

Some Moral Theories and What They Say about Lying

1. Acts and Moral Principles

Whenever someone claims that a certain action was wrong, it makes sense to ask that person why. What was it about the act that you think made it wrong? In asking this, we are asking the person to identify a *moral principle* or set of principles that apply to the action in question and imply that it is wrong. Suppose that Frank is someone who uses drugs like crystal meth when he parties with certain friends, that some of these friends call to see if he wants to party with them, and that Frank's girlfriend, Mary, answers the phone. Hearing who it is, she tells the caller that Frank is not home. When Frank asks her who just called, she lies and says that it was a telemarketer. She does this because she is trying to help Frank stop using drugs. She believes that he is trying to stop, but she also believes that if he goes partying with these friends he'll probably give in to temptation and resume his drug use. So she tells him a *paternalistic lie*, deceiving him for what she thinks is his own good. If you think that what Mary did was wrong, then you ought to be able to identify a moral principle that she violated. Perhaps you think that what Mary did was wrong simply because she told a lie and lying is *always* wrong. She violated a moral principle that *absolutely* forbids lying. Perhaps you think that not just Mary but anyone (yourself included) acts wrongly if they tell a lie for any reason whatsoever. An absolute prohibition on lying is a possible moral principle.

Suppose that Frank discovers the lie and objects. Mary's belief that Frank really wants to stop using crystal meth was based on his assurances that this was so, but he didn't mean it. *He* was lying to *her*, because he didn't want her to leave him. In fact, he routinely lies to Mary (and others) whenever he thinks it's necessary, and he sees nothing wrong with it. If Frank wants to condemn Mary for lying to him, he needs to cite moral principles that explain why. But any remotely plausible moral principle that forbids Mary's lying to Frank will probably also forbid Frank's lying to Mary. Of course one can imagine some bizarre moral principles that don't have that result—such as “It's always okay for drug users to lie to others, but never okay for a non-drug users to lie to drug users.” But if Frank appealed to that principle to defend his moral judgment about Mary, I think that we would doubt that he is being serious. The principle seems obviously contrived or *ad hoc*. (If he insisted that he really does believe that such a principle is true, we should ask him to *defend* it. Why should drug users have special moral privileges that are denied to others?) This doesn't mean that Frank can't condemn Mary. All it means is that he probably can't condemn her without admitting that what he frequently does to others is also wrong. We have to be *consistent* in our moral judgments. Whatever moral principles I cite to defend my judgment that you acted wrongly are principles that I have to be prepared to apply to my own behavior. The same requirement holds for all of us. (Moral philosophers often refer to this as “the universalizability requirement”).

2. Five moral theories.

What is the *correct* moral principle about lying? (Note: Lying is making a statement that one knows is false with the intent to deceive. It's possible to deceive

without lying. Suppose that I dress as a priest, even though I'm not one, because I know that you're a devout Catholic who'll confide in priests, and I want to find out some of your secrets. I've clearly deceived you, even if I never lie and say "I'm a priest.") Which moral principle about lying is the one that belongs to the correct "ordinary" morality? (Recall that "ordinary" morality is the morality that is supposed to apply to all of us, just because we are human beings. It can be contrasted with "role" morality, or the morality that applies only to those people who occupy certain social roles. Professionals are bound by the rules of ordinary morality, just as everyone else is, but they are also bound by the role morality of their profession, because they also occupy the social role of being a doctor, attorney, engineer, or whatever).

