Chapter 16
Subjectivism About Value

Suppose you think that you can spare Sally's feelings by lying to her about Roger. Not lying to her about Roger means that Sally will be terribly hurt.

Should you lie to Sally about Roger to spare Sally's feelings? You have reasons to lie to her and reasons not to lie to her. Lying to her would, you believe, spare her feelings. You are confident that Sally will not discover your lie, or at least won't discover it until the crisis has passed when it won't much matter to her. Chances are, then, you reason, that lying to her will be all for the best. And yet lying to her, you understand, is to be disrespectful toward her.

So, *should* you lie to Sally? Whether you in fact do lie to her might well depend on what you value more–her feelings or that she be treated with respect. If you value her feelings over that she be treated with respect, you will lie to her; if you value that she be treated with respect over her feelings, you will not lie to her. But all that's beside the point. Our question isn't whether you *will* lie to her. It is, rather, whether you *should* lie to her.

To answer the question should you lie to Sally about Roger, we would need to know whether you should value Sally's feelings above that Sally is treated respectfully. What this question is asking and how it is to be answered is the topic of the current Part of this book. Our concern is to understand the nature of value and the possibility of evaluative judgements of the value of things.

One answer to the question whether you should value Sally's feelings over that Sally be treated with respect is that it depends on which of the two, feelings or respect, is more valuable. If Sally's feelings are more important than that Sally be treated with respect, then you should lie to her to spare her feelings, even though lying to her is to treat her disrespectfully. But if it is better that she be treated with respect, then you shouldn't lie to her, even at the cost of her feelings. The idea here is that states of affairs have value or importance within themselves and we act well only when we value things according to the value they themselves have. This idea will be the focus of the next chapter.

Another answer is that what you should do depends on the quality of the various reasons you have for acting one way or another. You should do what the highest quality reason you possess is a reason for doing. This idea–that practical reasoning itself enjoins us to place certain values above others–will be the focus of the chapter after the next one.

A third answer to the question whether you should value Sally's feelings over that Sally be treated respectfully is that it depends on the mores embedded in the folkways of one's culture. How is it within one's culture with regard to feelings and respect? If the cultural traditions within which you live are more on the side of sparing feelings than on the side that people are treated with respect, then you should lie to spare Sally's feelings; if they are more on the side that people are treated with respect, then you shouldn't lie to her. The idea that what you should value and do depends on cultural norms is the focus of the fourth chapter of this Part.

A fourth response to our question isn't really an *answer* to it at all. This fourth response is to say that questions about what one should do–what one should value, what one should value more or less than another thing–can sensibly be asked only against the background of the values of the person asking the question. The idea is that nothing has value in itself–either objectively or relatively to a culture's mores–; there's just the activity or state of valuing things. Your sparing Sally's feelings is neither better nor worse, in itself, than your treating Sally respectfully. When you consider whether you should do the one rather than the other, you are simply considering which of the two you care about most. If you pause to ask whether you really should care, say, about Sally's feelings to the extent that you do you are asking simply how caring for them to that extent comports with all the things you like and dislike, love and hate, and so on. When I ask whether you should care about Sally's feelings to the extent you do, I am asking simply how your caring for them to that extent comports with *my* likes and dislikes, *my* hopes and fears–for myself, for you, for Sally, for anything else that might cross my mind. This response is the focus of this chapter.

***

No one likes to take the bus, and yet people all the time are taking the bus. No one likes taking the bus but taking the bus is often an effective and inexpensive and pleasant enough way of getting where one wants to be. Where one wants to be is at the pool swimming, at a restaurant with friends, in the library reading and thinking about the nature of value. These things–swimming, socializing, thinking–*are* things that one can like and enjoy for their own sakes. Taking the bus is a means to an end; the end for which taking the bus is a means might be an end in itself for you, something that matters to you directly, something that you value for its own sake. We can say that you value taking the bus instrumentally, while swimming, which you like to do for its own sake, is something you value intrinsically.

