

# 3

## Subjectivism

### 3.1 What is subjectivism?

Subjectivists take moral properties to be reducible to psychological properties and relations. They think that for an object to be good is for some person or group to have (or be disposed to have) some psychological attitude or reaction towards it. Here are a few examples of subjectivist theories; in each case, the '=' sign indicates that the expression on the right explains *what it is* for the expression on the left hand side to be true:<sup>1</sup>

- a)  $x$  is good = The speaker believes that  $x$  is good.
- b)  $x$  is good = The speaker approves of  $x$ .
- c)  $x$  is right = Society approves of  $x$ .
- d)  $x$  is right =  $x$  accords with God's wishes.<sup>2</sup>

Theories (a) and (b) make morality relative to the individual; that is, they imply that two different individuals may correctly call different things 'good'. According to (a), when I say something is good, I mean that *I* believe it is good; when you say something is good, you mean that *you* believe it is good. Theory (c) makes morality relative to a culture; it implies that members of different societies may correctly call different things 'right', though two members of a *single* society may not.

Note that each of these theories is a theory about what it is to *be* good or right, not about what it is to be *taken to be* good or right. Thus, (b) does not hold that for a thing to be *taken to be* good is for

it to be approved of. It holds that for a thing to *be* good is for it to be approved of.

One should not confuse subjectivism with the idea that there are no moral truths or moral facts. Take theory (c): it is definitely true that our society disapproves of theft; therefore, a cultural relativist should say it is *definitely true* that stealing is wrong (in our society). What the relativist denies is that this truth is *objective*, that is, independent of observers in the sense defined in section 1.2. But since there certainly are facts about observers, there is no difficulty in recognizing moral facts.

Subjectivism is a form of reductionism about value. The arguments I deploy in the following chapter serve to refute ethical reductionism in general. Nevertheless, subjectivism is such a common view that it is worth giving it an extra chapter of its own. There are special problems for subjectivism that do not apply to other versions of reductionism. Let us discuss each of the above four versions of subjectivism in turn.

### 3.2 Individualist subjectivism

Begin with theory (a):

- a)  $x$  is good = The speaker believes that  $x$  is good.

Its main problem is one of circularity: to know what it is for a thing to be believed to be  $F$ , we must first know what  $F$  is. A related problem is the infinite regress: if [ $x$  is good] = [I believe  $x$  is good], then [I believe  $x$  is good] = [I believe (I believe  $x$  is good)], and so on. In general, it is incoherent to postulate a proposition whose truth would consist solely in your believing it. The same point applies to the view that to be good is to be believed to be good by society, or by God, or by anyone else, and the same applies if we substitute 'perceived', 'known', 'asserted', or any other verb for 'believed'. The word 'good' should not appear within the explanation of what it is for  $x$  to be good.

This point is worth bearing in mind when we come to theory (b):

- b)  $x$  is good = The speaker approves of  $x$ .

If—as seems plausible—'approving' of a thing is having a moral belief about it, such as the belief that the thing is good, then (b) collapses into the benighted theory (a) with all its attendant incoherence. So the subjectivist will have to say that 'approval' denotes some kind of

attitude or emotion not analyzable in terms of moral beliefs. Even so, theory (b) faces serious problems.<sup>3</sup>

First, imagine we meet a neo-Nazi. He says:

1. I approve of killing Jews.
2. Killing Jews is good.

His first statement is true. But the second obviously is not. But theory (b) implies that both statements mean the same thing, so both are true; indeed, that we are committed by our admission that (1) was spoken truthfully to admitting that (2) was equally truthfully spoken. This seems absurd.

Of course, (b) does not commit *us* to saying 'Killing Jews is good'; since *we* disapprove of such killing, we may truthfully say it is evil. But in doing so, we do not contradict what the neo-Nazi said. We merely report that killing Jews is bad *for us*, but it is nevertheless good for Nazis. Theory (b) thus makes it impossible to disagree with anyone about morality—this is a second absurd consequence. In saying that killing Jews is evil, we obviously *are* intending to contradict, and are contradicting, the neo-Nazi's statement.