I'm going to assume that these questions have answers. *Moral subjectivism* is the view that right and wrong are all a "matter of opinion," where any moral opinion is just as valid or reasonable as its opposite. I'm going to assume that moral subjectivism is false. Questions about what's right and wrong, like questions in math and physics, have correct and incorrect answers. The answers may not always be easy to figure out or obvious. *Some* moral questions may be very difficult, leaving plenty of room for reasonable disagreement about what the correct answer to them is. But that's consistent with *moral objectivism's* being true. Moral objectivism holds that there are moral facts, just as there are math and physics facts, and that these facts are independent of people's opinions about or attitudes toward them. Sincerely and fervently believing that you did the morally right thing does not *guarantee* that you did the morally right thing (any more than sincerely believing that Mt. Everest is the tallest mountain in the world guarantees that it really is the tallest mountain). The terrorists who flew the planes into the World Trade Center towers and killed thousands were firmly convinced that they were doing the right thing. If you think that they were *mistaken* about this, then you have to agree with moral objectivism. If you think it makes sense to wonder if the moral principles that you currently accept are correct or not, then you must assume that moral objectivism is true.

Moral objectivism says that there are *correct* (and incorrect) principles about when it's okay to lie. There are three possibilities: i) it is morally wrong to lie in all circumstances, ii) it is morally permissible in all cases, and iii) it is sometimes permissible and sometimes wrong. Moral objectivism does *not* say that i) has to be the correct view. One could accept moral objectivism and think that iii) is correct. Of course, stated as it is iii) isn't very helpful. What one wants to know is what are the moral principles that explain when lying is okay and when it's wrong.

In what follows I'll describe five different *moral theories*, focusing on what each says about lying. The theories are:

- i) The divine command theory of right and wrong.
- ii) The natural law theory of right and wrong.
- iii) Act utilitarianism.
- iv) Kantian moral theory, and
- v). Rule utilitarianism.

Each of these theories offers its own answer to the question, "How should one determine what are the correct moral principles about anything, including lying?"

3. The divine command theory (or DCT).

This view is in a way quite simple. It holds that if an act is wrong, what makes it wrong is the fact that there is a God who commands us not to perform acts of that type. If lying is wrong, then it's wrong because God forbids it. His forbidding it is what makes it wrong. In order to know what's right and wrong, one has to know what God's commands are.

How does one figure out what God's commands are? Different world religions offer different answers to that question. According to the Jew one must turn to the Old Testament. According to the Moslem, it's the Koran. According to the Christian, it's (mainly) the New Testament. The fact that there are many different religions with different views here means that there are many different versions of DCT: Christian DCT, Moslem DCT, etc. How can anyone know which if any of these religions is correct in its claims about what God's commands are? If the answer is unclear, that creates a problem for anyone who wishes to defend some version of DCT.

What does Christian DCT say about lying? Well, one of the Ten Commandments, the ninth, forbids "bearing false witness" against others. The main thing to notice about this commandment (as well as the other ones) is that it includes no qualifications and thus would seem to be *absolute*. It seems to say that "bearing false witness" against others is *always* wrong—no matter how much good it would do or evil it would prevent, no matter whether one's objective in telling the lie is selfish or altruistic, no matter whether the lie would benefit the person being lied to, etc. If the Christian DCT is correct, then what Mary did in lying to Frank was wrong.

4. Natural Law Theory.

This theory assumes that everything has a function/purpose and "the good" for a thing is whatever helps realize that function, while "the bad" is what hinders or thwarts the fulfillment of that function. The theory assumes human beings as a whole have a function, as do their various parts (eyes, heart, sex organs, etc.); the functions of the parts complement the function of the whole. When human beings act in a way that is contrary to these functions, they are acting immorally. Since immoral acts are contrary to our function or nature, they may be said to be "unnatural." The most prominent proponent of natural law theory was the medieval Roman Catholic theologian Thomas Aquinas. The theory is accepted by the Roman Catholic Church, which uses it to defend its condemnation of fornication, homosexuality, and masturbation. Using natural law theory, Aquinas held that engaging in nonprocreative sex acts is wrong, because "unnatural" (i.e. contrary to the functions of us and our parts.) Masturbation, for example, is supposed to be wrong because the function of the sex organs is procreation and masturbation is a kind of sex act that cannot possibly lead to procreation.

What does Aquinas's theory say about lying? It says that lying is wrong because contrary to the function of speech, which is to communicate to others the ideas in our minds. It agrees with Christian DCT that lying is wrong not because of its *consequences*, but because of the intrinsic nature of the act.

