than some alternative open to her; (ii) she could³⁸ have foreseen this effect and (iii) could have avoided it then X is wrong in virtue of being done from M1.

The principle is focused on operative motives whose effects are foreseeable, or epistemically available. It is the epistemic availability of the effects of actions that is at

issue here; above it was largely the issue of the availability of normative judgments about actions that was at issue.

An operative (but unaffirmed) desire that causes some bad effect that was foreseeable can be thought of as involving a sort of negligence. It is comparable to an agent's doing something that has unforeseen, but foreseeable, bad effects that she does not consciously seek to produce. A standard example of this is a worker on a building who carelessly throws an object away that strikes a pedestrian below.³⁹ Deontic status in consequentialism is determined by which action open to an agent has optimal effects. DRUM states that sometimes the optimality of an action open to an agent is affected by its operative but unaffirmed motives. So, returning to our example, we can say that if S was able to foresee the effects of her competitive desire then she acted wrongly in, for example, serving the ball so forcefully. It is plausible to think that mature agents do act wrongly on occasion because of the effects of operative motives that they do not choose to act on.

C

Replacing Wrong-Making Motives

Finally, we can return to some questions left unanswered in Chapter 3. Actions have motives, but motives are not themselves actions. We do not choose our motives (or reasons) separately from actions. We choose to act for reasons. But these facts still leave room for the following two questions. The first is another psychological question: is it ever possible that at one time, a rational agent S can X for one reason or X for another

reason? The second question is a related substantive moral question: if it is ever possible at one time for a rational agent S to X for one reason or to X for another reason, is it also possible that S Xes permissibly if she does it for one reason and S Xes wrongly if she does it for the other reason? It would be surprising to find that this could be true.

As I mentioned before, I will come at these questions by reconsidering the scenario presented in Chapter 3. We can suppose that S is inclined to refuse to shake R's hand because R is black. In Chapter 3 I argued that, given the way the example was described, an extrinsic consequentialist can say that racism is here a wrong-making motive. Our next question is, can S refuse to shake R's hand from self-interest (say)?

Given what we have learned about the epistemic preconditions of all motives the answer is going to be 'yes' in certain cases. An agent may be aware at one time that two motives favor Xing, and she can then often choose to X for one reason, or choose to X for another reason. For example, S may be aware that Xing is obligatory and in her self-interest, and she could therefore choose to X only because it is obligatory or choose to X only because it is in her self-interest. (Whether only one motive would be operative in each case is another matter.) It follows that S might initially see only that X is in her self-interest but come to see, after reflection, that X is obligatory, and thereby acquire the ability to X from duty. Similarly, S might initially see only a racist reason to refuse to shake R's hand, but come to see, after reflection, that it is in her self-interest to do this, and thereby acquire the ability to choose to refuse to shake R's hand from self-interest.

Any given agent would also often be in situations where it is not possible at one time to choose to X from one reason or choose to X for another reason. In particular, it might not be possible for S to refuse to shake R's hand from racism or refuse to shake R's

hand from self-interest. This might not be possible because self-interest might not be epistemically available to S. This will be true in the cases where S has no reason to believe that refusing to shake R's hand would be in S's interest: she has no reason to believe that shaking R's hand would, say, cause her physical pain or illness. In these situations S cannot choose to refuse to shake R's hand from self-interest. Here again we see the crucial importance of a rational agent's objective situation.

It might be replied that S could engage in some form of self-deception and convince herself that shaking R's hand would be bad for her health, so that she could refuse to do so from self-interest. A critical problem with this suggestion is that, given our stipulations about the time frame of the problem and S's epistemic situation, S simply could not succeed in such outright and hurried self-deception. The issue here is basically that of belief and the will, and there are real limits to what a rational agent can will to believe, especially in the face of evidence.⁴⁰

If refusing to shake hands from racism is wrong in this case, what moral position is S in when no motive is affirmatively available to her that would make refusing to shake R's permissible? S is morally obligated to shake R's hand. She has no permissible way to refuse to do so.

Let us return to the cases where it would be possible for S to choose to refuse to shake R's hand from racism or choose to refuse to shake R's from self-interest. Does this mean that S can choose whether to refuse to shake hands permissibly or choose to shake hands wrongly? In addressing this further question we need to remind ourselves of the features of motives that make them relevant deontically in extrinsic consequentialism.

1. They can indirectly affect the intrinsic value of the action itself;

2. They can affect how an action is done, and thus indirectly affect the intrinsic value of an act's narrow consequences.
3. They can influence people's responses to the act because they care about why the agent did what she did. These responses in turn can affect the intrinsic value of the act's narrow consequences.

