2

ETHICS
THE BIG THREE THEORIES

In the last chapter we examined divine command theory, ethical egoism, and moral relativism. All of these theories have been historically influential, and each continues to have modern supporters who try to answer (or at least swallow the pain of) the objections presented. Nevertheless, they aren't really the big players in contemporary ethics. In this chapter we will examine the three 800-pound gorillas of ethical theory. 2.1

Utilitarianism (Is Morality Doing What I Can to Make This the Best World Possible?)

One of the criticisms of ethical egoism in the preceding chapter was that egoism meant that each person treated himself or herself differently than everyone else in the world, even if there was nothing that merited this differential treatment. That is, egoism violated the principle of equal treatment. A related problem cropped up in the case of moral relativism, in that if morality is restricted to cultures or societies, then how you should treat members of your own culture may be wildly different from how you should treat people in other cultures. The duties you have to fellow Roman citizens are completely unlike your duties to outsiders; in fact you might have no obligations to the barbarians at all. Here too there is a sort of equal treatment problem—the in-group/out-group distinction doesn't seem to be a relevant distinction for a difference in how you treat people, as moral relativism demands. 2.2

This Is Philosophy: An Introduction, First Edition. Steven D. Hales.

2.3 A moral theory that does treat everyone equally, without prejudice to personal standing, is utilitarianism. Utilitarianism is an enormously influential ethical theory. The basic idea is that moral action is all about producing good in the world; the more good your action produces, the better it is. Your moral duty is to perform whatever actions are the best ones in this sense. Utilitarianism is focused on the outcomes of action—will a possible action create happiness, produce pleasure, and improve the lives of those it will affect, or will it cause pain, harm, and make people worse off? There is a fundamental tie to what morality is intuitively all about, namely the improvement of our lot by increasing our well-being and easing our burdens.

2.4 In addition utilitarianism is able to provide, in principle, an answer to every moral question or ethical dilemma. Should we legalize drugs? Well, will doing so lead to net gain in our collective happiness (because people are freer to do as they wish, fewer people will be in prison, and we will have tax revenues from drug dealers) or will it lead to a decrease in our net happiness (because there will be more addicts, less productivity, and more DUIs)? All we need to do is settle the question about potential consequences and we automatically get an answer about the morally correct course of action. Should abortion remain legal? Should you steal music? Should we kill animals for food? Even if the answers to these questions are not obvious, utilitarianism still provides the means to answer them. In this way it is an incredibly powerful, flexible moral theory.

Consequentialism and hedonism

2.5 Utilitarianism is made of two parts: (1) a theory about the structure of morality, and (2) a theory about the object or end of morality—that is, what morality is aiming at. Let's look at the first idea. Utilitarians hold that the only thing that matters for morality is the consequences of what you do. So part of the utilitarian creed is that consequentialism is the correct structure of morality. We can put it like this:

> *Consequentialism*: All that morally matters is the consequences of action.

It doesn't matter what you say, what you plan, what you intend, or what you tried to do. From the perspective of morality, all that matters is what you actually did. To find out whether you did the right thing or the wrong thing, all we need to do is look at the consequences of your action. In fact, what you should do is produce the best consequences you possibly

can. For utilitarians, the bar is set high: you are always obligated to do the best that you can. Here are the principles that lay out obligation, permissibility, and impermissibility under classical utilitarianism.

- If an action X has better consequences than any other action you could perform instead, then your duty (moral obligation) is to do X.
- If an action X has better consequences than any other action you could perform instead, then you are morally forbidden from doing any action other than X. Doing something else is the wrong thing to do.
- If actions X and Y have better consequences than any other action you could perform instead, and X does not have better consequences than Y, but Y does not have better consequences than X either, you are obligated to perform one of the actions, but it is morally permissible for you to pick either one.

The idea is this. At any given time you are faced with a range of possible actions that you might perform. You could keep reading this chapter, go for a walk, take a nap, get a coffee to help you make it to the end, all kinds of things. Which thing you decide to do should be whatever has the best consequences. What you should *not* do is something suboptimal, something that doesn't have the best results. Suppose there is a tie at the top—you could donate $10 to UNICEF or to Oxfam but not both; those choices are tied with each other and both are superior to all other options. In that case it is morally indifferent which you do. Just pick one.

2.6 Obviously, there are many questions that immediately arise. One is *consequences for whom*? When you're contemplating the possible outcomes of various choices, who should you be thinking about? Should you only care about consequences for yourself, or do other people count too? What about future people, or nonhuman animals? Only short-term consequences, or do you have to weigh the long-term as well? The utilitarian answer is simple: *you have to consider the consequences for everyone affected by your action, not just now, but indefinitely into the future*. Morality is not all about you. If you only had to be concerned with the consequences for you, then utilitarianism would devolve to egoism. One of the big differences with egoism is that, for utilitarians, everyone is on equal footing. You count for moral assessment, but you don't count extra.

2.7 You may be asking *what about motives and intent*? Surely that matters too; if one intentionally does something immoral, isn't it worse than someone who does the wrong thing by accident? Doing the wrong thing in the

heat of the moment, or just messing up, well, that may be bad, but planning to do evil, intending to do so with malicious motives and a wicked heart, that's just so much worse. So consequentialism can't be correct—there must be more to morality than just consequences alone.

2.8 The standard utilitarian response is to make a sharp division between blameworthiness/praiseworthiness and right action/wrong action. The rightness and wrongness of actions, utilitarians continue to insist, just has to do with consequences. But whether someone is blameworthy or praiseworthy for his or her action, that has a good deal to do with motive and intent. For example, suppose Johnny Missalot tries to shoot you. Fortunately for you, Johnny's such a lousy shot that he couldn't hit the ground with his hat. Now, clearly he did the right thing by missing. Missing his shot had better consequences than other actions he might have taken, like aiming a little to the left and actually shooting you. But he's certainly not praiseworthy for what he did, since he tried to shoot you, even though in one sense he did the right thing.

2.9 Or, to take another example, imagine you try really hard to do the right thing in some situation. You think things through, and make every possible attempt to do right. But suppose that you screw things up anyway and do the wrong thing. An example would be a case of "**friendly fire**":[1] a soldier who carefully follows orders to bomb a target and scrupulously aims his missiles, but still kills his comrades, who he did not know were in the target zone. A utilitarian would say that you still did the wrong thing (an action with suboptimal consequences compared with other actions you might have performed), but you might not be blameworthy for it. In fact, you might even deserve praise for having tried your best to do the right thing.

2.10 The next obvious question is what are good consequences and what are bad ones? In answering this question we come to the other key component of utilitarianism: a theory of the highest good (the ***summum bonum***)[2] or what the aim of morality really is. Utilitarianism as such isn't committed to any particular theory of the good, and we'll get different versions of the theory depending on what is named as the highest good. For example, some contemporary proponents of the view argue that preference satisfaction is the *summum bonum*. Classical utilitarianism is *hedonistic utilitarianism*, according to which the highest good is pleasure. We'll focus on this traditional view. Thus,

Consequentialism + the highest good is pleasure = hedonistic utilitarianism

What does it mean to say that the highest good is pleasure? Well, it means that pleasure is intrinsically valuable. It is an important, valuable thing to possess in its own right, and not because of something it will produce or provide for us later on. Moreover, nothing is more valuable than pleasure. So the point of morality is for each person to produce the greatest amount of pleasure in the world with each action they perform. That sounds a bit daunting, or excessively lofty, but the idea is modest: you should always do the best you can. And the best you can do in any situation is whatever action will produce the best balance of pleasure over pain for everyone your action affects (including you).