Aquinas tells the story of a Christian in Roman times whom some Roman Centurions were looking for to arrest. The Christian is in a boat moving upstream on a river, and the Centurions are in a boat moving downstream. As the two boats pass, the Centurions ask the Christian if he knows where the man whom they're looking to arrest (namely, the man whom they are talking to; only they don't know this) is. His answer is

that if they continue going down the river, they will eventually have passed the man whom they're looking for in another boat. Aquinas claims this answer is not morally wrong. The Christian did not say anything false, and therefore, did not lie. As I recall, Aquinas also says that the Christian (who eventually was captured and martyred) *foresaw* but did not *intend* that the Centurions misunderstand his answer to their query. Hence, the Christian not only did not lie, but he also did not deceive the Centurions (because deception requires an intent to deceive). I'm not so sure I agree with Aquinas that there was no intent to deceive in this case.

Like Christian DCT, natural law theory supports a principle that forbids lying *absolutely*. There is, however, an important difference between the two theories worth mentioning. According to natural law theory, we can figure out by means of unaided human reason, without religion, that lying is wrong. Even people ignorant of divine revelation/God's commands are supposed to be able to figure out what the function of speech is and how lying is contrary to that function. For DCT, on the other hand, only people who are familiar with the word of God will know whether lying is ever morally permissible. Hence, DCT makes ethics dependent on religion in a way that natural law theory supposedly doesn't.

5. Act utilitarianism (AU).

This theory claims that an act is morally right if and only if it produces more total happiness (not just for oneself, but all of society) than any other act that one could have performed in the circumstances. The morally right act is the one that maximizes "utility," which just means "happiness" or "well-being." For AU happiness alone has intrinsic value. If anything else is good or valuable, such as telling the truth, that's only because it's a *means* to happiness. Assuming that honesty usually produces more happiness in the long run than dishonesty does, (AU) says that we should *usually* approve of honesty and disapprove of dishonesty. (AU) says that "don't lie to others" is a good "rule of thumb" for how to maximize utility. But it is only a rule of thumb, not an iron-clad rule that admits of no exceptions. If you're in one of the rather exceptional situations in which lying would produce more total happiness in the long run, then AU says that you have a *duty* to lie. Of course AU admits that people are far from infallible in their judgments about how much good or evil the different actions they could perform might produce. It says that we should decide what to do on the basis of the best predictions we can make about the future, predictions based on all of the information that's available to us.

One would think that lying to junkies for their own good would usually produce more total happiness than not lying to them would. But lying to people for their own good is certainly risky. If someone discovers that I've lied to him for his own good, it will probably weaken his trust in me, and that in turn may make it more difficult for me help him in the future. So it may not be entirely clear whether or not AU approves of Mary's lying to Frank (given what she knows or thinks she knows about him). Clearly the moral principle that AU supports here is that lying is wrong only when it produces less total happiness in the long run than honesty would. What's not so clear is how this principle applies in the Mary and Frank example.

Mary's lie may not be one of them, but there seem to be other cases in which there can be no doubt that lying is the way to maximize utility. Suppose that you live in the Netherlands in 1941. The Nazis come knocking on your door asking if you know where any Jews are hiding. You know that your neighbor is hiding some in his cellar.

Lying to the Nazis by saying that you don't know would clearly produce more utility than telling the truth or refusing to answer their questions would.

