

WHAT IS COMMONSENSE MORALITY?

Robert Frederick

ANN: Good question. I'm glad you asked. Commonsense morality is that system of moral rules that we use in everyday life to make judgments about the character and actions of other people. And of ourselves, of course.

BILL: Okay, but what is a moral rule? And what makes common sense moral rules into a system of rules instead of a mere list?

A: A moral rule is a rule that instructs us what to do, morally speaking. And a system of moral rules is, well, systematic instead of an eclectic hodgepodge.

B: But how do we figure out whether a rule is speaking morally or in some other tone of voice? And how do we figure out that the rules are systematic and not some hodgepodge?

A: Let's begin with the second question. For rules to be systematic they all have to be about the same thing. If you mix up the rules of the road, the rules of chess, and the rules of arithmetic, then you have a hodgepodge. Ergo, to be systematic the rules of commonsense morality can't be all mixed up. They all have to be about the same thing, namely, morality.

B: So a rule like 'never discuss morality after two a.m. or two six packs of beer, whichever comes first,' is a moral rule since it's about morality?

A: Not exactly. True, it's about morality, and yes, it tells us what to do, but it's not a moral rule because it doesn't tell us what to do, morally speaking. It's a sort of prudential rule, or more likely, someone's feeble idea of a joke.

B: For commonsense morality to be systematic, then, it must not be mixed up with other kinds of rules, such as rules of prudence or etiquette.

A: Now you are catching on. Maybe an example would help. A commonsense moral rule is 'it's not permissible to break a promise.' A rule of etiquette is 'it's not permissible to drink from the fingerbowl.' Clearly, there's a difference.

B: So all we need to do to have systematic set of commonsense moral rules is make a list and separate out all the non-moral rules? Aren't we now back to the question: what is a moral rule?

A: Let's stick with the system problem for a minute since for the rules to be systematic there needs to be other restrictions. For instance, the rules must be related in the right way. Not any old list of commonsense moral rules is a systematic list.

B: How about 'abortion is always wrong' and 'abortion is not always wrong'? Since they are systematically related by negation, are they a part of the commonsense system of rules?

A: That's a misunderstanding. The rules must be appropriately related, and negation is not an appropriate relation. Moreover, they must be related in such a way that following a specific rule doesn't commit us to breaking other rules. So for any two rules to be part of the system, they must not only be logically consistent, they must be compatible in the sense that observing one of them is compatible with not violating the other.

B: But suppose that to prevent serious harm someone must break a promise to have lunch with another person? Since it's impossible in this situation to simultaneously observe the rules 'keep your promises' and 'prevent serious harm whenever possible,' does that mean that these commonsense rules aren't compatible and so can't both be a part of commonsense morality?

A: No, because in the situation described the rule to prevent harm overrides the rule to keep promises. That preserves their compatibility.

B: Why does the rule to prevent harm take override the rule to keep promises?

A: Because in the case you describe it would be morally much worse to allow the harm than to break the promise.

Breaking the promise is not insignificant, but it's relatively trivial compared the harm.

B: Does the rule to prevent harm always override the rule to keep promises?

A: No, sometimes a promise must be kept even if keeping it allows harm to occur. It depends of the importance of the promise and the severity of the harm.

B: How do we judge the importance of a promise? How do we determine the severity of harm? And how is it possible to balance the importance of one against the severity of the other?

A: Those are difficult questions, but in general we need criteria for judging importance, determining severity of harm, and weighing importance against severity. Without such criteria we would have no principled way to decide what to do when these rules conflict.

B: Isn't a criterion just a rule?

A: Well, yes.

B: So to use a commonsense system of moral rules we need other rules to tell us what to do when the moral rules seem to conflict?

A: In a sense, yes.

B: Are these new rules also moral rules?

A: Not the ones I have in mind since they don't tell us what to do, morally speaking. Rather, they help us correctly apply the moral rules.