2.11 You might wonder what you should do when you're in a lousy situation and there aren't any pleasure-producing options available to you. Your car is hydroplaning in a storm and you could bring it to a stop either by rear-ending the motorcycle in front of you or steering it into the guardrail on your right and grinding along until friction slows you down. Neither choice is one that is going to be producing much pleasure. However, one is definitely worse: hitting the motorcycle. Turning into the guardrail will tear up your car, but hitting the motorcycle will kill its driver. The utilitarian judgment is to choose the lesser of two evils: hit the guardrail. The total amount of happiness in the world will be higher with that choice than it would be with the choice to hit the motorcycle.

Measuring pains and pleasures

2.12 You might think that all this sounds fine, but rather abstract. How are we supposed to measure the prospective pains and pleasures of our actions, so that we know which things to do and which to avoid? The utilitarian will first note that most of the time you don't need to spend a lot of thought on this issue. Mostly it will be rather obvious—you should use your hammer to hit the nail instead of hit yourself in the head. Did you really need to sit down with pen and paper and calculate the relative values of the prospective pains and pleasures for those choices? Of course, there are complicated questions. Should the death penalty be abolished? Is it morally proper to legalize marijuana? Is there anything wrong with assisted suicide? In these cases, we need to think things through.

2.13 In the eighteenth century, Jeremy Bentham, an early and influential utilitarian, proposed a way to figure out what the consequences of our actions will be, that is, how we can measure the future pains and pleasures

our actions might cause. Bentham's proposal was the *felicific calculus*. He claimed a pleasure or a pain may be measured by its

1. Intensity

   How powerful or intense is the pleasure or pain? Some pleasures are mild, like a tasty apple. Some are great, like the joy of graduation, or a wedding. Some pains are mild, like a papercut, others are strong, like a migraine headache.

2. Duration

   How long will the pain or pleasure last? Obviously, you want pains to be brief and pleasures to be long lasting.

3. Certainty or uncertainty

   How likely is it that the possible pain or pleasure that we're considering really will occur? An action with a high likelihood of pleasure to follow and a very low risk of pain looks like a better choice, all other things being equal, than an action with a low chance of pleasure and a high risk of pain. When you're playing cards, bet high on good hands.

4. Propinquity or remoteness

   How soon is the pain or pleasure? Is it going to happen right away, or is it years in the distance? The pleasures of education may be a long time coming—learning to play guitar is a slow process, and the joy of mastery is remote in time. The pleasure of an afternoon nap is imminent. The further away a sensation is, the more intervening factors there may be that prevent it, and so the less likely it is that it will ever happen.

5. Fecundity

   A sensation is fecund just in case it tends to be followed by the same type of sensation. For example, the pleasure of learning to read tends to lead to other pleasures, such as reading a good book. So the pleasure of learning to read is fecund. The pain of food poisoning often follows the unpleasantness of eating bad seafood, so the latter is a fecund pain. Clearly you would prefer your pleasures to be fecund and your pains not.

6. Purity

   A sensation is impure just in case it tends to be followed by the opposite type of sensation, otherwise it is pure. For example, drinking a lot of alcohol is an impure pleasure, since it tends to be followed by the pain of a hangover. Working out at the gym is an impure pain since it tends to be followed by

the pleasure of fitness. So you should want your pleasures to be pure and your pains impure.

7. Extent

   How many people will be affected by your action? To what extent will the pains or pleasures produced by your behavior spread out to other people? Those have to be taken into account and added up too.

Again, you needn't work through the felicific calculus every time you act. But it is there, waiting in the wings, for those problematic cases in which it's not obvious what the right action and the optimal consequences really are.

Quality and quantity

The root notion of utilitarianism is that we should act in such way as to maximize the quantity of pleasure in the world. You might be concerned that utilitarians make no mention of the *quality* of pleasures and pains. Indeed, Bentham was quite clear about that, writing, "Prejudice apart, the game of push-pin is of equal value with the arts and sciences of music and poetry. If the game of push-pin furnish more pleasure, it is more valuable than either" (Bentham, 1825, bk. III, ch. 1). Push-pin was a child's game much like tiddlywinks. For Bentham, it was mere snobbery to suppose that the pleasures of art museums, classical music, and fine literature are any better than cheap beer, horror movies, and NASCAR. The quantity of pleasure is all that matters, and it is just pompous moralizing to declaim that this or that pleasure is somehow superior in any way other than its amount. 2.14

Not all utilitarians have agreed with Bentham. Later, in the nineteenth century, John Stuart Mill tried to develop a way for utilitarianism to accommodate the idea that some pleasures are of higher quality than others, in a way that just measuring their quantity could not capture. Pleasures might be better or worse in some way besides mere amount. It may initially seem that one can't aim to maximize pleasure in general while at the same time maximizing high quality pleasures in particular. However, even under the fundamental position of hedonistic utilitarianism that our moral duty is to produce as much pleasure in the world as is possible by our actions, there is still room for promoting quality as well. Suppose that you could perform either an action X or an action Y, and both are superior to any other action you might do, but are tied with each other. Commonly utilitarians say that it is then morally indifferent which you do; as long as you do either X or 2.15

Y, you should just pick one. However, if quality matters too, then when there is a tie in quantity of pleasure produced, we ought to choose the action that produces the higher quality pleasure. The promotion of quality is far from innocuous. In fact, the idea that we should promote and appreciate higher quality pleasures is a substantive and radical proposal about how we ought to live.

2.16 Let us understand quality as the *density of pleasure per unit of delivery*. Consider two fishing trips. On fishing trip A you fish all day, pulling up one modest fish after the next. There is always something on the line, so you never get bored and there is always a little thrill. But at the same time you don't really catch anything particularly noteworthy. On fishing trip B you fish all day and only catch one fish—but it is a monster. It takes all your skill and cunning to boat the giant lunker, but you eventually do. It does not take much imagination to suppose that the total amount of pleasure attached to both fishing trips is the same; we can even suppose that the total weight of edible meat is identical. The quantity of pleasure associated with the string of fish from trip A is identical with the quantity of pleasure represented by the string of fish from trip B; it is just that there is only one fish on the string in the latter case. Other things being equal, A and B are equally good choices as far as the *quantity* of pleasure is concerned.

2.17 Trip B has one key thing going for it: the giant lunker. This is a higher quality fish than any of the ones caught on trip A, in fact that single fish is as good as the entire string from trip A. How should we understand this higher quality? Precisely as the density of pleasure: there is more pleasure concentrated in the lunker than in any of the other fish. This interpretation of quality well accords with our ordinary intuitions and once we start thinking about quality in this way, we can see that it is ubiquitous. However, the pursuit of high quality pleasures has its risks.

2.18 Suppose that Jane has $30 to spend on beer. Jane is debating whether to spend her $30 on two cases of Coors Extra Gold pilsener or one case of Pilsener Urquell. According to *The Beer Lover's Rating Guide* (Klein, 2000, p. 102), Coors Extra Gold is "sharp, light, and tasteless . . . it quickly subsides into a typical pedestrian brew, even on a summer picnic with cold cuts and salads. Touted as a 'full-bodied beer'—yes, in comparison to Coors' regular pilsener." On a scale of 0 to 5, Klein rates Coors Extra Gold 1.8, which means it is below average and suitable only for the extremely thirsty. Pilsener Urquell, on the other hand, Klein describes as "crisp, fresh, and mustily hoppy pleasant, understated aroma; intensely carbon-

ated; floral mouthfeel contains some bitterness, but it is subtle and well-calibrated; admirable textural strength; slides into tempered sweetness with spicy foods; a first-class beer to be enjoyed in multiples" (p. 242). Klein rates it 3.5, which is the middle of the above average range. If we assume that taste is objective, Klein is a competent judge of beer, and that Klein's rating system is linear, then Pilsener Urquell is about twice as good as Coors Extra Gold. Under these assumptions, Jane's choice is to buy two cases of Coors or half as much Pilsener Urquell, which tastes twice as good. The cost is the same, and the total quantity of pleasure to be produced is the same. How is Pilsener Urquell a higher quality beer than Coors? There is twice as much pleasure per bottle.