(I defined 'narrow consequence' to include only those events that occur after an act, qua intentional bodily movement. 'Narrow consequences' include two kinds of events (or facts): first, the causal effects of an action, considered as an intentional bodily movement; second, any states of the world following the action that would have been different had the agent chosen to act differently.) It is the third feature that is most important here. I originally described R's reaction to S's refusal to shake hands as being humiliation and anger because it was obvious to her that S's motive was racism. This point was the main reason that the extrinsic consequentialist asserts that S's refusal to shake hands is wrong.

If we are now granting that at one time S can refuse to shake hands for one reason or refuse to shake hands for another reason we are re-describing S's options in a critically important way. We are supposing that she has three options, not two: shaking hands; not shaking hands from self-interest; not shaking hands from racism. In order for the extrinsic consequentialist to say that S's refusing to shake hands from racism is wrong and her refusing to shake hands from self-interest is permissible it must turn out that refusing to shake hands from racism produces less good than one of the other two options, and refusing to shake hands from self-interest produces at least as much good as any other option.

It does not follow from the fact that at one time S can refuse to shake R's hand from self-interest or refuse to shake R's hand from racism that S would act wrongly if she did it for one reason and permissibly if she did it for another. This does not follow even if one of the reasons that S has for acting is a weakly wrong-making motive like racism. The explanation is that R might not perceive the difference in S's motives.

To see this, let us first suppose that the self-interested reason that S becomes aware of is also evident to R. A simple way to flesh out this supposition is to borrow a feature of the other agent described in Chapter 3, and suppose that S injured her right hand and now has it bandaged. As she considers her own situation she would then be able to see that refusing to shake R's hand is in her self-interest. Thus, S could choose to refuse to shake hands only from self-interest. But if S has a bandage on her hand then it no longer is plausible to suppose that R would be upset if S refuses to shake her hand from racism, since this would be hard to distinguish from S's refusal to shake her hand from self-interest.

There is an objection that could be made to this assertion about the imperceptibility of S's motive. Motives affect how an act is done, so that it is possible that S's refusal to shake hands from racism would be perceptibly different from her refusal to shake hands from self-interest. However, if R and S were personally unacquainted and R's reactions were based just on what she could infer from their quick interaction, it seems unlikely that she would be as upset by the refusal of someone with a bandaged hand to shake her hand (from racism) as she would be upset by the refusal of someone with an unbandaged hand to shake her hand (from racism). But it was the reaction we supposed that R would have to a refusal to shake hands by someone without

bandages that led us to say that that such an action is wrong in virtue of the agent's motive. So the very supposition about the objective circumstances that we introduced to make it possible that S can refuse to shake hands for two different reasons makes it less likely that the difference in her reasons would be detectable by R.

This is not to deny that racism is a weakly wrong-making motive in extrinsic consequentialism. But what we can now say is this: when a less problematic motive like self-interest is epistemically available to the agent and others as a reason for her to do what racism also favors, this is less likely to be a situation where racism is in fact wrong-making. Racism is more likely to be wrong-making when there is no other motive epistemically available to all concerned for S's doing what it favors. In other words, if R can see that S's self-interest also favors not shaking hands, it is less likely that she will be disturbed by S's refusal to do this—whatever S's reason is.

We have a different sort of situation when S has a reason other than racism to refuse to shake R's hand that is epistemically available to her but is not evident to R. It could be painful for S to shake R's hand even though S does not have a bandage on her hand. In this situation some features of S's circumstances—which include, after all, S's own psychological states—that S finds provide her with a normative reason to act are not evident to other people. This sort of situation is more open to misunderstanding, and hence to the production of the bad effects that are relevant to an extrinsic consequentialist.

Thus, it is also possible that R is misinformed or makes mistaken inferences about S's motives: R may be angered in a case where S chooses to act only from self-interest. So the bad effects of motives that we are discussing can be brought about when an agent

acts from a generally acceptable reason like self-interest. The fact that R's mistaken belief about S's reason for acting can upset or anger her is indirectly a sort of confirmation of the position I have been trying to defend.

But, to repeat the surprising main conclusion we have reached: it is sometimes possible that at one time S Xes permissibly if she does it for one reason and S Xes wrongly if she does it for another reason.