B: Tell me more about these other rules. Do they have to be systematic as well?

A: I see that I need to say more about this, so here goes. When we make judgments about the importance of a promise or the severity of harm we do so against a background of commonsense beliefs about promises and harms. Some of these beliefs are about the relative importance of promises and the relative severity of harms. For example, most of us believe that it's more important to keep a promise to pay a big debt than to keep a promise to have lunch with a casual friend.

B: That seems right so far.

A: The same for harms. Most of us believe that a broken leg is a much more severe harm than a paper-cut.

B: Where is this going?

A: Since we have this background of beliefs about promises and harms, and I should mention that we may not all have the same beliefs about these things, the judgments we make about the importance of a promise or the severity of a harm are judgments of fact, of a sort. Since these facts are relevant for our application of moral rules, let's call them 'morally relevant facts.'

B: I'm not sure I follow.

A: Think of it like this. Commonsense background beliefs about the importance of promises and the severity of harms give us a method of measurement – they are the criteria against which we make judgments of relative importance and severity. They allow us to place promises on a rough scale of relative importance and harms on a rough scale of relative severity, and this helps us see which promises are more important and which harms more severe.

B: I get it. Now all we have to do is weigh them against each other and that will tell us what to do in cases of conflict.

A: Not exactly. These 'weights' aren't really comparable in any straightforward sense. It's more like when we judge the relative redness of two apples, and the relative sweetness of their taste. We might decide that this one is redder and that one is sweeter, but it doesn't make any sense to then compare the redness of one against the sweetness of the other. Redness and sweetness are different things. Something like that is going on with the importance of promises and the severity of harms. They are different things.

B: So how do we decide between keeping the promise and preventing the harm? Is it an arbitrary choice?

A: Not necessarily; we'll get to that, but there are a few other details that we need to discuss first.

B: There's more?

A: Yes. We have these 'facts' about the importance of promises and the severity of harms, but we don't yet know

what to do with them. By themselves they don't instruct us what to do, morally speaking. They are just morally relevant facts, not moral instructions or rules.

B: So we need moral rules. Fortunately, we have them.

A: True, but not the right kind. So far we only have 'keep your promises' and 'prevent harm whenever possible'. We have no rules about what to do when these rules conflict.

B: Are these additional rules also commonsense moral rules?

A: In the case you describe I think so, but I'm not as confident about other cases that might arise. However, the rule we need is 'when the choice is between preventing a severe harm and breaking a relatively unimportant promise, one should prevent the harm'. That's a moral rule since it tells us what to do, morally speaking. It preserves compatibility, and it prevents us from making arbitrary choices about keeping promises versus preventing harms.

B: Good; now we've solved the problem. We've got the morally relevant facts: the promise is unimportant and the harm severe. All we need to do is follow the new rule. Easy.

A: Not as easy as we might wish. Call the new rule 'R1.' Consider another rule, R2: when the choice is between preventing a severe harm and breaking a relatively unimportant promise, one should keep the promise. Most of us, I assume, would say that R1 is the right rule, and R2 the wrong one. But using only the tools of commonsense morality it's not easy to say precisely why.

B: Sure it is. R1 is the right rule and R2 is wrong one because if we follow R1 we prevent a great deal of harm at a small cost, but if we follow R2 we allow great harm for only a small gain.

A: It might seem that way, but note that what you are doing is making direct comparisons between preventing harms and keeping promises, which you can't do because they aren't comparable in that way.

B: So what do we do?

A: There's more. R1 and R2 are rules about what to do if there is a conflict between the commonsense rules like

'keep promises' and 'prevent harm.' We need another rule to tell us what to do when there is conflict between other commonsense rules. For example, the rules 'do not kill' and 'do your duty' can conflict in wartime, so we need a rule to tell us what to do in that situation. This will be another R rule, and it might be very complicated. Moreover, these rules, 'R-level rules,' we can call them, might also come into conflict.