Theory (b)'s third absurd consequence is that my moral attitudes are infallible.<sup>4</sup> That is, I would be correct to assert, 'Anything that I approve of is automatically good'. Only a dogmatic egotist would think this, but theory (b) implies that this is necessarily true, since something's being good just consists in my approving of it. If so, I can find out what is good and bad by simply introspecting my own attitudes. Utterances like the following would thus be a mark of confusion: 'I approve of it, but is it really good?' Yet this does *not* seem to be a confused remark; just as I can know that I believe something but still have some doubts about whether it is true, I can know that I approve of something but still have some doubts about whether it is good.

Fourth, consider the question, *why* do I approve of the things I approve of? If there is some reason why I approve of things, then it would seem that that reason, and not the mere fact of my approval, explains why they are good. If I approve of *x* because of some feature *x* has that makes it desirable, admirable, etc., in some respect, then *x*'s desirability (etc.) would be an evaluative fact existing prior to my approval. On the other hand, if I approve of *x* for no reason, or for some reason that does *not* show *x* to be desirable (etc.) in any respect, then my approval is merely arbitrary. And why would someone's *arbitrarily* approving of something render that thing good?

These reasons may explain why few philosophers have embraced theory (b). Let us turn to the more popular theory (c).

### 3.3 Cultural relativism

Theories like (c) are called 'cultural relativism':

- c) *x* is right = Society approves of *x*.

Cultural relativism faces essentially the same problems as individual subjectivism, with minor modifications.

First, imagine that you are a German citizen during World War II. A Nazi officer asserts, 'Killing Jews is right', and most of his and your society agrees.<sup>5</sup> You, however, are a resister: like Oscar Schindler, you are secretly engaged in trying to prevent as many Jews from being killed as possible. According to theory (c), the Nazi officer is correct and you are wrong, for he and not you is acting in accord with the prevailing social norms. Far from being a moral hero, the cultural relativist must say that Schindler was immoral. It is difficult to see the plausibility of this view.

Of course, the relativist may say that Schindler was a hero according to *our* moral standards. But a large part of the *point* of being a cultural relativist—perhaps the central motivation of relativism—has traditionally been that one opposes judging other cultures by our moral standards, or indeed judging other cultures at all. So the fact that the Final Solution was evil according to *our* standards was irrelevant for people living in Nazi society—it gave them no reason for acting in any particular way—and the cultural relativist has no criticism to make of the moral standards that allowed the Nazis to do what they did. Indeed, if our own society's conventions should one day change in such a way as to endorse the mass murder of some minority group, the cultural relativist can only advise us that, on that day, it will be morally right to partake of the bloodshed.

Second, how does theory (c) explain moral disagreements? If one person says 'Abortion is wrong' and another member of the same society says 'Abortion is not wrong', then they are disagreeing with each other. But, according to (c), they are only disagreeing about what the conventions of their society are—they are disputing whether abortion is in fact allowed in their society. It is odd to think that this is all that moral debates are about. Furthermore, if one person says '*x* is right' while a member of *another* society says '*x* is not right', then the cultural relativist must say that they are not disagreeing with each other, since each is only reporting the conventions of his own

society.<sup>6</sup> Even if I say, 'No, I don't just mean that members of *my* society should not do *x*; I mean that people in *your* society should not do *x* either. Your society's practices are wrong'—I still do not thereby disagree with the foreigner who says *x* is right. For in saying the foreign society's practices are 'morally wrong', according to theory (c), all I could mean is that they conflict with my society's rules. But here again the relativist is wrong: when I say that Nazi actions were evil, I obviously am not merely saying that the Nazis had divergent conventions from ours, nor am I saying something unintelligible.

Third, although theory (c) does not imply that any individual is morally infallible, it does imply that *society* is morally infallible. According to (c), whatever society endorses is right, automatically. Thus, everyone should accept such remarks as the following:

If our society approves of abortion, then abortion is right.  
 If our society does not approve of feeding the homeless, then feeding the homeless is not right.  
 If our society approves of torturing babies, then torturing babies is right.

I don't think any of those statements are true, and neither should you. Even if you happen to be pro-choice, I don't think you should endorse the remark that if society approves of abortion, then, automatically, it is right.<sup>7</sup> Our ordinary moral conception does not treat social endorsement as decretory.

Fourth and finally, why should we obey social customs? Either there are *good reasons* for the customs—that is, reasons that show the customs or the behavior they endorse to be good in some way—or there are no such reasons. If there are such reasons, then at least some evaluative facts exist prior to the customs. If there are no such reasons, then the customs are merely arbitrary rules, and why should we obey arbitrary rules?