Our question is how we can and should think–and think critically–about those things, whatever they happen to be, that we value intrinsically. Should you care as much about swimming as you do? Maybe you should care more about it; maybe you shouldn't care about it at all. Let's suppose you presently don't enjoy viewing abstract painting. Works by the likes of Rothko and de Kooning leave you cold. Perhaps you should, though, acquire a taste for them. To return to ethics: Sally is moved to act by her desire that people be in good spirits and by her fear of their being upset or hurt. These things matter to her for their own sakes and matter to her greatly. Sally will opt to lie to a person to spare that person's feelings whenever that person's feelings are at risk and lying will set things right. But should Sally care so much for people's feelings? After all, by caring as much as she does, Sally is prone to treat people disrespectfully. Perhaps she should care more that people are treated with respect.

We are supposing, though, in this chapter, that while we value some things intrinsically–for their own sakes, that is–the things we value intrinsically are not themselves intrinsically valuable. We are supposing that nothing possesses value as one of its properties. Things are red or blue, heavy or light, long lasting or fleeting, salty or bitter, exciting or dull, painful or pleasant, heathful or unhealthful in themselves or through their effects–but nothing itself is good or bad, or right or wrong, or better or worse than anything else.

Our question is whether we can and how we should determine what to value intrinsically–what to care about for its own sake. Given our supposition that nothing is

intrinsically valuable, we can't answer our question by saying that we should value intrinsically that or only that which is intrinsically valuable.

One suggestion is that our basic tastes and our deep attachments are beyond evaluation. We neither should nor shouldn't care as much as we do about people's feelings; we feel the way we do and that is that. To ask whether we should care as we do is to suppose that we have available to us standards by which to measure that which we care about. But the tastes and values we are now discussing are basic–we have no deeper or more basic tastes or values that could serve as standards by which to evaluate them. These tastes and values can be used as standards to evaluate other of our tastes and values, but they themselves lie beneath all other tastes and standards and, so, cannot themselves be criticised or commended.

The basic idea in this argument why some of the things we value intrinsically are not themselves subject to evaluation is sound. That basic idea is that all evaluation requires standards. If we are to evaluate something, we need to do so against something else. Where there is nothing else, there can be no evaluation. The basic idea is sound. What is false is the claim that when it comes to something we value intrinsically, there is nothing available to us by which to evaluate whether we should value it intrinsically. If there is any problem with regard to determining whether we should care deeply about something about which we do care deeply, it is not that we have no standards available to us by which to evaluate it. The problem, rather, if there is one, is that we have an excess of standards by which to evaluate it.

We can evaluate our caring intrinsically about something–about swimming, about justice, about children, about our children–against anything else we care about at all. Sally can evaluate her concern for people's feelings against her concern that people be treated with respect. Should she do so, she will find that her concern for feelings leads her sometimes to treat people disrespectfully, to lie to them, for instance, in order to spare their feelings. Coming to this result might give Sally a reason to modify the degree to which she is concerned for people's feelings. But it needn't. Sally might accept that she will sometimes treat people disrespectfully as the cost of caring about people's feelings to the degree that she does.

Now it might be that for most people most of the time, life goes well to the extent that the various things they value intrinsically don't get in each other's way. A person whose deep loves and likes are in harmony with each other will be able to serve his loves more successfully than will a person who must often sacrifice one thing she loves in order to pursue something else she loves. Likely enough, a person who encounters less frustration in life is happier than one who encounters more. Most of us, certainly, seek to be happy and none of us wants to be miserable. This might incline us to conclude that a person's deepest concern is to be happy. Should we draw the conclusion that we all seek happiness over everything else, we will think that we have available a final standard by which to evaluate any deep care or concern a person might have. We can ask whether caring for that thing in the way that the person does (in her material circumstances and given reasonable expectations for the future) conduces to her happiness overall or tends to frustrate it.

We must be careful when we draw these inferences. Even though we all seek happiness and none of us wishes to be miserable, we might not all want happiness above all other things. Some of us might accept–and quite properly accept–unhappiness as the cost of serving those things that matter the most to us.

We've been asking whether it is possible for a person to evaluate critically her loves and hates, likes and dislikes, with an eye toward finding reasons for altering them in one way or another or instead maintaining them as they are. We've asked this question while supposing that nothing in the world is in itself or intrinsically lovable or hateful, likable or despicable. We've discovered that indeed it is possible for a person to evaluate critically that which she values intrinsically. A person can ask of anything she values whether her valuing it as she does hinders her pursuit or enjoyment of anything else she values.