2.19 Let's stick with the beer example for a moment. As one becomes more informed and more expert about any subject—food, antiques, literature, tennis racquets, movies, travel, romantic trysts, jazz, or Platonic dialogues—one gains a finer appreciation for the high end while losing the ability to be satisfied with the low end. The recognition of quality comes at a cost. In the case of beer, a casual beer drinker will be more willing to knock back a corporate brew, and more likely to get a little pleasure out of it, than someone who consumes only cask-conditioned ales pulled from an English beer engine.

2.20 Suppose that Jane Pivo, a beer enthusiast, and Joe Sixpack, who is just enthusiastic, decide to drink beer together every night for a month. Their financial resources are limited, so they cannot afford artisanal craft beer every night. Most nights they will be forced to drink mass-produced beer, but once in awhile they splurge and drink the top-shelf stuff. Jane gets very little pleasure on the nights when they drink Rolling Rock Light and very great pleasure the evenings they share a **Brooklyn Black Ops** Imperial Russian stout[3] aged in bourbon barrels. Joe likes Rolling Rock just fine, although he is not a complete idiot and enjoys the Brooklyn a bit more. Their month of tasting can be presented graphically in Figure 2.1.

2.21 For the month, Jane totaled 300 units of pleasure and so did Joe. Thus, from a purely quantitative standpoint, it is no better to be informed and knowledgeable about beer than not. Jane received no more pleasure than did Joe over the course of the month. Joe's pleasure was more frequent and more evenly spread out, whereas Jane's beer-induced pleasure was rarer and more concentrated. The Millian view on quality is that we should live our lives like Jane Pivo—we should become knowledgeable about various pleasures, pursuing and promoting them. When confronted with two courses of action that produce the same quantity of pleasure, we ought to

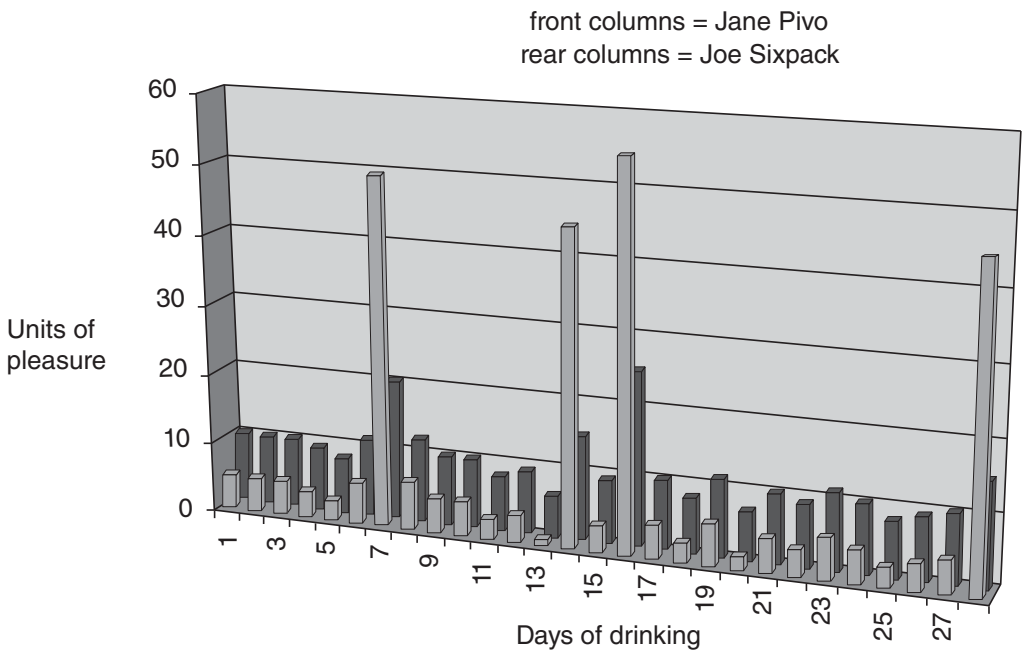

Figure 2.1 Two ways to experience the same total pleasure over a month of drinking beer

pursue the one with the higher quality, concentrated pleasures, even knowing that it is at the expense of enjoying lower-quality ones.

2.22 Mill's recommendation here should be tremendously controversial; it is not some innocuous, modest view that every utilitarian should obviously hold. Consider Jane. She might well wonder whether becoming a beer aficionado was worth it; after all she and Joe Sixpack drank all the same beers and on most nights Joe had a better time. Why isn't it perfectly reasonable for her to wish, as she sips an Old Milwaukee, that she could enjoy it as much as Joe?

2.23 One lesson here is that the appreciation of high quality pleasures is certainly worthwhile when the cost of gaining those pleasures is low. If Jane and Joe both had unlimited resources and could afford to drink only the finest beers every night, then clearly Jane is better off. She will get more pleasure out of each beer, and since she will never drink a low grade beer again, will end up with more total pleasure than Joe. Some pleasures are like this, even for the poor. Fine literature, for example, is in great abundance at public libraries and is available for low or no cost. There is no concomitant downside to learning to appreciate great literature, since it is free for the taking and in a supply greater than anyone could read in a

lifetime. With the advent of digital music files that are easily shared, music is becoming like literature, where the abundance of inexpensive music is so considerable that we are well advised to seek out and grasp the higher quality. Music and literature are a vast prix fixe buffet—there is no point in loading up on the Jell-O with mini marshmallows when one could have the lobster instead. In these contexts, the pursuit of high quality pleasures will lead to greater overall quantity, and the fundamental tenet of hedonistic utilitarianism is that we should perform those actions that produce as much pleasure as possible.

2.24 The controversy is in cases where either (1) the high quality pleasures are in short supply, or (2) they are expensive or difficult to obtain. In such instances one might prefer to remain in ignorance and not become sensitive to and appreciative of the subtle nuances that make for fine art, desirable first editions, Highland single malts, or super sports cars.

Objections to utilitarianism

2.25 There are six primary objections to utilitarianism, which will be addressed in roughly ascending order of seriousness.

2.26 *Objection 1: Practicality* The first objection is that one of the things we want out of morality is real guidance about what we should do and how we ought to act. Utilitarianism, though, is not a very practical ethical theory, since there is no way that we can perform the requisite calculations. We could make an educated guess or a decent calculation for the short term, but there is no way that we can predict all of the outcomes of our actions to the end of time, which is what the theory demands. If we don't know the ultimate result of a butterfly flapping its wings on the opposite side of the world, how can you possibly know whether some action will eventually lead to more pleasure than pain, or vice versa?

2.27 Utilitarians respond that no one said that doing the right thing was easy. Recall the discussion of praiseworthiness and blameworthiness earlier. You might select one college over another, there meet your eventual spouse, and go on to have children with that spouse. There is no way to tell in advance if one of your children will become a serial killer or the winner of the Nobel Peace Prize. But if your son becomes a killer, then you could reasonably judge that you should have gone to a different college, thus leading to a different spouse and children. You might have done the wrong thing, but you're not blameworthy for it. All we can do is the best we can; we have no

control over the final consequences of our actions. Utilitarianism tells us that tells us that *in fact* the right thing to do is whichever act maximizes the good in the world, and even if we are not sure how to hit that target, it is still what we should aim for.