III

Summary

I have argued that the issue of the 'availability of motives' covers four distinct problems: the availability of normative considerations to an agent; the ability of that agent to choose to act on those considerations; the ability of the agent to determine which desires move her to act; and her ability to call forth certain distinctive feelings. I called these four aspects of the availability of motives 'epistemic', 'affirmative', 'operative' and 'affective'. While I granted that there are some important differences in which motives are available to an agent at a given time in a given respect, I argued that there are fundamental similarities among all motives with regard to epistemic, affirmative and operative availability.

In the second part of the chapter we examined three moral issues that arise once the moral psychology of the availability of motives has been understood. I argued that people affected by the actions of others would be unlikely to be interested in the agents' feelings as such; that extrinsic consequentialism will focus its attention on operative

motives—which often coincide with an agent’s affirmed motives; and that extrinsic consequentialism will acknowledge that it is possible at one time for an agent to act wrongly if she performs a certain action from one motive and permissibly if she does it from another. Motives are most likely to be wrong-making precisely in those cases where the agent is believed by the people affected by her action to have no other reason epistemically available for acting as she does.

¹ G 66 (398-399). A similar point is made about a man naturally “cold in temperament and indifferent to the sufferings of others”. *Ibid.* Marcia Baron reasonably objects to Paton’s translation here of “aus Pflicht” as “for the sake of duty”. She suggests “from duty” or “out of duty”. *Kantian Ethics Almost Without Apology* (Ithaca, New York: Cornell University Press, 1995), p. 13. I emend Paton’s translation accordingly.

² G 67 (399).

³ See Chapter 2, n. 6 above.

⁴ RG 5.

⁵ RG 23.

⁶ Blum, *Friendship*, *op. cit.*, pp. 30-36, at 30. Blum credits Schopenhauer, perhaps incorrectly, with first pointing this out in Kant’s writings. Schopenhauer wrote of “a categorical imperative that is always at our service.” *On the Basis of Morality*, trans. E. F. J. Payne (Indianapolis: The Library of Liberal Arts, 1965), p. 42. He did indeed criticize Kantian claims about the sense of duty on p. 75, but it is not clear that it is availability that is at issue. John Rawls uses the term ‘available’ to discuss a closely-related issue. TJ 177. The whole passage on “the strains of commitment” of various conceptions of justice, and their “psychological stability” on pp. 176-8 is relevant. Rawls makes the Kantian-sounding criticism of the utilitarian conception of justice that it “is threatened with instability unless sympathy and benevolence can be widely and intensely cultivated”. 178. This is because utilitarianism may require some people to make great sacrifices for their entire lives in order to increase the total happiness. It is surprising that Rawls does not consider the possibility that a utilitarian sense of duty would be available when sympathy gives out. See also Timmons, “Motive,” *op. cit.*, p. 270.

⁷ G 57-8 (390); 66 (398); 68 (400); 74 (406).

⁸ Bernard Williams, “Morality,” *op. cit.*, p. 228.

⁹ For Kant, see MetM 241 (447); 196 (392). For a gloss see CKE 21, 84. MetM 250-251 (457) is moving and interesting. For Ross, see RG 5.

¹⁰ RG 5.

¹¹ See Millgram, *op. cit.*, ch. 2. He thinks that such methods, if used by an agent who knows there is no rational support for the desire to be produced, will generally fail.

¹² See John Campbell, “Kantian Conceptions of Moral Worth,” *Canadian Journal of Philosophy* 13 (1983), pp. 535-540, for this idea.

¹³ Ross, *Foundations*, *op. cit.*, p. 168.

¹⁴ Ibid. John Rawls criticizes a doctrine about the value of motives that pertains to this point. He calls it “the doctrine of the purely conscientious act”. It “holds...that the highest moral motive is the desire to do what is right and just simply because it is right and just, no other description being appropriate.” TJ 477 He says: “...on this interpretation the sense of right lacks any apparent reason; it resembles a preference for tea rather than coffee.” TJ 478. Cp. Reath, op. cit., pp. 138-43 and 169 n. 39.

¹⁵ Cp. Walter Schaller, “Should Kantians Care about Moral Worth?” Dialogue 32 (1993), pp. 31-32.

¹⁶ G 74-5 (406-407).

¹⁷ Sidgwick agrees. ME 203. When Sidgwick discusses availability it is this sort of case he has in mind. Perhaps he rejected the very idea of unconscious motives as we now think of them. See also ME 363, n. 1; 368. In this sense Kant’s discussion of “secret impulsions”, G 74-5 (406-7), is more modern. Cp. RG 170.

¹⁸ G 87 (419); MetM 241 (447); 196 (392).

¹⁹ Other cases might be used to clarify the issues about availability that we are concerned with. For example, we might consider the case where S chooses to X from duty and self-interest. But, as I explain, the case in the text is important for various reasons. One reason is mentioned in the next note.