B: What then? Might we need 'S-level' rules to tell us what to do when R-level rules conflict?

A: It's possible.

B: That's discouraging.

A: One last thing. Remember I said two people might have different rankings of the importance of promises and the severity of harms?

B: Yes; so what?

A: Suppose the people are you and me, and suppose you think promises to go to lunch are much more important than I do.

B: Alright, suppose I do.

A: Then we might both accept R1 and still disagree about what to do because we disagree about the importance of the promise.

B: You mean I might think that the promise is so important that I'm willing to allow the harm to occur?

A: Right. If we don't rank promises and harms the same way, we will disagree about the application of R1.

B: Can't we come to an agreement about how to rank promises and harms?

A: Possibly, but not necessarily. We can discuss our rankings – if we can get clear about them ourselves – but even then we might not resolve our differences.

B: This is very distressing. I thought commonsense moral rules were something ordinary people could use in their ordinary lives, but all of these rules about systems and rules about rules and who knows what else seem much too complicated for that. Could something be wrong with the whole idea of commonsense morality?

A: I can see that our progress has slowed. Maybe it's time to bring in other resources.

B: I'm game. What are they?

A: Suppose we had a moral theory, something grander, more abstract, than mere commonsense moral rules. Then, if we have the right theory, maybe we could get commonsense rules from the theory. Maybe even rules like R1.

B: Let me get this straight. We have a moral theory of some kind, and we logically deduce commonsense moral rules from the theory.

A: No, nothing quite as formal as that. Let's say that commonsense rules 'follow from' the theory, but I don't want to say that they are strictly logically implied by the theory.

B: So they bear some sort of logical relation to the theory, but you don't want to say what it is.

A: Something like that. I will say that commonsense rules follow from the theory, but the theory doesn't follow from the rules.

B: No doubt you'll clear things up later. For now, how are we better off with this theory?

A: Several ways. First, it should give us a way to distinguish between moral rules and other kinds of rules. If a purported rule follows from a moral theory, then it's okay. But if it doesn't, then it's some other kind of rule, such as a rule of prudence or etiquette. Also, the theory may help us make a principled choice between R1 and R2 and other rules like them.

B: So genuine moral rules follow from a moral theory, but rules of etiquette don't.

A: Right; an advantage of having a moral theory is that it gives us a more comprehensive and organized way to view commonsense morality.

B: Maybe, but that depends on the theory. What is it?

A: Sorry?

B: The theory – the moral theory you mentioned. I want to see if it will work as you say it will.

A: Yes, that brings up a little problem.

B: I should have seen this coming.

A: Actually, there are several theories. Maybe half a dozen or so, for starters.

B: Half a dozen.

A: Or so. It might help if we divide them into three categories of moral theories: *deontological* theories, *consequentialist* theories, and *virtue* theories. There may be one or two others, but these will do for now.

B: Fine; let's stick with these three. Now, as I understand it, commonsense rules follow from these three different kinds of theories.

A: Not exactly; each category might imply different set of rules. For example, one set of commonsense rules might follow from deontological theories, and a different set from consequentialist theories. There will be overlap, but it won't necessarily be complete.

B: Now I really am confused. If I follow, there could be three sets of commonsense moral rules, each set following from a different category of ethical theory. That just doesn't make sense. How could three sets of different rules all be 'commonsense' rules? Isn't there just one set of commonsense rules?

A: I was hoping to put this off, but it's time to face facts. It's unlikely that there is a single set of logically consistent, compatible, and systematic commonsense rules. Instead, we have a commonsense moral rule hodgepodge, which is made worse by the fact that people often disagree about how to rank things like harms and promises. What people call 'commonsense morality' is really not much more than a little bit of this and a little bit of that we collected together over the centuries. It doesn't have unified basis in a single moral theory or viewpoint.