To illustrate the dilemma, suppose the relativist says that customs exist to promote the stability and survival of social groups. Well, is the stability and survival of social groups *good*? If so, then that is at least one evaluative fact that is independent of customs. But if not, then the existence of a custom so far creates no reason to act in the way the custom requires. One might say that customs create non-moral reasons for actions—perhaps you should obey them in order to avoid punishment by other people. But, unless one posits an obligation to avoid punishment, this would not show why a person is morally obligated to follow social customs, but only why it is

prudent to do so. And one may very well be able to get away with violating customs on some occasions, in which case the present view gives us no reason to act 'morally'.

I now turn to two responses that relativists may make to some of these problems. The first response I call 'the Rigidifying Move'. Traditionally, cultural relativists have been charged with endorsing such statements as,

If society were to approve of eating children, then eating children would be good.

which is clearly false. The Rigidifying Move claims that the word 'good' at the end of the above sentence refers to what society *presently*, in fact, approves of (as opposed to what it would approve of in the hypothetical situation described in the beginning of the sentence).<sup>8</sup> Since society presently does not approve of eating children, the sentence as a whole is false. To illustrate the idea, imagine that I am in Boulder when I utter the following statement: 'If I had flown to Minneapolis today, then Minneapolis would have been here'. This sentence is false. Minneapolis would still have been where it is; it would not have somehow moved to Colorado. 'Here' at the end of the quoted sentence refers to the place where the speaker in fact is, not where he would have been in the hypothetical situation described. According to the Rigidifying Move, the word 'good' works like 'here'. Thus, the relativist is not committed to the absurd statement that if we approved of cannibalism, then cannibalism would be good.

My response is that the relativist is still committed to endorsing:

If society *approves* of eating children, then eating children *is* good.

This slight modification—converting the offending sentence from the subjunctive to the indicative mood—is enough to defeat the Rigidifying Move. Invoking the analogy to the word 'here' again, consider the sentence, 'If I have flown to Minneapolis (such that I am now in Minneapolis), then Minneapolis is here'. This sentence is true, wherever the speaker is when he says it. Thus, the proposal that 'good' works analogously to 'here' offers the relativist no comfort this time.<sup>9</sup>

The second response on behalf of relativism I call the 'Fundamental Values' response, and it applies to the Nazi and Fallibility problems. The Fundamental Values response maintains that *some* socially accepted practices might be morally wrong in virtue of their

failure to cohere with *other*, more fundamental social practices. For instance, suppose that early American society accepted a general norm of respect for persons but that, due to false factual beliefs, it failed to classify black people as persons. In that case, the relativist can argue that slavery was wrong in early America, notwithstanding its social acceptance—it was wrong in virtue of its conflict with the more fundamental norm of respect for persons.<sup>10</sup>

This reply does little to reduce the unpalatability of relativism. The view still implies that a society is infallible in its *fundamental* values. Moreover, there is no reason to assume that past acceptance of slavery was due to factual (non-moral) errors. Perhaps, instead, past acceptance of slavery was due to a *fundamental* antipathy felt towards members of other races, for no reason other than that they belonged to other races. Perhaps arguments about the inferior intellectual abilities of Africans were disingenuous rationalizations for this underlying racist hostility. And perhaps this antipathy was so deep and intense that it would have overridden any other values that conflicted with it. I do not know whether this is the case. But it hardly seems that this being the case would place the institution of slavery in early America beyond reproach. If anything, this would make it even more reprehensible. Yet the Fundamental Values reply entails the reverse. According to the relativist, provided that racist attitudes are sufficiently fundamental, they escape all possible moral criticism.

### 3.4 The divine command theory

We often hear warnings to the effect that the decline of religious belief undermines morality. There is more than one way of interpreting this concern, but here I shall focus on the suggestion that there could be no moral truths if God did not exist—as Dostoyevsky says, if there is no God, everything is permitted.<sup>11</sup>

The Divine Command Theory of ethics holds that an action is morally right if and only if it is of a kind that God commands (or approves of, or wants us to perform).<sup>12</sup> The version of the theory I want to discuss holds a similar view about all other evaluative properties, including goodness, justice, and so on—that is, that all of these properties depend on God in such a way that nothing could have any evaluative property if God did not exist; however, in the following, I shall focus on the property of rightness. Since it takes rightness to be reducible and dependent on the attitudes of an observer (God), the Divine Command Theory is a form of subjectivism. This is worth pointing out, since the theory is often seen as the

arch-nemesis of cultural relativism, whereas in fact the two are variants on the same basic metaethical approach.