A third person, an observer, can, as well, evaluate critically a person's values just as that first person can herself. A third person's findings, though, can have significance for the first person only if they touch one or another of the first person's own desires or emotions. Roger might conclude after study that Sally should be kinder or more generous, for instance, but his conclusion expresses nothing more than his wish that Sally be kinder or more generous unless it is possible for Sally to see, from her perspective, some point to being kinder or more generous.

Let us now consider more fully the supposition we have made, that nothing possesses value intrinsically. What, exactly, have we been supposing? We've been supposing that the things of the world–objects, states of affairs, events, institutions, practices, actions–have no value properties. We can call this idea subjectivism about value. Subjectivism about value includes the contention that the things of the world possess no features or qualities in virtue of which they are good or bad, or right or wrong. If "good" and "bad" are taken to name features or qualities that might exist along with the other features or qualities of things, then, since there are no such features or qualities, sentences attributing goodness or badness to states of affairs must all be false, reasons the subjectivist. "Teasing children for fun is wrong" is false, for it attributes a feature, wrongness, to instances of a type of practice, teasing children for fun, while instances of that practice lack that feature. "Teasing children for fun is sometimes all right" is also false, for it attributes a feature, being all right ethically, to some instances of a practice while no instances of that practice possess that feature. "Justice is good" is false, and so is "Justice is bad."

Consider now not ethical value but other types of value. We laugh at certain jokes, but do we laugh at them because they are funny? No, says the subjectivist. It is not some quality, funniness, inherent in the joke that causes us to laugh. We laugh because of the irony or the incongruity or whatever. The joke itself possesses no quality of being funny. The joke tickles us and we laugh. In calling the joke funny we express our amusement.

Andrea and Boris often laugh at the same jokes but sometimes Andrea laughs at a joke that leaves Boris cold and sometimes Boris laughs at a joke that leaves Andrea cold. Suppose Andrea and Boris together hear a joke and Andrea laughs while Boris doesn't. Has one or the other made a mistake? Has Boris failed to perceive something in the joke that's actually there? Has Andrea seen in the joke something that's not actually there, as one might see a pool of water on the horizon over the sand? No, says the subjectivist about value. So long as both Andrea and Boris understand the joke and see how it is supposed to work, Andrea is not mistaking some quality of the joke for funniness and laughing in error, nor is Boris failing to perceive the funniness that really is there in the joke. It is just that different things amuse them.

Suppose that Boris, confiding in Sally, says, "Andrea is very sexy." Sally says, "No, she isn't." Do Boris and Sally disagree about Andrea? No, according to the subjectivist about value, they don't. They agree about everything there is to agree about–the colour of Andrea's hair, the shape of her nose, the way that she walks. Sexiness, says the subjectivist, is not a quality that Andrea might have in addition to her other qualities. Nor is it a quality that might emerge from those other qualities or supervene on them. Boris is attracted sexually to Andrea, Sally isn't.

The subjectivist about value says similar things about beauty, comfort, being delicious–nothing is either beautiful or ugly, nothing is either comfortable or uncomfortable, nothing is either delicious or terrible. A thing might be such that some of us–or, indeed, all of us–are captivated by it, or sit comfortably in it, or love its taste. But things do not possess value properties, neither alongside their actual properties nor supervening upon them.

Back to ethical value, then. Taken as a purported description, "love is better than hate" is false, as are both "hate is better than love" and "love and hate are of equal value." But a person who says "love is better than hate," though she fails to describe anything accurately and, for that reason, has said something false, might be doing, or seeking to do, a whole range of things. She is at least, for one, expressing her approval of loving over hating. Another way in which she could express he approval would be to say "Yes to love and no to hate." She might also be expressing a commitment she has to loving over hating. "To love rather than to hate–that is my commitment." ("Boldly to go were no man has gone before.") She might be issuing a command or a recommendation or a suggestion. "You (me, all of us)–*love* people rather than hating them!" Or she might be expressing a wish or a hope. "Would that we were to love rather than to hate." There are at least four different things she could be doing by speaking that false sentence, at least if she is sincere in speaking it. She is at least expressing her approval of love over hate. But she might also be expressing one or more of a commitment, command, or wish. For the subjectivist, what she isn't doing is describing accurately how something in the world stands.