AU is a *consequentialist* or *teleological* moral theory. Such theories claim that right acts are ones that maximize what is intrinsically good. Because AU identifies the intrinsically good with happiness, it must say that honesty is valuable only as a means. Other teleological theories needn't agree with AU on this point. There is no reason why a teleological theory can't take a more expansive view about what's intrinsically good, claiming that not just happiness, but other things such as knowledge, beauty, play, friendship, rationality, or even honesty itself have intrinsic value. Would such a theory agree with Christian DCT and natural law theory that lying is always wrong? The answer is "no." All consequentialist theories assume that "the end justifies the means." Imagine a consequentialist theory which held that honesty is the *only* thing that's intrinsically good and dishonesty is the *only* thing that's intrinsically bad. (Of course such a theory would be extremely implausible, and I'm aware of no one who has ever accepted it). Not even *this* theory would agree that it's *always* wrong to lie. Why? Because it's possible to imagine cases in which if you tell a small lie, that will prevent a much larger number of much bigger lies being told by others, and your telling the lie is the only way to prevent those other lies. In such cases this theory must approve of your lying, because even though it is evil, it would prevent much greater evil later. Teleological theories assume that if something very bad will happen unless you do X, then even if you're not the "cause" of that very bad thing, you're just as much to blame for it if it happens (because you could have prevented it by doing X but didn't) as you would be if you were its cause. Any teleological theory must approve of doing evil acts whenever doing them is the only way to prevent greater evil. By contrast, both Christian DCT and natural law theory deny that preventing greater evil by others could ever justify doing evil oneself, and thus, that preventing deception by others could ever justify acts of deception by oneself.

This doesn't mean that Christian DCT and natural law theory never evaluate acts on the basis of their consequences. Both theories forbid "fornication" (heterosexual intercourse between unmarried adults) as bad for its own sake. (By contrast, AU condemns fornication only if/insofar as it produces bad consequences like children being born out of wedlock and not properly raised and cared for). Because these theories regard fornication as intrinsically wrong, they should probably call for the FDA not to approve sale of the "morning after pill" to the general public, on the grounds that the pill's widespread availability would increase fornication. They would object to FDA approval on account of what they think would be its bad consequences.

For AU, there is no difference between the moral rules that apply to professionals and the moral rules that apply to nonprofessionals. There is no difference between "ordinary morality" and "role morality." Instead, there is only one fundamental moral rule, and it applies to all of us all of the time: "Do the act that produces the most happiness, taking everyone's interests into account equally." Imagine an attorney who knows where there is evidence (his client told him where it is) that would convict his client, a vicious criminal, of the crime of which he is accused. Assume that if the attorney respects client confidentiality, then it is likely that his client will be acquitted (the prosecutor right now has a weak case) and it is likely that he will commit more vicious crimes as a free man. If the attorney secretly or anonymously informs the prosecutor of where this evidence is located, then it is likely that his client will be

convicted. Given these assumptions, secretly violating client confidentiality is likely to produce more total utility for society. Therefore, AU claims that at least in this case, secretly violating one of the principles of professional legal ethics is the morally right thing to do.

Probably the most significant objection to AU alleges that it can't be the correct moral theory because it requires injustice in certain circumstances. If the way to produce the most utility is to perform an act that violates the rights of a few, then AU says that it is morally right to perform the act. But shouldn't individuals' rights be respected, even if violating them would do more good for society as a whole? An example that illustrates this objection: Imagine a doctor with one healthy and five very sick patients. These five will die unless they receive organ transplants. One of them needs a new liver, another a new kidney, another a new heart, etc. It just so happens that the healthy patient would be a suitable organ donor for all five of them. If the surgeon carves him up and redistributes his organs, he can save the five lives. If he does nothing, the one healthy patient lives, the five sick ones die. Since saving 5 and killing 1 is likely to produce more total utility than letting 1 live and 5 die, AU says that the surgeon should perform the transplants—even if the healthy patient does not consent to parting with his organs. But isn't it wrong to violate the healthy patient's right to life (by killing him without his consent) even if violating that right is the way to maximize utility?

There are many sorts of cases, ones where nobody's rights would be violated no matter what we decide to do, where it seems reasonable to use AU as the method for determining what should be done. Suppose, for example, that we want to build a highway that connects two cities. The shorter route would require destroying some scenic wilderness that is enjoyed by some nature lovers. The longer route avoids that but entails a longer driving time for people who commute between the two cities. Which route should the highway be built on? Here it's plausible to think that we should make the decision on the basis of utilitarian considerations. We look at all the costs and benefits of both alternatives and pick the one with the most favorable benefit to cost ratio.