Note that the theory is not merely that it is always right to obey God's commands, or that we can find out what is right by consulting God's commands. Those views would be consistent with the idea that right actions are right independently of God but that God (being all-knowing) always knows which actions are right and (being all-good) always approves of them. Rather, the Divine Command Theory holds that right actions are right only *because* God commands them.

Does the divine command theory face the same problems as other forms of subjectivism? The answer is that there are two problems that are peculiar to the Divine Command Theory; on the other hand, the theory avoids at least one problem afflicting secular subjectivism. Consider six potential problems.

*First*, does God exist? If there is no God, the divine command theory is unattractive. Other subjectivist theories avoid this problem, since there is no doubt that human individuals and societies exist. There are serious grounds for doubting God's existence; however, they are beyond the scope of this book.

*Second*, how do we know what God approves of? Again, this problem has no analog for secular forms of subjectivism, since it is easy to know what human beings approve of. According to Western religious tradition, we can know God's desires from such sacred texts as the Bible or the Koran. But in fact, this source of guidance is notoriously unreliable. The moral guidance to be found in the Bible includes the following imperatives (my paraphrases):

Kill anyone who curses their parents. (*Leviticus* 20:9)

Kill anyone who commits adultery. (*Leviticus* 20:10)

Kill homosexuals. (*Leviticus* 20:13)

Kill women who have premarital sex. (*Deuteronomy* 22:20-1)

Kill people who work on the Sabbath. (*Exodus* 35:2)

Make war on the people occupying the promised land. Show no mercy; kill every man, woman, and child. (*Deuteronomy* 7:1-2; 20:16-17)

Slavery is okay. (*Leviticus* 25:44-5)

It is okay to beat your slaves, as long as they don't die. (*Exodus* 21:20-1)

Not to be provincial, I note that the Koran advises us:

Make war on and kill those who don't accept Islam. (*Sura* 9.5, 9.29-31)

Anyone who dies fighting for Allah will go to heaven. (47.4-6)  
 Unbelievers are the vilest of all creatures. (8.55)  
 They will suffer gruesome torture in hell, forever. (22.19-22)  
 Appropriate punishments for opposing Islam include execution,  
 crucifixion, the cutting off of hands and feet, and exile. (5.33)  
 Cut off the hands of thieves. (5.38)

Few today, even among believers, would defend the above ideas and precepts. An obvious alternative for some theists would be to hold that although there is a God, He did not inspire the Bible or the Koran. The problem is that this would leave it mysterious how we can know what God wants, thereby threatening our capacity for moral knowledge. But it seems clear that we have some moral knowledge, as nearly any theist would agree.

Another alternative, for a Christian or Jew, would be to hold that the Bible contains some errors introduced by human transmitters of God's word, who may have misunderstood parts of God's message, inserted some of their own prejudices, or made mistakes in transmitting the scriptures down the generations. This would enable one consistently to reject the precepts I have listed above while maintaining that the Bible contains much that is true, important, and inspired by God. I cannot refute this position (not here, at any rate), nor do I claim that there is nothing of value to be found in the Bible or the Koran. But notice that if one takes this position, one cannot simply appeal to the Bible to determine what is right and wrong. The Divine Command theorist would have to posit some sort of access to moral facts independent of our knowledge of God's will,<sup>13</sup> which would seem to surrender one of the major advantages of a Divine Command theory—namely, its ability to explain moral knowledge. For example, the Divine Command theorist might be driven to posit a faculty of moral intuition—but then it is unclear how his theory would be better than traditional ethical intuitionism. Or the theorist might adopt one of the proposals of recent ethical naturalists (see the following chapter, where these proposals are also criticized).

*Third*, the problem of Nazis does not afflict the Divine Command theorist as it afflicts other subjectivists—the divine command theorist need not grant the truth of Nazi moral assertions, since the Nazis did not act in accord with God's will. But a similar problem arises nonetheless. The divine command theorist must hold that if God were to command us, for no particular reason, to torture and murder our children, then such actions would be moral.