2.28 *Objection 2: Invasiveness* The second objection is that, under utilitarianism, morality is just too invasive. Now every single aspect of our lives has moral weight. Whether you take out the garbage before or after dinner is now a moral issue. What you have for breakfast is laden with moral choices. You probably have a moral duty to get out on one side of the bed rather than the other. If getting out on the left side of the bed puts you that much closer to the bathroom, or your closet, or wherever you first go when you get up, then that's the side you *should* get up on. There's just a tiny bit less hassle in your life getting up on the left side of the bed, just a little bit less pain. So now it is your positive moral duty to get up on the left side of the bed. If someone gets up on the wrong side of the bed, well, that's no longer a figure of speech. Then might have gotten up on the morally wrong side. But that's crazy, goes the objection, morality has no business telling me how to get out of bed. Morality should be about the big issues—how we treat others, things like that.

2.29 Utilitarians reply that, yes, maybe you should take the garbage out after dinner and get up on the left side of the bed. But that's nothing to get too excited about; those are small potatoes sorts of actions. Every action has moral properties like every object has mass. Feathers aren't as likely to have much impact as bowling balls, but technically they have mass too. Utilitarianism shouldn't be seen as invasive, but merely comprehensive.

2.30 *Objection 3: Supererogation* The third objection is that under utilitarianism there is no such thing as supererogation. "Supererogation" refers to actions that are good actions, but greater than what duty requires. Recall the case of Private McGinnis, who threw himself on an Iraqi grenade to save his fellow soldiers. You may think that what he did was the very epitome of heroic, noble self-sacrifice, above and beyond the call of duty. Not utilitarians. If jumping on a grenade produced better consequences for everyone involved than any other action Pvt. McGinnis might have taken, then doing so was no more than his moral duty.

2.31 Utilitarians don't deny that, strictly speaking, there are no supererogatory acts. You are always obligated to perform the best action you possibly can. In some circumstances, like those of Pvt. McGinnis, doing the best

thing may be very difficult, or come with great personal sacrifice. In those cases it is quite reasonable to regard those who did the right thing as especially praiseworthy or admirable. Utilitarianism may mean that there are no supererogatory actions, but that does not mean that there are no morally heroic actions. There might be situations in which you're a hero just for doing your duty.

Objection 4: ***Simpson's Paradox***[4] The fourth objection is deeply puzzling, 2.32 
and some regard it as a showstopper for utilitarianism. The problem arises out of a statistical oddity called Simpson's Paradox. Simpson's Paradox is when a set can be partitioned into subsets that each have a property opposite to that of the superset. That sounds a little technical, but there are familiar examples. In the **2009 Wimbledon finals**,[5] Roger Federer beat 
Andy Roddick by a score of 5–7, 7–6 (8–6), 7–6 (7–5), 3–6, 16–14. Even though Roddick won most of the games (39 versus Federer's 38), he still lost the match. In 2003 the New York Yankees finished the regular season with 10 more wins than the Florida Marlins. The two teams met in the World Series and the Yankees outscored the Marlins 21–17 over the course of the series. Nevertheless, the Marlins won the World Series by four games to two. These cases are examples of Simpson's Paradox. The problem for utilitarianism is that we may be obligated to make every person alive less happy, because it will increase the total global amount of happiness. Consider the following two scenarios (Figures 2.2 and 2.3).

In Scenario 1, imagine that there are two people alone on a desert 2.33
island. It isn't a paradise; there's limited food, water, and shelter, and the two people have to struggle for survival. But suppose that nonetheless they are reasonably happy. Let's say that each person has a total of 100 units of

Figure 2.2 Two people on a desert island

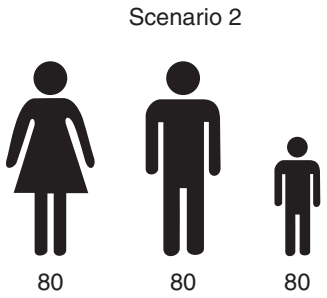

Figure 2.3 Three people on a desert island

happiness at the end of their life. The numbers don't matter; they're just placeholders to indicate relative values. Now, suppose that the couple is considering having a child, and creating Scenario 2. In this condition each adult is a bit less happy (remember that children make parents less happy), and they have to work that much harder to provide for their child. But the child has a fairly happy life, and the parents, while less happy, are still in the positive numbers for lifetime happiness.

2.34 Which is the morally preferable world according to utilitarianism? The answer is Scenario 2, because it is an overall happier world than Scenario 1, totaling 240 happiness units to 200. In this case, the couple on the desert island is morally obligated to create more people, even though it makes everyone there less happy. The desert island scenarios, while somewhat abstract, are not that far removed from reality. It is not hard to imagine that the entire planet is like the desert island, and that we might under utilitarianism be obligated to keep increasing the human population until we reach a tipping point, even if by doing so we make every living person less happy. This very counterintuitive result is a reason to reject utilitarianism.

2.35 *Objection 5: Agent-relative intuitions* The fifth objection to utilitarianism concerns, surprisingly, one of the initially attractive features of the theory, namely its respect for the principle of equal treatment. Utilitarianism is an agent-neutral moral theory, that is, one according to which everyone has the same duties and moral aims, no matter what their personal interests or interpersonal relationships. Theories like egoism are agent-relative, meaning that your moral obligations and goals may be completely different from mine. Consider the following two cases, which are designed to cut against agent-neutral intuitions.

Drowning swimmers. You can save one of two drowning swimmers but not both. You are the only means of rescue. One of the drowning swimmers is your child and the other is a stranger, but still a child with equal life prospects and a comparable network of family and friends. From the perspective of total happiness in the world it is morally indifferent which child you save. You may have a preference to save your own child, but in addition do you think you have a special duty to your own child. Are you obligated to save your own child?

A friend in need. Your best friend is down on her luck. Her husband left her and she is looking for work. She is having trouble making ends meet and you are considering buying some groceries for her to tide her over. As you enter the grocery store, you see a charity collecting outside to send money to aid famine victims in a distant land you have never visited. The charity claims that you will save lives through the donation of a few dollars. Not being terribly flush yourself, you can either buy groceries for your best friend or donate to the charity, but not both. Even though donating to the charity will produce more global happiness, do you think you have a duty to help your friend instead?

A utilitarian faced with the drowning swimmers case might choose his or her own child, but would have no obligation to do so. Or perhaps the utilitarian might decide that flipping a coin is the fairest way to decide who gets to live and who dies in that instance. The friend in need would just plain be out of luck. The distant charity should get the money. If you're inclined to think that those actions are not what you should do, then you might doubt that utilitarianism is the correct moral theory.

Objection 6: Nothing is absolutely wrong The final objection is that under utilitarianism there is no act so heinous, so terrible, that it is utterly unconscionable. In fact, think of the worst possible action you can, and there is an imaginable scenario under which it is your utilitarian moral duty to perform that very action. Think that no one should ever own slaves? Knowingly convict an innocent person of a crime? Kill the innocent? Torture political enemies? For utilitarians, all of those actions might be your moral duty, given the right set of circumstances. For instance, see what you think about this case. 2.36

The organ-robber. Imagine you are an attending physician in a busy emergency room. You're particularly worried this evening, because there has been

a train wreck, and not only are all ten victims your patients, but each of them needs an organ transplant—kidneys, livers, hearts, lungs, they each need something different. You've seen the wait list for new organs, and you know they're not going to make it. As you ponder this sad situation, Joe Klutz walks into the ER. Joe has fallen off of his ladder and has broken his arm. You set Joe's arm, and decide to run a couple of routine tests while he is there. You look at the test results and realize that, incredibly, Joe is a perfect donor match for all ten of your patients. You start thinking it over, and realize that if you just slipped Joe some chloroform, well, you could harvest his organs and save the lives of all ten of your patients.