B: I just don't think you're right. The core of commonsense morality isn't difficult. I can sum it up in two rules: do no harm, and don't break trust. Any act that violates these rules needs a strong reason. Without it, the act is immoral according to commonsense morality. That's not so hard.

A: Had we started here things might have gone more smoothly, but I see that being forthright is not your style.

Still, the rules you mention are a major part of common-sense morality. Not all of it since your two rules only tell us what to avoid doing, not what we should do in a more positive way. But let's stay with them for now. You see that these rules can conflict.

B: You mean that to do no harm I might have to break trust, or that to keep trust I might have to do harm. Now I suppose we need R-level rules, along with rankings of harm and breaking trust, to tell us what to do when this happens.

A: Right again, but now we are in a better position to say how these different kinds of rules might work together with moral theories and morally relevant facts.

B: Go on.

A: Okay, you said that to permissibly break one of the rules requires a strong reason.

B: Yes; you have to have a really good reason to break a moral rule. If you don't, you act immorally.

A: Where does this requirement to 'have a good reason' come from?

B: I have no idea what you mean.

A: You said that the rule is 'do no harm,' not 'do no harm unless you have a good reason.' So where does the 'unless you have a good reason' part come from?

B: I hadn't thought of the rules that way before. I suppose it's conceivable that moral rules never conflict – as a matter of fact, keeping one of them never causes us to break another. That might be nice, but it's not our world. In our world, conflict happens. Maybe, then, the 'good reason' clause comes from our moral experience. We're all familiar with conflicting obligations. The only way out is to think it through, to decide as best we can what we have most reason to do, morally speaking, and then do it. So moral experience tells us that reasons are essential for moral action, especially in cases where our obligations conflict.

A: Suppose you're right. Then it would follow that all commonsense moral rules have a 'unless you have a good reason' clause attached to them, including R-level rules.

B: Maybe some of them don't, but I can't think of any at the moment.

A: So we need to know what counts as a good reason. Any ideas?

B: A good reason is, well, you know, good. It's a reason that other reasonable people would regard as good.

A: All of them?

B: Well, most of them. The really reasonable ones, anyway.

A: You are on to something here, so let's work on it. Here's one possible way to understand the notion of a good reason. First, at the most basic level, something is a reason only if it is a belief or has a belief as a component. If you don't believe something it can't be a reason for you. Second, a good reason is a reason you are willing to act on, other things being equal, and it's a reason you believe other people that you regard as reasonable would agree is good.

B: Is that all?

A: Not quite; there are two other pieces. The first is that the act in question, i.e. the act to which the reason in question applies, must fall under a moral rule. In other words, it must be an act such that, if you did it, we could properly say that it was either the morally right thing to do or the morally wrong thing to do.

B: And the second?

A: This one is tricky. One can do the morally right thing for bad reasons. For example, you might prevent harm to someone solely because you think you will get a reward. Something is a good reason for an act that falls under a moral rule, let's say, only if it would be considered a good reason by all fully rational and emotionally stable people who accept the moral point of view.

B: What's the moral point of view?

A: Someone who accepts the moral point of view typically makes decisions about the morally right thing to do on the basis of good reasons.

B: This feels so circular that my head is spinning.

A: It could be the beer.

B: Maybe, but you've characterized a good reason for moral action by reference to someone who accepts the moral point of view, and someone who accepts the moral point of view as someone who acts morally for good reasons. Isn't that too circular even for you?

A: You promised not to be harsh. I included the idea of the moral point of view because it seems possible that there are people who do not accept the moral point of view. They wouldn't think, for example, that making a promise to repay a debt is a good reason to repay it.

B: So they are 'amoralists' rather like those who reject religion are 'atheists'?

A: That's the general idea. If there are people like that, they won't be much use in helping decide what a good moral reason is. But let me suggest one final point. Good reasons for moral action are related to moral theories. That is, one marker of a good reason for moral action is that it would be seen as a good from the point of view of a particular moral theory. For example, from the point of view of divine command theories of morality the fact that God commands something is a good reason to do it, but for most consequentialist theories, God's commands aren't terribly relevant.