A person who thinks subjectivism about value is true will not suppose that she is attempting to get right how things in the world stand when she says "It's wrong to tease children for fun" or "helping people is always a good thing to do." She will use such expressions, if she uses them at all, to express her approval and, perhaps, also to express her commitments, to issue orders, or to express wishes–but not to describe how things are.

Each of us values many things intrinsically. That is, we like or love or otherwise have a positive affective attitude toward them for their own sakes. We might love honesty for its own sake, or the music of Miles Davis, or carpentry, or children, or our children. Though we value things intrinsically, says the subjectivist about value, the things we value intrinsically do not themselves possess value, for there is no such property as value.

Subjectivists about value have at least three different arguments to the conclusion that there are no value-properties. One is the argument from the fact that people differ greatly in what they value. Another is the argument from the idea that we need not posit intrinsic value in order to explain anything. And another is the argument from the fact that the value judgements we accept tend to move us to action.

1) The metaphysical argument. Take two judgements plainly regarding a matter of fact that are inconsistent with each other. "A female ancestor shared by all contemporary

humans lived no more than 125,000 year ago" could be one (the "Out-of-Africa" hypothesis); "The most recent common female ancestor of contemporary humans lived long, long ago in the past, at least 300,000 years ago, maybe 500,000 years ago, but certainly more than 125,000 year ago" could be the other (the "Regional Diffusion" hypothesis). These two hypotheses are currently subjects of investigation and debate among anthropologists and biologists. Both hypotheses cannot be true. Which, then, if either, is true? Each side has what it takes to be evidence in favour of its view and thinks it can explain away at least some of what the other side takes to be evidence in favour of its view. What the two sides take to be evidence consists in observations that can be made right now, primarily observations about the genetic similarities and differences among people around the world and observations concerning old bones. Now it might be the case that the matter will never be settled. Perhaps the evidence is just too scattered and thin these days, as at least tens of thousands of years have passed since the events that make one of the hypotheses true and the other false. Still, researchers understand more or less what would count as relevant observations and are working hard to put themselves into places where they can make relevant observations.

The idea here is that events leave traces. When something happens, the world is different in virtue of its happening from how it would have been had that something not happened. That is, the occurrence of particular event A changes the world not just in that had A not occurred, the world would be different in lacking A. The occurrence of A changes the world in that what happens after A and next after that is different from what would have happened had A not occurred. We can put this point this way: Were any one fact different, lots of other facts would be different as well.

We can conceive of what the world would be like were the Out-of-Africa hypothesis true, and we can conceive of what the world would be like were the Regional Diffusion hypothesis true. Armed with these conceptions, we observe our world to see whether it is like the world in which the Out-of-Africa hypothesis is true or whether it is like the world in which the Regional Diffusion hypothesis is true. There is disagreement and controversy in anthropology and biology regarding these hypotheses, but researchers have ideas about what traces would be left and they busy themselves trying to find them.

Now take two judgements about a matter of value that are inconsistent with each other. "Equality is better than freedom," say, and "Freedom is better than equality." We can rewrite these judgements as "When freedom threatens equality, one should act to preserve equality even at the cost of rendering people less free" and "When equality threatens freedom, one should act to preserve freedom even at the cost of rendering people less equal." Or take the judgements we encountered earlier, that you should lie to Sally about Roger rather than be honest with her and that you should be honest with Sally about Roger rather than lie to her.

The conflicting judgements here cannot both be true. For each pair, if one of the pair is true, then the other is false. How are we to seek to determine which of each pair is true, if either is true? By collecting evidence for and against each and then weighing that evidence, of course, just as people investigating the origins of anatomically modern humans are currently doing with regard to the Out-of-Africa and the Regional Diffusion hypotheses. We ask what the world would be like were the first judgement true and what the world would be like were the second judgement true. The world will look different in some ways, of course, should the

one judgement be true from how it will look should the other judgement be true. No fact could be or could have been different than how it is or was without other facts being different.

Now it is possible, we should be aware, just as it is possible in the case of the origins of anatomically modern humans, that evidence has dissipated, or that taking an observation as evidence one way or the other requires reading it through a controversial theory or principle. So we cannot be confident that our quest to determine which of the two is true will succeed. We should be prepared to find our investigation either inconclusive or interminable, just as scientists always are. Well, okay, we're not confident we will succeed and we're prepared to leave the matter open if it cannot be settled on the evidence. Let's get going! Let's generate some ideas about how the world would look should the one judgement be true and how the world would look should the other judgement be true, and then let's compare our ideas with what we see in this world.