6. Kantian moral theory.

This view gets its name from the 18th century German philosopher who expounded it, Immanuel Kant. Kantian ethics has many defenders among contemporary moral philosophers, because it seems to have two strengths: i) it does not depend on religion or any religious assumptions, and ii) it demands respect for individuals and their rights, and thus, avoids the main objection to AU.

Kantians believe that it is important to treat people with *dignity* or *respect* because they are moral agents with capacities for free will and reason. (Since animals lack those capacities, they are nonpersons and the Kantian thinks that means that they may be treated as things or mere resources). Kant insisted that we must always treat humanity, both oneself and others, as “ends-in-themselves,” never as “mere means.” This principle is called the Principle of Humanity. Kant claims that it is one way of stating the supreme principle of morality, the “Categorical Imperative.” He claims that another, equivalent way of stating the supreme principle is that one must act only on maxims that one could will be universal laws that apply to all. That's the “universal law” formula, which we'll ignore.

You treat someone as an end when you treat him in a way that he would not object to if he were morally reasonable, thinking clearly, and well-informed about relevant factual matters. When society puts a thief in prison, it's treating him as an end, because even though he might object to being incarcerated, he wouldn't object if he were morally reasonable; he would admit that thieves deserve to go to jail. Suppose, to take another case, that you knock a cup of coffee out of my hands before I can drink from it. I object that you have no right to interfere with my freedom, but only because I assume, incorrectly, that there's nothing wrong with the coffee; in fact, it contains poison. In that case too you have treated me as an end, because I would not object to your interference if I knew this about the coffee. Suppose, finally, that a woman persuades me to sell her my car because she tells me a lie that I believe ("males who drive this type of car are 20 times more likely to develop testicular cancer than males who don't"). In that case she runs afoul of the principle of humanity because she manipulates me or treats me as a "mere means" to achieving her own goals. She fails to respect my status as an autonomous being with capacities for free will and rationality.

Kantian ethics has a direct bearing on lots of professional-client interactions. There are some cases where a doctor's lying to or withholding information from a patient would probably be better for him in the long run than full disclosure would. In these cases there is a conflict between the value of beneficence ("do what's in the patient's long term best interests") and the value of respect for the patient's autonomy. The Kantian says that the duty to respect autonomy should take priority, while the utilitarian says that beneficence should. Kantian ethics supports the idea that doctors have a duty to obtain "informed consent" from a patient before treating him for his condition.

Kant held that all deception violates the principle of humanity, treating the person to whom the lie is told as a "mere means." Even if someone wishing to commit a murder asks you the whereabouts of his intended victim and you know, you may not lie to throw him off track. (Of course, that doesn't mean that you must answer him honestly. You can try to be evasive or simply refuse to answer). Whereas AU says that a lie is permissible whenever telling one would do more good or prevent more harm than honesty would, Kant held that telling a lie is never morally permitted.

7. Rule utilitarianism (RU).

(RU) accepts the basic utilitarian idea that things like murder, rape, and robbery are wrong because of their bad consequences in terms of reducing the total amount of happiness. But it gives a slightly different definition of "morally right" from (AU). According to (RU):

An action is morally permissible only if it is permitted by that public set of rules the general acceptance of which by the members of a community would maximize total utility in that community. An act is wrong if it is forbidden by the same rules. This form of utilitarianism tells us not to look at the consequences of particular actions to determine their rightness or wrongness. We should proceed more *indirectly*. We should look at the consequences of different sets of rules; find the set of rules which, if people accept them (i.e. for the most part try to conform to them), will produce more utility than their acceptance of any other set of rules would; then ask if that set of rules (the "ideal moral code") permits the action in question.