B: Somehow I doubt that God would see it that way. By the way, I thought you said that a good reason for moral action would be so regarded by rational and emotionally stable people. Do you mean to say that these people have to know moral theory?

A: That's another difficult question. No, I don't really want to say that. But moral theory is a deliberate and organized attempt to say something useful about moral right and wrong, and that's got to help when judgments need to be made.

B: It's getting pretty late. I have to get up early tomorrow. Can we wrap this up soon?

A: Be patient; here's my final suggestion. A good reason is not a simple fact, or a rule, or even a theory. We act on

reasons, and considered individually these things are not the kinds of things that we act on. Rather, a good reason is a complex thing that has facts and rules as parts. For instance, the fact that an action would prevent harm along with the rule that one ought to prevent harm, is a good reason for performing the action.

B: When we deliberate about the morally right thing to do, we appeal, at least implicitly, to moral rules and morally relevant facts?

A: Right. Commonsense morality doesn't consist only of commonsense moral rules; it's an enormously complex network composed of ordinary moral rules, rules for cases of conflict, beliefs about morally relevant facts, and, not least, beliefs about the world and human nature. Commonsense morality depends essentially on all of these components. Any attempt to separate out the rules and make sense of them in isolation from the others is doomed to failure. By themselves the rules are useless. They only work as part of a system, a web of moral belief, to modify a phrase. Commonsense morality is the whole package.

B: Remind me again what moral theory has to do with all this.

A: My contention is that the package is riddled with inconsistencies, incompatibilities, and incoherence. It's not one neat bundle, and in its present state can't be sorted out in any rational way that all rational people would agree to. Or so I claim. The only way to make any progress sorting things out is by using moral theory to help begin to make sense of it. I can't see any other way.

B: So the idea is to use moral theory to help distinguish, from the point of view of the theory, good reasons from bad ones. But I still don't see clearly how this works. Can you give an example?

A: Remember your original case – you have to break a promise to prevent harm, and imagine that you are explaining what you did to a friend you regard as reasonable and morally astute. You might say, 'I couldn't prevent the harm

without breaking the promise, and I judged that the harm was severe and the promise relatively trivial. Since I accept R1, I acted to prevent the harm.' No doubt your friend would accept this as a good explanation for your action, and so would you if you and your friend's positions were reversed.

B: Do you mean that a good reason is a good explanation?

A: Close, but not exactly. However, it's true, I think, that if you give a sincere and truthful explanation for your moral action, you give the reasons you had at the time; and if you give the reasons you had for moral action to someone, you are giving an explanation of your action.

B: But didn't you say there were half a dozen or so moral theories?

A: I did.

B: I suppose they all have different views of what's moral and what isn't.

A: I'm afraid so. There are similarities, of course, but there are also important differences.

B: Then all we need do is pick the right one and get started.

A: Started on what?

B: Sorting out commonsense morality, of course.

A: Yes, well, picking the right one is a bit thorny.

B: Let me guess. Because we need criteria to make the choice?

A: Right.

B: And nobody agrees on what those criteria are?

A: Something like that.

B: In other words, at the moment we have no way to make a choice that everyone would agree to, and no prospect of achieving such a happy result?

A: Right again.

B: What do we do?

A: Well, we could begin a philosophical study of moral theories. That might help.

B: How would it help?

A: Maybe if we understood the theories better we would be able to make the choice, or at least identify the criteria for making the choice.

B: But haven't philosophers been studying moral theories for a long time?

A: They have.

B: Have they settled on a choice, or the criteria for making a choice?

A: Let's just say it's an ongoing process.

B: That's just great. It's after 2:00. I'm exhausted.

A: Me too. Any of that beer left?

Robert Frederick teaches at the Department of Philosophy, Bentley College.