Hmmm. How *would* the world differ were equality better than freedom rather than were freedom better than equality? What should we expect to find different in the world if you should lie to Sally about Roger than if you should be honest to Sally about Roger? The subjectivist about value says in answer to the first question "the world would differ in no way at all," and she says in answer to the second question "you should expect to find nothing different." The problem is not that there's a paucity of evidence or that the evidence is ambiguous or dependent for its status as evidence on some background theory. It is, rather, that we cannot even conceive of what would be evidence one way or another. Nothing is different whether you should lie to Sally or whether be honest with her.

This indicates, says the subjectivist, that there is no fact of the matter whether equality is better than freedom or freedom better than equality and no fact of the matter whether you should lie to Sally or whether you should be honest with her. Thus, all four judgements here are false.

Value judgements cannot express truths, for there are no value properties in virtue of which a value judgement could accurately describe how things are. There are no value properties for, if there were, they would make a difference to how other properties stand and to what the facts are, but we cannot conceive of any such differences.

The subjectivist is not saying that it makes no difference whether you lie to Sally or are honest with her. It will make a great difference to many things! What the subjectivist is saying that it makes no difference to how things are whether it is true that you should lie to Sally or whether it is true that you should be honest to Sally. Since it makes no difference, neither judgement is true.

To summarize this argument: For any fact, were that fact different, many other facts would be different. But we cannot conceive how anything would be different were some judgement about the intrinsic value of something true rather than false or false rather than true. Thus, there are no facts about the intrinsic value of things. Or: nothing in the world is either good or bad, noble or ignoble, right or wrong.

2. The argument from explanation. This argument, though different from the metaphysical argument, might rest on many of the same presuppositions. Indeed, it might be true that if either of these two arguments is successful, the other one is as well.

How might we go about determining whether something exists? One idea is that if we don't need to suppose the thing exists (or to suppose that things of a sort exist) in order to

explain anything, the thing doesn't exist. We don't, for instance, need to suppose that unicorns exist in order to explain anyone's behaviour, even the behaviour of those who think they've seen or ridden a unicorn. On the other hand, if our explanations of phenomena are simplified or expanded by supposing something to exist, then we have reason to believe that that thing exists. Maybe the idea that there's a fourth primary colour, one humans cannot see, is an example. Apparently, with this assumption much that is otherwise puzzling about the ability of birds to navigate the skies as precisely as they do is resolved. (There is, by the way, independent evidence from bird physiology and anatomy that birds see one more colour than we do.) Finally, if we cannot get away at all without supposing that something exists, then we have fine reason to believe that that thing exists. Do numbers and mathematical relations exist? Well, our best physical theories employ numbers and mathematics in their explanations and they use mathematics to generate predictions. If these theories cannot be recast so as not to involve mathematics, then we are entitled to hold that numbers and mathematical relations are real features of our world.

The argument from explanation for the unreality of value properties proposes that nothing–no event, no phenomenon–is explained by noting that something was in itself good or bad or right or wrong. It's not the goodness of charity that explains why a person donates clothes to charity, but rather her belief that her donated clothes will help someone to stay warm and her desire that people aren't cold. The idea is that nothing comes about because it is good or right that it come about, and nothing fails to come about because it is bad or wrong that it come about. We don't need to suppose that there are value properties inherent in things or states of affairs in order to explain whatever happens or to predict successfully what will happen.

We are never required to suppose that an event or a phenomenon has any sort or degree of intrinsic value to explain that event or phenomenon, or at least so begins the argument from explanation. But that we are not required to suppose a property exists or that things of a certain type exist in order to explain well what goes on, then we have evidence that that property does not exist or that nothing is of that certain type. Thus, we have evidence that no events or states of affairs have whatever properties would be denoted by "good," "bad," "right," or "wrong." Therefore, nothing is really, in itself, good, bad, right, or wrong.