Joe might live another 50 years without your intervention, but each of those patients would easily survive at least 20 years apiece with Joe's organs. 200 years of life versus 50: it's a utilitarian no-brainer, right? Sure, there would be family grieving for Joe, but there would be ten times as many people grieving for your other patients, should you not harvest Joe's organs. All in all, the world will be a much better place if you piecemeal Joe to save the lives of ten other people.

If you think that murdering an innocent person in order to cut them up and steal their body parts is wrong, no matter what good may come of it, then that is a reason to reject utilitarianism. Let's move on to consider the second major player among ethical theories, namely deontology, or Kantian ethics.

Deontology, or Kantianism (Is There an Absolute Moral Law?)

2.37 The last complaint against utilitarianism was that there are no actions that are absolutely morally forbidden. Every possible action—killing the innocent, rape, torture—could be your moral duty if the stakes are high enough. This goes against the intuition that some things are so terrible that it is always wrong to do them, no matter what the practical results. The German philosopher **Immanuel Kant** (1724–1804)[6] agreed, and maintained that the morality of actions does not vary from circumstance to circumstance, but instead there is an absolute moral law which applies to everyone at all times. The behavior that we owe to each other does not vary, and it is this idea that is behind the notion of moral rights. Of all the moral theories discussed so far, it is Kantianism alone that underwrites the possibility of you possessing moral rights.

Imperatives and good motives

Kant thought that utilitarians had things backwards. For utilitarians, things 2.38
like intent and motive are completely divorced from what makes an action a right action. Kant, on the other hand, argued that good intention, or a good will, is the only thing that is good without any qualification. In contrast, lots of character traits like courage, resolution, and constancy of purpose are not good or bad in themselves. Same thing for the gifts of fortune like power, fame, and wealth. It all depends what you use those talents and gifts to do. If you use your wealth and forceful personality to establish a religious cult whose sacraments are crack cocaine and kinky acts with the deacons, then maybe those weren't good personal traits to have. But noble motives and intentions, the desire to do the right thing, that's invariably good. Of course, we all know what road is paved with good intentions. Having a good will does not guarantee that one will do the right thing. Often our attempts to do the right thing are thwarted, or we're unlucky, or we act on bad information, or we just don't understand what the right thing to do really is. So exactly how can we distinguish having good intentions from having bad ones?

Kant's answer is that you have a good will if you try to do what's right, 2.39
if you try to follow the moral law. Unlike utilitarians, who think that you could accidentally do the right thing by bringing about the best overall consequences, for Kant you do the right thing only if you do it out of a sense of duty. Having a good will is the only thing good in itself. Good motives are part of what it is to do right; the other part is the successful following of the moral law. Thus:

> You have good motives + you follow the moral law = you do the right thing

What exactly is this moral law idea? It is a certain kind of imperative about what you should do. Of course, there are lots of imperatives, for example:

- If you want to pass the test, then you should study.
- If you're going to drink, then don't drive.
- If you can't make our meeting, then be sure to call.
- If you're planning to read Kant, then drink plenty of coffee.
- If you're the last one out, then you should turn off the lights.

Notice that all of these imperatives have a conditional structure, that is, an if . . . then . . . form. None of them tell you what you should do, come what may. You're not instructed to study, give up your keys, call, guzzle some coffee, or leave. These instructions tell you to do those things under certain conditions: if you want to pass, if you're going to drink, if you can't make our meeting, if you're going to read Kant, if you're the last one out. But maybe you're not going to do any of those things, in which case the imperatives don't have any force. They just don't apply to you. These if . . . then . . . instructions are *hypothetical imperatives*. Hypothetical as in "hypothetically speaking, if you were to scratch Sara's new car, then she would be very angry with you." None of that means that Sara is in fact angry with you.

2.40 All those hypothetical imperatives may be true of everyone, but for most people they are trivially true. Even for teetotalers it is true that "if you're going to drink, then don't drive"; it's just not a rule that affects their behavior. The moral law can't be one of these hypothetical imperatives, since it governs everyone's behavior, no matter what their own plans or personal situation may be. The moral law is therefore a *categorical imperative*—it tells you what you should do irrespective of idiosyncratic facts about you.

2.41 Kant gives a couple of different formulations of the categorical imperative, which he believed to be in some sense equivalent. Working out what Kant really meant is a task for Kant scholars. Here let's just take a look at the principles he gives.

Categorical imperative (version 1)

> The *categorical imperative* (version 1, universalizability): Act only according to those principles of action that you could will to be a universal law of nature.

2.42 The basic notion isn't as complicated as it first appears. When you were a kid, and wanted to do something your parents disapproved of, did they ever tell you, "what if everybody did that!?" If so, then your parents were closet Kantians.

2.43 Every time you do something, we can describe your reasons for acting in terms of a general principle, or maxim, of action. For example, suppose you're driving down the road, and you've just polished off a Big Mac, fries, and a Diet Coke from Mickey D's. You wad up your garbage and toss it out of the window. In this case, your principle of action is something like

whenever there is garbage in the car, get rid of it in the easiest way possible. This, of course, is when your mom pipes up with "what if everybody did that!?" She obviously doesn't mean to suggest that everyone could perform the exact action you just did since, well, you've already done it, and your car is free of trash. What she means is what if everyone acted in the same way for the same reasons you just did—what if your principle of action was made into a law that everyone followed, and everyone was chucking their garbage out of the window? You wouldn't want that. Kant is after the same idea as your mom; he just spells it out it more detail.

There are two ways that a principle of action can violate the categorical 2.44
imperative and thereby be a morally wrong principle to act upon.

Inconsistency Suppose that you have a serious gambling problem, and 2.45
you've already blown all of your money at the racetrack. But you're absolutely sure that Plato's Beard is going to win in the seventh race. So you go up to one of your friends and say, "Bob, can I borrow $50? I promise to pay you back." But in fact your plan is to go gamble with Bob's money. Actually, even if Plato's Beard does win, you know that you'll stay at the racetrack, betting on everything in sight, until you're broke again. You have no intention of paying Bob back. You probably find deceitful promise-making to be intuitively wrong. But just how does it come into conflict with the categorical imperative?

In this case, your principle of action is something like "promise to achieve 2.46
your own advantage, even when you know that you will not keep the promise." Let's universalize that. Now everyone acts on exactly the same principle. What will happen to promising? It will become meaningless; everyone will know that promises aren't worth spit. Which means that your false promise to Bob will get you nothing. Bob's not about to give you $50, knowing that promising is just some empty convention. In this case, universalizing makes your act of promising worthless. False promising is effective only against a backdrop of general honesty; if everyone is dishonest, then deceit won't work. In sum,

1. If your principle of action were universalized, then it would make your own action an impossible or fruitless one.
2. Thus your act could not be the result of a principle of action that you could will to be a universal law of nature.
3. Thus your act violates the categorical imperative.
4. Thus your act is immoral.

It is because false promising violates the categorical imperative that it is morally wrong.

2.47 *Inconsistent willing* Here's the second case. Suppose that no matter what, you always look out for yourself. If some politician wants to raise your taxes to provide health care to poor people, you'll vote them out of office. Whenever there's a food drive, the only canned goods you'll contribute are that tin of pickled beets that has been in the back of your pantry for two years. If somebody is poor, or sick, or uneducated, it's either his or her own fault or just bad luck. Either way it's not your problem. Everyone in this world has to take care of himself or herself, that's just the way it is.