A defender of the theory might observe that, since God is morally perfect, he would never command such things.<sup>14</sup> But this does not

engage the objection. Even if we know that someone would not do *A*, we can still ask what would follow *if he were* to do it. For instance, I am not going to drink an entire bottle of Everclear. However, we can still discuss what would follow if I were to do so; for instance, if I drank a bottle of Everclear, I should immediately call the hospital. Likewise, we can talk about what would follow if God were to command us to murder children, even if we are sure he will not do so. Furthermore, the reasoning for the claim that God would never make such a command is fallacious. According to the divine command theory, whatever God commands is right, so if God were to command murder, murder would be right. Therefore, God's moral perfection is no obstacle to his issuing such a command: a morally perfect being is not debarred from commanding people to do things that are morally correct as of the time he commands them.<sup>15</sup>

Robert Adams proposes another response to the problem. He says that the obligation to obey God's commands is contingent on there being a *loving* God. If God were to command the killing of children, this would show that He was not loving as Christians believe him in fact to be. Since the rightness of an action consists in its conforming to the will of a loving God, child killing would *not* be right in that situation.<sup>16</sup> But nor, on Adams' view, would it be wrong. Since on Adams' view, an act's being commanded (forbidden) by a loving God is a necessary condition on its being right (wrong), Adams must say that if God were not loving, if He refrained from forbidding the killing of children, or if He simply did not exist, then there would be *nothing wrong* with killing children. This is hardly more plausible than the position that God's commanding child murder would make it right.

*Fourth*, the divine command theory is less subject to the problem of disagreement than other forms of subjectivism are, because it is not absurd to suppose that people disagreeing about ethical questions are sometimes disagreeing about what God would approve of. It seems odd, nonetheless, to maintain that *everyone* who argues about moral questions, including atheists, agnostics, and polytheists, is arguing about what God would approve of.<sup>17</sup>

*Fifth*, there is one problem afflicting other forms of subjectivism that does *not* apply to the divine command theory. This is the problem that subjectivists must hold some individual or group to be morally infallible, which is implausible. But in the case of God, it is not implausible—indeed, it is true by definition—that he is infallible.

*Sixth*, and most importantly, even if there is a God, why should we do what He says? Someone's telling you to do something does not in general create a moral obligation to do it. If your next door neighbor

tells you to kill your son,<sup>18</sup> this creates no obligation at all, even *prima facie*, for you to do so. If Satan tells you to avoid eating pork, this creates no obligation to obey. There must be something special about God, if *His* commands are to create moral obligations. What is it?

I argue that the divine command theorist has no satisfactory answer to this question. In outline:

1. If no characteristics of God ground an obligation to obey God's commands, then there is no obligation to obey God's commands.
2. The morally neutral characteristics of God do not ground an obligation to obey God's commands.
3. If the morally significant characteristics of God ground an obligation to obey his commands, then some moral facts are independent of God's commands and attitudes.
4. If either (a) there is no obligation to obey God's commands or (b) some moral facts are independent of God's commands and attitudes, then the divine command theory is false.
5. Therefore, the divine command theory is false.

Premise (1) has already been motivated. If there is nothing special about God that sets him apart from most other beings, explaining why we must obey him in particular, then, since we are not *in general* obligated to obey other beings, we are not obligated to obey God.

In premises (2) and (3), a 'morally neutral' characteristic is one that it is possible to have without having an evaluative property; a morally significant characteristic is one that one cannot have without having an evaluative property. For instance, being powerful and being the creator of the Earth are morally neutral characteristics. Being just and being cruel are morally significant characteristics, since one cannot be just without being good in some respect, nor can one be cruel without being bad in some respect.

(2) can be supported by hypothetical examples. For any morally neutral characteristic of God, it is possible to imagine an evil being—Satan, say—having that characteristic. In such a situation, intuitively, there would be no moral obligation to follow Satan's commands. For example, it might be said that we should obey God because He is our creator. But imagine that you found out that you had actually been specially created by Satan. In that case, would you be under a general obligation to follow Satan's will? It is also said that God has unlimited power over us, and can bestow eternal punishments and rewards. Imagine that Satan had unlimited power over you, and could bestow eternal punishments and rewards. In that case, no doubt it would be *prudent* to obey Satan's commands. But would it be *morally obligatory*

to do so; would you be *immoral* if you were to resist Satan's evil plans? Lastly, suppose Satan were all-knowing and all-powerful (but still thoroughly evil). Would it then be morally obligatory to follow his will? In all of these cases, the answer is surely that you would not be obligated to follow the will of Satan. Therefore, none of these properties—being our creator, having the ability to reward and punish us, being all-knowing, being all-powerful, or any combination of these—can ground an obligation of obedience on our part towards another being.