Let's return to the example of the doctor with the 5 sick and 1 healthy patient to see how RU's approach differs from AU's. According to RU, the doctor needs to ask himself which set of publicly accepted moral rules would produce more utility—one which gave doctors the discretion to use the organs of healthy people, without their permission, whenever doctors thought that doing so would benefit society the most, or one which forbade doctors to take living donors' organs without their consent? The answer, says the defender of RU, is clearly the second. If society adopted the first set of rules, then people would be reluctant to go the doctor for treatment of their curable illnesses and injuries, because they would fear that the doctor would use the power that those rules give her to act like God and harvest some of their organs for the greater good of society. Clearly people being afraid to visit their doctors when they are ill would have very bad consequences for society as whole. So RU judges the act of harvesting the healthy patient's organs without his consent, so that 5 other lives may be saved, wrong because forbidden by that set of rules the general acceptance of which would maximize utility. Secretly breaking those rules is not permitted, even if doing so would maximize utility.

RU should not be confused with “ethical relativism.” Ethical relativists say that people have a duty to uphold the moral rules that are *actually* accepted by the majority in their society. If the majority in my society believed that enslaving racial minorities is morally okay, then relativists say that I ought to accept and support that view. (RU) says that I should try to figure out what would be the *ideal* ordinary morality for my society using utilitarian criteria, and then I should obey that *ideal* ordinary morality. If the majority in my society don't accept that ideal morality, that doesn't matter. (RU: “You're supposed to do what really is right, not what everyone else in your society thinks is right.”)

RU provides a method for determining not just what it says are the correct principles of ordinary morality, but also what are supposed to be the correct principles of professional ethics. If one code of professional ethics would produce more total happiness for a society than a different code, then RU says that professionals have a duty to obey the rules of that code, even if the professional societies have not recognized all the rules of that code. (The professional societies *ought* to recognize that code, according to RU). Professionals ought to obey the rules of the ideal professional code, even if violating one of them on some occasion would produce more total utility. (By contrast, AU says that the professional should violate the rule in that case). Whereas AU tells the attorney in the example given earlier to violate client confidentiality, RU probably tells him not to. After all, if society had a rule allowing attorneys to rat out any client whom they thought was guilty, what would be gained? Knowing that such a rule was in existence would lead guilty clients not to divulge their incriminating secrets to their attorneys. The main drawback of the rule is that it would lead some clients who *mistakenly* believe that they have committed a crime not to divulge all relevant information to their attorneys. (Their mistake is due to their ignorance or misunderstanding of the law, something about which their attorney is an expert and they are not). These clients will receive worse legal representation than they deserve and would get if there were a rule in place requiring attorneys to uphold confidentiality in almost all circumstances.

What sort of rule or principle about deception is the one that would maximize utility if everyone accepted and tried to conform to it? There are many possibilities, including:

1. Never deceive others, under any circumstances whatsoever.
2. It's okay to deceive evil people if that's the only way to thwart their evil plans, but it's never permissible to deceive others.
3. It's okay to deceive anyone, evil or not, if that's the only way to prevent a great evil (e.g. a death) from occurring.
4. It's okay to deceive anyone about anything if the deception would produce more total happiness in the long run than nondeception would.
5. Not all deceptions are equal; some are more serious than others. "Big" lies are permissible only if necessary to prevent very big evils from happening. But "small" lies (or "white" lies) are justified if they will prevent smaller evils, such as hurt feelings. (Example: A proud mother shows me her newborn and asks "don't you think he's beautiful?" I reply "yes," even though I think he's one of the uglier babies I've ever seen).

I can't believe that the ideal ordinary morality of RU would include 1 or 4. But which of the other 3 is best may not be so clear. One problem with 5 is that it may not be entirely clear how to distinguish "big" from "small" lies.

A final example: Suppose that Joe is on his deathbed, but he is lucid during the last few moments of his life. He asks me how his adult daughter whom he loves dearly is doing. Instead of telling him the truth—that she and her children were killed by a drunk driver last night—I lie to him, telling him that she just won the lottery. I do this so that he can "die happy." Is that a "big" or a "small" lie? Christian DCT, natural law theory, and Kant would all disapprove of this lie, claiming that it is morally wrong. AU would surely approve of it. What do you think?