2.48 We can characterize your principle of action in this case as "act selfishly." How does it stack up against the categorical imperative? Unlike the false promising case, act selfishly is a principle of action that is consistently universalizable. It could be the case that everyone always acts solely in his or her self-interest. However, "act selfishly" is not a principle that you can consistently *want*. It's easy to advocate selfish action when everything is going well for you, when you're young, healthy, strong, and have money in your pocket. But suppose that your car breaks down in the middle of a blizzard out in the boonies someplace and your cell phone's dead. If "act selfishly" is a universal law of nature, then no one is going to stop and help you. Then you'll be extremely sorry that everybody only cares about themselves. In other words, you can't consistently want that "act selfishly" be universalized. Here's the argument:

1. You could not consistently will that your principle of action be universalized.
2. Thus your act could not be the result of a principle of action that you could will to be a universal law of nature.
3. Thus your act violates the categorical imperative.
4. Thus your act is immoral.

Relentlessly selfish action is immoral because it violates the categorical imperative, as just discussed. It's another way to understand the violation of equal treatment complaint against ethical egoism—egoism is incompatible with the categorical imperative because it can't be universalized. Selfish action is appealing when you are on top of the world, otherwise, not so much.

Categorical imperative (version 2)

As mentioned, Kant gives another formulation of the categorical imperative too. Here it is. 2.49

> The *categorical imperative* (version 2, treating others): Treat other people as ends in themselves and never merely as means to your own ends.

What is this business about ends, or ends in themselves? The basic idea is pretty simple: treat other people with respect for their own goals, values, and interests; recognize the inherent dignity in others. In other words, don't just use people to get what you want. Immoral action comes from treating others as merely there for your use, just objects to be manipulated to your advantage, bodies to be stepped on as you climb the corporate ladder.

Sometimes people get the categorical imperative confused with the Golden Rule, which says to treat others as you would like to be treated. They are not quite the same principle, though, and differ in this important aspect: the Golden Rule assumes that everyone has the same aims and preferences. Your goals are those of your neighbor's too, so you should act in a way to further her achievement of those goals (because that's how you would like your neighbor to treat you). But, as was noted at the beginning of this chapter, people don't have the same preferences. The categorical imperative does not assume that everyone has the same values and interests, in fact it demands that we treat others with respect for *their own* goals, ones which may turn out to be radically different from or even dramatically opposed to ours. 2.50

Actions that we ordinarily take to be immoral are easily shown to be wrong under the categorical imperative. Consider theft, rape, killing, and fraud. All of these actions treat other people as merely as means to one's own ends. If you knock over a liquor store, you treat the store's owner and employees as just bodies who stock the shelves and fill the till. When you steal, you fail to treat them with respect for their own goals, values, and interests, which surely include the interests they have in retaining their property and remaining in business. Thus the categorical imperative offers a theoretical explanation of why those actions are wrong. 2.51

Here's a possible concern. Suppose you go into Wal-Mart to buy some tennis balls. You go up the cashier, hand her some money, get your receipt, and leave with a can of tennis balls. Presumably you just used the cashier 2.52

to get what you want, namely some tennis balls. She used you too; by servicing customers she gets to keep her job. Since the categorical imperative tells us that using people is wrong, purchasing tennis balls is immoral. It is a short step to showing that every kind of transaction must be morally wrong. A result that counterintuitive shows that the categorical imperative must be mistaken.

2.53 The proper response is to note the little word "merely" that appears in the categorical imperative. It is wrong to treat other people *merely* as a means to your own ends. But you and the cashier, while you do indeed use each other to further your own ends (getting tennis balls and getting a paycheck, respectively), do not use each other *merely* as a means to your own ends. After all, you did pay her for the balls, and she did give them to you. If you had stolen the balls, that would have been to treat the cashier without respect for her own goals and interests and would therefore have violated the categorical imperative and have been the wrong thing to do.

2.54 You might claim that a shopkeeper has a *right* not to have his or her goods stolen. In fact, the very idea of moral rights comes from Kantian ethics. None of the moral theories so far considered in this chapter—religious moralism, ethical egoism, moral relativism, or utilitarianism have any truck with the idea of moral rights. Utilitarians, for example, think that whatever situation one is in, the morally correct action is the one that produces the greatest amount of pleasure in the world. No one has a right to anything. In fact, that was one of the objections to utilitarianism, that no action, no matter how intuitively horrible, is ever absolutely forbidden.

2.55 Recall the organ-robber case. Under utilitarianism, it looked like the morally correct thing to do was to butcher Joe Klutz and redistribute his organs. If you thought that result was completely mistaken, it is likely that you thought that Joe had a right not to be a killed, and a right not to have his bodily organs stolen, no matter how noble the purpose. Under Kantianism, the moral law is universally applicable and exceptionless, and the organ-robber case treats Joe merely as a means, and not as an end-in-himself. Precious little respect for Joe's interests is shown by carving him up against his will, which is why it is wrong to do so. Kantianism thereby gives the intuitively correct answer in the organ-robber case.

2.56 The categorical imperative sets out a claim against the behavior of other people, that is, others are obliged to treat you in a certain way— they should treat you with respect for your own dignity and interests. Rights too are claims against the behavior of others. For example, you have a right not to

be killed, which means that you have a claim on the behavior of others that they refrain from killing you. This is a right that would hold against everyone at large. Other kinds of rights may only hold against specific individuals. For example, you may have a right that your priest keep the confidence of the confessional, even though you have no claim against others that they keep your secrets. The priest, however, promised confidentiality. We saw earlier how false promising violated the first formulation of the categorical imperative, which explains why the priest owes you silence. While a detailed discussion of different kinds of rights and how to understand them is beyond the scope of this chapter, it is worth noticing that the origin of the contemporary conception of rights is in the categorical imperative. If you think that people have moral rights that others must respect, then you might wish to side with the Kantians.

Objections to deontology

Objection 1: Generality The first objection to Kantianism focuses on the 2.57
first formulation of the categorical imperative, the idea that our principles of action must be consistently generalizable. The problem has to do with the right way to describe one's principle of action. Consider the organ-robber case again. All of the following are plausible candidates for being the operant principle of action.

1. If you want someone's bodily organs, just kill them and help yourself.
2. If you are a physician, you should save the lives of as many of your patients as possible.
3. You should act in such as way as to produce as much pleasure in the world as you can.

Number (1) fails the test of the categorical imperative. You couldn't consistently want everyone to act on that principle, not when needy patients start to stare hungrily at your young and healthy heart, liver, and lungs. Since the principle of action in (1) fails the categorical imperative, organ-robbing is immoral.

However, (2) apparently conforms to the categorical imperative. It does 2.58
not generate any inconsistencies to will that it be a universal law for physicians to save the lives of as many of their patients as possible. Nor does it

seem implausible that you could consistently will such a principle to be a universal law. In fact, some principle like this one underwrites the idea of triage—battlefield and ER doctors tend first to the patients most likely to survive only with immediate treatment at the expense of the less injured. While principle of action (2) leads to organ-robbing in unusual circumstances, it is seemingly compatible with the categorical imperative. Therefore organ-robbing is morally permissible.