²⁰ It is important here, because Barbara Herman, in a well-known essay on this issue, claims that Kant’s doctrine is precisely that an action has moral worth only when it is performed from the sense of duty alone. PMJ, ch. 1. My “Kant, Nonaccidentalness”, op. cit., examines this claim.

²¹ Daniel Wegner’s, The Illusion of Conscious Will (Cambridge, Mass.: MIT Press, 2002), argues for a dubious form of epiphenomenalism about conscious intentions and motives, but it has some fascinating material pertinent to our question. See esp. ch. 5.

²² I take Wilson’s book, Strangers to Ourselves (Cambridge, Mass.: Harvard University Press, 2002), to summarize our current knowledge about unconscious psychological states and processes in general. The quote occurs on p.194.

²³ See Wegner, op. cit., pp. 149-51.

²⁴ There are other interpretations we could give of this kind of case, though. Perhaps hypnotism temporarily induced very strange beliefs and desires in the subject. It is worth mentioning that Wegner reports that there are serious disagreements among experts about what the hypnotic state is, and even whether it exists. Ibid., chap. 8.

²⁵ MetM 241 (446). Cp. 196 (393).

²⁶ CKE 21. Cp. 84.

²⁷ G 66 (398). Translation again emended.

²⁸ Cp. Blum, op. cit., p. 16; Martha Nussbaum, Upheavals of Thought (New York: Cambridge University Press, 2001), ch. 6, esp. pp. 306; 321; 336-7.

²⁹ This idea is related to the so-called ‘veneer theory’ of human nature. Nomy Arpaly in effect takes issue with it. She suggests that an agent can unconsciously respond to morally significant facts. Nomy Arpaly, Unprincipled Virtue (Oxford: Oxford University Press, 2003), pp. 8-10; 75-9.

³⁰ See William Lyons, Emotion (New York: Cambridge University Press, 1980), chs. 2 and 3; Elster, Alchemies, op. cit., pp. 244-283, esp. 247-8; Robert Roberts, Emotions (New York: Cambridge University Press, 2003), pp. 83-106. Theorists like Lyons hold

that having an emotion also involves other things, such as distinctive physiological changes.

³¹ Michael Stocker, "Emotional Thoughts," American Philosophical Quarterly 24 (1987), pp. 59-69. See also Stocker and Hegeman, Valuing Emotions, op. cit., ch. 1; Lyons, op. cit., pp. 57-62. Compare the criticism by the Stoic Posidonius of Chrysippus' thesis that emotions are simply certain kinds of judgments. See Richard Sorabji, Emotion and Peace of Mind (Oxford: Oxford University Press, 2000), ch. 7.

³² Details on request from the author.

³³ Cp. Stocker and Hegeman, Valuing Emotions, op. cit., pp. 158-9.

³⁴ I am here indebted to Martha Satz.

³⁵ Op. cit., pp. 33-4.

³⁶ But even Kant's views do not clearly favor focusing on the deontic significance of consciously chosen reasons. Recall that Kant asserts that it is not clear that any action has moral worth, since we cannot be sure that there are no "secret impulses" of self-love operating. G 74-5 (406-7) This implies that even unconscious motives can affect moral worth. Deontic concepts are not the same as that of the moral worth. Still, Kant would need to explain why unconscious operative motives can affect the application of one sort of moral concept and not the other.

³⁷ The 'roughness' here has to do with delicate conceptual issues that I am not addressing. These concern which factors are relevant to the very existence of a deontic status like that of moral obligation, and which are relevant to judgments of responsibility, praise and blame. Some facts about what S could know (and hence, do) may only be relevant to whether she is blameworthy. So it may be true that S had an obligation to X, but is not to blame for failing to carry it out. For our purposes the crucial point is that sometimes an operative motive that is not affirmed may be relevant to the deontic status of X.

³⁸ My paper "Pure Negligence," American Philosophical Quarterly 30 (1993), pp. 137-49, considered analogous cases like those in Hart, Punishment, op. cit., p. 147.

I argued that the 'could' involved had to be interpreted in a way that requires us to accept incompatibilism. I now think that this claim is mistaken. For an excellent discussion of this question see Michael Smith, Ethics and the A Priori (New York: Cambridge University Press, 2004), ch. 6, pp. 114-35.

³⁹ Hart, Punishment, op. cit., p. 147.

⁴⁰ A well-known discussion of this issue is in Bernard Williams, "Deciding to Believe," in Problems, op. cit., pp. 136-51.