There is one crucial respect, or family of respects, in which God differs from Satan in my hypothetical scenarios. This is that Satan is evil, whereas God is purely good, just, benevolent, and so forth. It is plausible that if a *supremely good* being tells you to do something, then you are thereupon obligated to do it. But here is where premise (3) comes in. If the reason we must obey God turns on God's goodness, then some moral facts must exist prior to God's commands. By hypothesis, the capacity of God's commands to generate moral truths depends upon God's goodness; on pain of circularity, therefore, it cannot also be the case that God's goodness is generated by his commands. (It cannot be, for example, that God is good because he approves of himself, or because he commands us to worship him. Imagine Satan likewise approving of himself and commanding us to worship *him*.) God must *already* be good, so there must already be a standard of value in place—for example, that lovingness, mercy, and justice are good—prior to God's commands.

Apropos of this, recall Robert Adams' solution to the problem of God's possible command to kill children. Adams finds God's lovingness relevant to the obligation to obey Him and argues that a loving God could not order such killing. But what motivates Adams to single out *lovingness*? Is it not that he considers this trait a virtue, or in some way *good*?<sup>19</sup> It is hard to imagine someone who placed no value on love giving Adams' reply. Thus, it seems that the ability of a loving God's commands to create obligations depends upon the independently assumed *value* of love.

Finally, premise (4) is true because the divine command theory holds that what is morally obligatory is so in virtue of God's commanding it, and that something similar holds for all other moral properties (such as goodness and wrongness). This implies both that it is obligatory to obey God's commands, and that nothing has any moral property independently of God's commands.

The conclusion (5) follows from premises (1)–(4): God's non-moral characteristics cannot ground an obligation of obedience. God's moral characteristics can ground such an obligation only if there are

moral facts independent of God's commands. So either there is no reason to obey God, or there are moral facts independent of God's commands. Either of these alternatives conflicts with the divine command theory. So that theory must be rejected.

A closely related point can be made more succinctly. Why does God command what he does? If God has no moral reasons for his commands, then they are merely arbitrary—and why should we obey arbitrary commands? But if God *has* moral reasons for his commands, then some moral truths must exist independently of his commands. Either way, the divine command theory is false.<sup>20</sup>

One response to this is that God's commands may escape arbitrariness by virtue of God's having *non-moral* reasons for his commands.<sup>21</sup> I don't think this helps, any more than the suggestion that God has *bad* reasons for his commands would help. How would God's having a reason *with no moral weight* for a command serve to render that command non-arbitrary in a morally relevant sense, that is, in such a way that we morally ought to follow it?

### \*3.5 The ideal observer theory

Ideal Observer theories say that a thing is good, or an action right, if an 'ideal observer' would take some pro-attitude towards it.<sup>22</sup> This is similar to the divine command theory (God being the ultimate in ideal observers), but it does not require any godlike being to actually exist; instead, we are to answer moral questions by reasoning hypothetically about what the ideal observer *would* favor, if there were one. I classify this as a form of subjectivism, since it makes value depend constitutively upon a (hypothetical) observer's attitudes.<sup>23</sup>

This statement of the theory raises three questions. First, what sort of 'pro-attitude' is the ideal observer supposed to take? Second, what is meant by an 'ideal observer'? And third, is the good what *all* ideal observers would necessarily favor, what *an* ideal observer might favor, or what?

I won't dwell on the first question. The pro-attitude might be desire, approval, or something else. In the following, I shall speak in terms of desire, but analogous arguments can be made regardless of which pro-attitude the ideal observer theorist makes use of.

In response to the second question, here are some traits that have been attributed to the 'ideal observer':

He knows all the non-moral facts.  
He imagines all relevant states of affairs perfectly vividly.  
He is perfectly rational.

He is disinterested and dispassionate; that is, he lacks desires and emotions that refer to specific individuals.

He is benevolent.