Those who reject this argument are not convinced that explanation can always proceed without normative or value terms. It all depends on what one wants to explain. Perhaps to explain why some event occurred one need not advert to any value anything might have. But to explain why something is in itself good, one would need to use evaluative terminology. So, to explain why friendship is good, one would have to cite the goodness of its components or effects or nature. Maybe in explaining why the universe exists, or why conscious beings exist in it, one would have to cite the goodness of existence or refer to the pointlessness and, thus, the lack of worth, of a universe without anything to contemplate it or enjoy it.

3. The psychological argument. "It's wrong to tease children for fun." Someone who says this and says it sincerely would seem to be averse to teasing children for fun. Teasing children for fun, we're inclined to say, is either something that this person doesn't go in for–or, if he does go in for it, it is something he doesn't want to go in for. It is, at least, something that disturbs him and something that he would rather did not occur.

"Toronto is to the west of Montreal." Someone who says this and says it sincerely would seem to be averse to ...well, who knows? We don't seem to be able to conclude anything about her likes or dislikes, what she goes in for or doesn't go in for. That she believes that Toronto is to the west of Montreal might be great significance in explaining what she does or doesn't do, but not all by itself. We need to have some idea of what she wants to do in order to come to a conclusion about what she will or should do given her belief about where Toronto is. If she is in Montreal and wants to get to Toronto, we could conclude that she will try to travel westward if she can. But without independent information about her desires and emotions, we remain at a loss regarding what effect believing that Toronto is to the west of Montreal will have on her behaviour.

That's not surprising, that knowing that a person believes something leaves us in the dark with regard to what he might do, one might think. In order to have a practical reason, a reason to do something, one needs to have both a belief and a desire or emotion. Beliefs guide us in our projects to our goals; we cannot discover a person's goals just from information about his beliefs. On the other hand, a person's desires and emotions, her values, do set her goals for her, at least her goals generally or vaguely conceived. That a person fears bears, for instance, implies that she will avoid bears if she can. That a person loves Toronto implies that he will try to spend time in Toronto, so long as doing so is consistent with the rest of his loves and hates.

Beliefs, then, it appears, do not motivate. Only desires or emotions do that. But someone who says sincerely that it's wrong to tease children for fun has revealed to us, at least vaguely, a motivation. He's motivated, at least a little, not to tease children for fun and to try to prevent children from being teased for fun. Or he is unhappy with and made uneasy by his proclivity to tease children for fun. We seem to have stumbled on a contradiction. Beliefs do not motivate, but the belief that teasing children for fun is wrong does motivate. The way out, says the subjectivist, is to suppose that "teasing children for fun is wrong" expresses a desire or an emotion. It is because it expresses a desire or emotion that holding that judgement has an effect on one's will.

Perhaps "it's wrong to tease children for fun" can also express a belief. But whether it does is of little moment in understanding what we do or why we do it. What's really at play here is a desire or an emotion.

The psychological argument concludes that value is really about desires and emotions, not about how the things in the world stand. It's about the activity of valuing, not about discerning something in the world. The psychological argument begins with the premise that a belief by itself cannot motivate. It adds that judging that it is wrong to tease children for the fun of it *can* motivate one to do something. From this, the argument concludes that one's judgement that it is wrong to tease children for the fun of it expresses a motivation. But only a conative or affective attitude can move one to action. Therefore, one's judgement that it is wrong to tease children for the fun of it expresses either a conative or an affective attitude. As such, it is evaluable only in the ways that conative and affectives attitudes are evaluable. Such attitudes are not evaluable through evidence or other reasoning meant to establish the facts of the matter. In the end, then, evaluative or normative judgements, though they can be expressed using indicative sentences, are without cognitive significance in our lives. They

just don't play the roles that beliefs play and, so, it is pointless to investigate or discuss them as though they express beliefs.

The premises and strict inferences in this argument might be consistent with the claim that evaluative judgements, judgements of good and bad and right and wrong and the rest, can and maybe always do express beliefs. When one says it is good to be concerned for the stranger among us, one might be expressing a belief after all, and perhaps a true one. The argument concludes, though, that that is not where the action is when it comes to evaluative judgements. The belief content of the judgement is irrelevant to whatever interest we have in thinking about and judging evaluative judgements. Our interest in them is furthered only when we consider them to be expressions of desires or emotions. Desires and emotions–that's where the action is when it comes to ethics and the other departments of value.