2.59 Principle of action (3) is just the fundamental utilitarian directive. Is this principle of action something that can be universalized in accordance with the categorical imperative? Perhaps it can. Kantians will certainly reject the idea that morality is all about you, or all about me. Remember that earlier Kant explicitly argued that the egoist position of "act selfishly" violated the categorical imperative and was thus an immoral principle of action. So Kantians must be prepared to make some personal sacrifices in order to do the right thing, something obviously believed by utilitarians as well. One can't reasonably argue that (3) can't be universalized because you couldn't consistently will it to be a universal law when the time comes for you to sacrifice for the greater good. To do so is to slide towards egoism. But if (3) can be universalized, then organ-robbing in the case described is the right thing to do.

2.60 The objection to Kantianism is that each of three principles of action just listed can be cited as the principle of action in the organ-robber case. But (1) violates the categorical imperative, whereas (2) and (3) apparently do not. Therefore the categorical imperative leads to inconsistent moral judgments.

2.61 *Objection 2: Agent-neutral intuitions* One of the objections to utilitarianism was that it was incompatible with agent-relative intuitions. One of the selling points of Kantianism is that it *is* an agent-relative theory. All well and good. However, there are agent-neutral intuitions too, and those cut against Kantianism. Consider this case. You're a manager at a large company, and you're hiring for a new entry-level position. You know that your niece Sylvia is looking for a work, so you blithely decide that you'll give the job to her. There are other applicants for the job, equally as qualified as Sylvia.

2.62 Kantianism is an agent-relative moral theory according to which our interpersonal relationships can impose particular moral obligations that we do not have to others. Promising is an obvious example: you're obliged to give Tim a cup of coffee because you promised him one, although you're not obliged to give a cup of coffee to everyone who wants one. In the nepo-

tism case you may judge that it is wrong to dismiss the other qualified applicants for the position. Sylvia's hire was a foregone conclusion; the other applicants never really had a chance at all and suffer unfair opportunity costs by applying. If nepotism is wrong, then the morally correct thing to do would be to treat all of the candidates equally, without prejudice, so that any of them have a fair shot at being hired. Agent neutrality is the morally correct stance.

The drowning swimmers case prompted agent-relative intuitions (you should save your own child from drowning), and so was an objection against utilitarianism. Yet the nepotism case motivates agent-neutral intuitions (you should give all applicants an equal chance at the job) and so counts against Kantianism. 2.63

Objection 3: Horrible consequences The Kantian bumper sticker is damn the consequences, abide by the categorical imperative! Like false promising, Kantians think that lying fails the categorical imperative and is for that reason always wrong. But how about this case? You're at home watching the Cubs break your heart again when there is a frantic pounding on your door. You open it to find your friend Maria. She tells you that her crazy boyfriend is trying to kill her and begs you to hide her. You tell her to go hide in the bedroom, and you lock the door. A few minutes later there is more frantic knocking on the door. You open it and there is Dangerous Dan, holding a 10-inch combat knife with a tanto point and a serrated recurve. He has bloodlust in his eyes. He asks you if you've seen Maria. Do you say: 2.64

1. "Why, yes. She's in the bedroom."

or

2. "No, I have no idea where she is."

If you choose (2) and decide to lie, then you are using Dangerous Dan merely as a means to your own ends, in this case the end of protecting Maria. You are not demonstrating any respect for Dan's own values and goals (murderous though they are). Kantianism does not allow violations of the categorical imperative in order to prevent other violations of the categorical imperative. To do so would mean that the moral law is not truly categorical after all; instead it is just another hypothetical imperative, along

the lines of "treat others as ends in themselves and not merely as means unless by doing so you enable others to treat others merely as means." Not only does this move give up the absolute, categorical nature of the moral law, but it looks suspiciously utilitarian. If Kantianism just capitulates to utilitarianism when the going gets tough, then it's not really offering an alternative moral theory.

2.65 Thus according to Kantianism, your moral duty is to tell the truth. So in accordance with the categorical imperative, you tell Dangerous Dan that Maria's in the bedroom. If you're inclined to think that's completely absurd, and that clearly you should lie to protect Maria, then that's a reason to reject Kantianism. One could only wonder what would have happened to the Jews who worked at **Oskar Schindler's munitions factory** if Schindler had been a good Kantian and never lied to his Nazi bosses.[7] A stern and inflexible absolute moral law is difficult to square with the lavish and unforeseeable variety of human situations.

Virtue Ethics (Is Morality All about Having a Virtuous Character?)

2.66 The final ethical theory on deck is also the most ancient, endorsed by religions such as Islam and with a pedigree going back to Aristotle's discussion in *Nicomachean Ethics*. It is the idea that morality isn't about outcomes (like utilitarians think), or rule following (like Kant and Christians think), but about being a certain type of person. Instead of obsessing about good *actions*, we ought to focus on what it is to be a good *person*. This is the idea of virtue ethics. Virtue is an appealing way to understand morality. Instead of a daunting calculus that must take every actual and future person into account, or some rigid and abstract moral law, virtue is within the grasp of everyone. We can become the people we ought to be though the development of our own characters, without needing some God's eye perspective on the human condition. Morality ceases to be another imposition or just a bunch of rules you're supposed to follow, but instead naturally arises out of your emotional motivations working in harmony with rational reasons for acting. A virtuous person does good in the world because she wants to.

2.67 There are two central components of virtue ethics: the concept of virtue, and the concept of character. A good person is a virtuous person, one with a certain sort of character.

What is virtue?

Virtues are good qualities or characteristics. Some virtues are narrow ones, 2.68
specific to particular tasks or professions. For example, a good trait for football linemen is to be big, with bulging muscles and excellent short-term speed and power. Someone weighing 140 pounds with thin arms and a scrawny chest would get broken in two in the NFL. However, professional cyclists need the opposite qualities. Weighing 300 pounds with a muscular upper body is a serious vice for a cyclist, who would lose badly trying to cart all that weight up a mountain. Nor do cyclists need brief explosive speed nearly as much as they need steady endurance, which is largely unimportant for football players.

Moral virtues are those qualities of personality that are valuable for 2.69
everyone to have, whether they are an offensive tackle, cyclist, or anyone else. Possessing and acting on the virtues amounts to living a morally worthwhile and flourishing life for a human being. Moral virtue does not guarantee that your life will go well in the sense that you will be immune to bad luck, or you will never make mistakes. Messing up and suffering ill fortune is the stuff of tragedy. Yet being a virtuous person assures that you are doing the best that you can and that you are living the best life possible given whatever situation you are in.

Here is a partial list of typical moral virtues: loyalty, honesty, fair- 2.70
ness, kindness, courage, considerateness, civility, compassion, friendliness, patience, self-reliance, generosity, and dependability. The opposite of virtue is vice. According to Aristotle, many virtues are the midpoint, or "**golden mean**"[8] between related vices of deficiency or excess. Virtue is a sort of Goldilocks zone. Here are some examples.

Too little	*Just right*	*Too much*
Stingy	Generous	Wasteful
Cowardly	Courageous	Reckless
Cranky	Friendly	Sucking up/brown-nosing
False modesty	Honesty about oneself	Bragging
Anorexia	Moderate consumption	Gluttony
Sloth	Ambition	Workaholism

It is courageous to defend your country from an invading army, and cowardly to run from the battle. But it is also reckless or foolhardy for a soldier with a pistol to singlehandedly attack a platoon armed with machine guns.

It is good to be ambitious and bad to be lazy or slothful. Yet it is also a bad thing to be a workaholic who ignores all other valuable things in life to take as much overtime as possible and work seven days a week.

2.71 When the Roman statesman **Marcus Tullius Cicero** (106–43 BCE)[9] wrote that, "Extremism in the defense of liberty is no vice; moderation in the pursuit of justice is no virtue," this would have been startling to his audience. It would have been startling because Cicero was explicitly rejecting the Aristotelian idea that virtue is (generally) to be located at the golden mean, a traditional view that would have been well known among the Romans.