He is in other respects normal (for instance, he's not crazy).<sup>24</sup>

Different philosophers may pick slightly different collections of traits, leading to different versions of ideal observer theory. We needn't worry about this. However, there are two important points about what the ideal observer can *not* be taken to be: The ideal observer should not be assigned any *moral* properties; for instance, he is not to be understood as a perfectly virtuous being. This would land the theory in circularity, given its aim to explain all moral properties. Also, though the ideal observer may be granted full awareness of non-moral facts, he should not be ascribed *moral* knowledge, or any traits entailing a capacity for moral knowledge (such as a faculty of ethical intuition). Again, the concern would be one of circularity; such a capacity would presuppose the existence of moral facts independent of the observer.

In response to our third question, there are two kinds of ideal observer theory. One takes the good to be that which every ideal observer would necessarily approve of.<sup>25</sup> The problem is that this seems to result in nothing's being good. The traits listed above (apart from benevolence) do not entail that the observer so described has any desires at all, let alone feelings of approval or disapproval. We might appeal to the 'normalcy' condition to deal with this—a being without desires would be abnormal. Perhaps there are some desires that all normal people must have. Now the problem is that the definition of the ideal observer seems to illicitly include an evaluative term: when 'normal' is used to mean 'has normal desires', it is evaluative. For example, people who think homosexuality is wrong refuse to call homosexual desires 'normal'; those who consider it permissible call those desires 'normal'. We call pedophilia 'abnormal' because it is bad. The defender of the ideal observer theory would probably want to exclude observers with sadistic desires; if this is done on the basis of such desires' 'abnormality', this would seem to be because 'abnormal' functions evaluatively. A similar point would apply to attempting to use the 'rationality' requirement to rule out apathetic ideal observers.

We might try simply stipulating what desires the ideal observer has; for instance, that he is benevolent (desires the welfare of people in general).<sup>26</sup> One problem with this is that, in stipulating the ideal observer's motivations, we would simply be stipulating the content of moral theory. If we say, 'Ideal observers desire *x*, *y*, and *z*', then we

are in effect stipulating that  $x$ ,  $y$ , and  $z$  are the sole goods. This is because nothing *else* in the definition of the ideal observer entails any motivations, so the only things that *all* ideal observers would desire would be  $x$ ,  $y$ , and  $z$ . But then we might as well take the simpler theory, 'For a thing to be good is for it to be  $x$ ,  $y$ , or  $z$ ', and call ourselves ethical naturalists. No work is done by adding the rest of the ideal-observer apparatus.

Richard Brandt offers a better approach to our third question: he relativizes value to the individual, and says that what is good *for* a given person is what *that person* would desire, if *he* were fully rational, fully informed, and so on.<sup>27</sup> This enables us to start from our actual desires, and merely consider what changes would be made if we were to become fully rational. Presumably, my becoming fully rational would not require my losing all my desires, so some things will be good (relative to me) on this theory.

All of that was by way of identifying what I take to be the best version of ideal observer theory. Now for what is wrong with it.

The problems are familiar. First, the *problem of horrible desires*: Suppose I want to see lots of death and destruction for its own sake. Vividly imagining the destruction only increases my desire for it. The desire was not caused by any false beliefs of mine, nor have I been brainwashed into it, etc.; I'm just a nasty person.<sup>28</sup> No further non-moral information would change my mind about this. According to the Ideal Observer Theory, it seems, death and destruction is good (for me), and I should aim at it. Brandt might respond that no actual person is like that. I see no reason for assuming that, but we need not debate that. Here is a more realistic example: there is an individual who has a basic aversion to people of dark-skinned races. If given full non-moral information, but no moral arguments or intuitions, he would continue to feel this aversion. The aversion, let's suppose, is a genetically programmed hostility that some people bear for people of other races and social groups. In this situation, the ideal observer theory would apparently counsel that this person should favor racist actions and policies, perhaps even slavery for the members of the despised race.

Brandt might claim that actual racists are not like this, that racist attitudes are in fact based on false non-moral beliefs about other races. I see little reason to believe this; at least equally plausible is that racist beliefs about the personality traits and intellectual and other abilities of minority races are accepted as rationalizations for preexisting racist feelings. Regardless, the ideal observer theory is committed to holding that *if* racist attitudes are fundamental and not caused by factual errors, brainwashing, and so on, then racist policies

are good (for the racists), and people with those attitudes *should* seek the unequal treatment of minorities.

Second, the *Problem of Disagreement*: On Brandt's version of the ideal observer theory, people who appear to be disagreeing about morals are not really disagreeing. This is true regardless of whether most people are similar enough that the ideally rational versions of them would desire the same things. It is true in virtue of the logical form of moral statements according to the theory: each person is talking about what *he* would desire if fully informed, etc. So 'x is good' and 'x is not good' do not contradict each other, if said by different people.

Third, the *Problem of Fallibility*: the ideal observer theory can explain how my actual preferences can be bad, since they may fail to match those I would have were I fully informed, etc. What it cannot accommodate is the fact that even a fully informed person may be morally fallible—correctness about all *non*-moral matters does not, conceptually, rule out moral error. If there were an ideal observer, he could intelligibly ask himself, 'Are the things I want really good?'

Fourth and finally, the *Problem of Arbitrariness*: Why does the ideal observer desire what he does? If he desires it because it is good, or for any other moral reason, then there are moral facts independent of his attitudes. But if he has no moral reason for desiring what he does, then his desires are arbitrary and have no moral import. Why would someone's desiring or approving of something *for no reason* make that thing good? This is essentially an appeal to the ethical intuition that *being desired* does not convert a morally neutral or bad object into a good one (not to be confused with the claim that *desire-satisfaction* isn't good).

### 3.6 The subjectivist fallacy

Why has anyone held a subjective theory of value? Individual subjectivists and cultural relativists commonly attempt to support their views with one or both of the following kinds of evidence: (i) evidence that moral beliefs vary among individuals or cultures; (ii) evidence that moral beliefs are influenced by emotions, peer pressure, and other non-rational causes. The argument seems to be:

1. Different people/cultures have many different sets of moral beliefs.
2. Therefore, there are many different moral truths (many different moral perspectives are equally correct).

Or:



3. People's moral beliefs are generally caused by emotions and/or their society's conventions.
4. Therefore, a moral belief is just a belief about emotions and/or conventions.

It is hard to believe that anyone needs to have the fallacies in the above pointed out to them. The first argument is invalid as it stands; for (1) to support (2), one would have to add the premise, 'Whatever people believe is true'. That premise is false, and well known to be false. The second argument is equally invalid; for (3) to support (4), one would have to add the premise, 'People's beliefs are always caused by what they are about'. That is also well known to be false. People are often wrong, and they often adopt beliefs for irrational reasons.

Arguments of this kind deserve the label 'the Subjectivist Fallacy', since they are virtually the only arguments given by subjectivists, and since they are given by nearly all subjectivists.<sup>29</sup> The underlying mistake is simple: it consists in confusing representations with reality. The first argument above overtly confuses *beliefs* with *truths*. The second confuses the way people form a belief with what the belief is about.

To illustrate, suppose three gamblers at the racetrack all think, for purely emotional reasons, that a different horse is going to win. It does not follow that the race will have three different outcomes, one in which each gambler is right. Nor does it follow that to say a horse is going to win the race is just to say that one has a certain emotion. Yet that is the sort of reasoning involved in the Subjectivist Fallacy.

In sum, the most commonly held forms of subjectivism are motivated by obvious confusions. Furthermore, all ethical subjectivists face at least three major problems:

- i) *The Problem of Horrible Attitudes*: Subjectivists must hold that seemingly horrible actions, such as murder and torture, are morally right as long as the appropriate person or group endorses them.
- ii) *The Problem of Disagreement*: Subjectivists cannot plausibly account for moral disagreement.
- iii) *The Problem of Arbitrariness*: Subjectivists must hold that things are made good, or actions made obligatory, by merely arbitrary rules or attitudes.

In addition, individual subjectivists, cultural relativists, and ideal observer theorists face:

- iv) *The Problem of Fallibility*: They cannot explain the intuition that the relevant person or group could have morally incorrect attitudes.

And divine command theorists face:

- v) *The Problem of God's Existence*: It is doubtful whether God exists.
- vi) *The Problem of Knowing God's Wishes*: Even if He exists, it is unclear what God wants.

These reasons seem more than sufficient for rejecting subjectivism. But, even if I am wrong about all of these problems, subjectivists will still face the problems raised for all reductionists in the following chapter.