What is character?

2.72 For virtue theory, the virtues are not specific behaviors or actions so much as they are habits of character. No single action demonstrates much, if anything, about one's character. Rather, character is a tendency to act in certain sorts of ways. Someone with an honest character will routinely tell the truth, even when it may be unpopular or difficult to do so. That doesn't mean that an honest person absolutely never lies, but to do so is uncharacteristic of them, or goes against their native feelings. Character reflects a kind of steadiness of behavior.

2.73 Character is the sort of thing that comes in degrees. So someone with a strong character is dependable, steady, unflinching, unwavering, steadfast, and reliable. Someone with a weak character is fickle, weak, faithless, irresolute, erratic, capricious, and incontinent. Strength of character does not alone guarantee that a person is virtuous or honorable. A vicious person might be dependably wicked—the schoolyard bully is reliably cruel. To lead the morally good life you need to cultivate a character that is both strong and virtuous. Moreover, it is not enough to be a loyal friend but also stingy and cheap, or to be compassionate but a workaholic. A genuinely flourishing life requires *personal integrity*: a unity of the virtues, made habitual, and leading to action.

Objections to virtue ethics

2.74 *Objection 1: Virtue is compatible with evil* The first objection to virtue theory is that it seems entirely possible to cultivate and endorse the classic moral virtues and still participate in considerable wickedness. Consider Cosa Nostra, also known as the Mafia. When Salvatore Lo Piccolo, the capo

of the **Sicilian Mafia**, was arrested in 2007, the police found a written code of behavior for the mob.[10] Here's the Mafia's "**ten commandments**."

1. No one can present himself directly to another of our friends. There must be a third person to do it.
2. Never look at the wives of friends.
3. Never be seen with cops.
4. Don't go to pubs and clubs.
5. Always being available for Cosa Nostra is a duty—even if your wife is about to give birth.
6. Appointments must absolutely be respected.
7. Wives must be treated with respect.
8. When asked for any information, the answer must be the truth.
9. Money cannot be appropriated if it belongs to others or to other families.
10. People who can't be part of Cosa Nostra: anyone who has a close relative in the police, anyone with a two-timing relative in the family, anyone who behaves badly and doesn't hold to moral values.

Lo Piccolo's list emphasizes honesty, respect, duty, and moral values. Traditional Mafia morality also includes absolute loyalty to the family, the **omertà code of silence**,[11] and the prohibition on harming a made man. (Remember: a friend will help you move, but a true friend will help you move a body.) Despite the Mafia's promotion of such traditional virtues among its members, nevertheless it is a criminal organization that also engages in extortion and murder. You may think that a Mafioso who is scrupulous about never speaking to the police and is completely loyal to his bosses is doing the wrong thing. It would be better if he were a stool pigeon who ratted out his fellow criminals to the cops. Genovese family soldier **Joseph Valachi**,[12] who broke the code of omertà and squealed to the FBI and the US Senate on the inner workings of the Mafia, was a disloyal man who violated his own blood oath. Still, you might reasonably hold that doing so made him morally superior to his mob bosses.

The problem for virtue theory is that it looks like less "virtue" might result in better behavior in the Mafia case. A virtue theorist may rejoin that mafiosos only have some of the virtues, but clearly lack some major virtues as well. The good life consists in possessing all the virtues and 2.75

having an integrated, unified moral character. Having some virtues and some vices might lead to worse consequences than not having those virtues at all.

2.76 Unfortunately, if having virtues and vices is to be evaluated in terms of the behavioral consequences they lead to, then virtue ethics seems to devolve into just a cumbersome utilitarianism. When is loyalty good? When it leads to good consequences. When is remaining loyal bad? When it leads to bad consequences. Or you might see it as a form of Kantianism. Having a virtue is a good thing if it leads to respecting the categorical imperative, but a bad thing if it leads to rights violations. No trait of character is virtuous in itself—you always need to look at consequences to tell. But if the very same trait can be either virtuous or vicious, then it does not look like virtue ethics is bringing anything new to the table.

2.77 If virtue theory is to remain a distinctively different moral theory, it needs to prevent assimilation into either utilitarianism or Kantianism, while at the same time explaining how it is possible that having some good character traits might lead to a worse life or worse results for others than lacking those qualities.

2.78 *Objection 2: Clashing virtues* Virtue ethics advises that the good life consists in cultivating *all* the virtues. It is a good thing to be honest, but it is even better if you are a kind person too. The present objection is that some virtues apparently conflict with each other. For example, suppose that your girlfriend is showing off her new hairstyle. She asks, "How do you like my new haircut?" In truth, it's awful. You say,

1. Babe, that looks great!
2. Whoa! Did you use a weed whacker on that?

Option (1) is a much kinder and considerate response. But (2) is more honest. Which wins out, kindness or honesty? Here's another example. Imagine that your favorite indie band is now donating all of their profits to charity. That's pretty virtuous, right? Since they want to give as much to charity as possible, the band starts writing radio-friendly pop drivel and commercial jingles which earn a lot of money. In other words, **the band just sells out**.[13] Should you be disappointed? Selling out is surely a vice of rock bands. On the other hand they are giving piles of money to save baby seals. One more example: workaholism is a vice of excess, as discussed earlier. Suppose, though, that you're a workaholic to support your large and

otherwise impoverished family. Taking care of your family is a virtue. So are you doing the right thing or the wrong thing?

If the virtues can conflict with each other, as in these examples, then 2.79
virtue theory is offering no guidance about how we should live or what we should do. Should you be an honest person or a kind one? Should your band sell out or be true to your muse? There is no broader, overarching guidance other than the familiar "just be virtuous," which is no help in these conflict cases. It's like telling someone to just do the right thing when they are confronted with a moral dilemma and have no idea what the right thing is.

Objection 3: Relativism about virtues The third objection is that what 2.80
qualities of character count as virtues and which count as vices seems to be heavily dependent on the culture in which you are raised. For example, consider the Amish. The Amish are a religious sect that rejects modernity in order to live lives much like their eighteenth-century forebears. They emphasize humility, modesty, and plainness as cardinal virtues. In order to avoid their clothing from drawing attention to themselves, the Amish shun ostentation such as buttons, zippers, or even Velcro, fastening their clothes with straight pins and snaps. The clothing itself includes bonnets, long dresses, and capes in dark colors and body-covering styles.

Now contrast the Amish with the participants in **the World Naked** 2.81
Bike Ride (WNBR).[14] The WNBR is a series of organized bicycle rides held in major cities in countries all over the world. The participants ride nude in order to draw media attention to their cause, which is protesting oil dependency and promoting cycling as an environmentally friendly alterative to automobiles. The riders in the WNBR celebrate freedom, body confidence, healthy lifestyles, and environmental awareness as important virtues.

It is inconceivable for an Amish man or woman to ride in the World 2.82
Naked Bike Ride, just as it is impossible for a WNBR rider to endorse the virtues held by the Amish. So which are the right virtues to live by? Is it better to be modest, plain, and humble before God, or better to flamboyantly reject an uptight society that's addicted to oil?

Aristotle argued that **human beings have a function**,[15] and perhaps if 2.83
we understood the nature of this function that we could decide which virtues are best to adopt. Just as the function of the opposable thumb is to grasp, and the function of the eyes is to see, Aristotle thought that there must be a function, or a right way of functioning, for human beings. It

can't simply be to take in nutrition or to grow, since we share those traits with plants. Nor can it be merely having sense perceptions, or gratifying our base desires, since we share those urges with nonhuman animals. There must be some unique function for human beings that we alone have. Aristotle argued that reason and rational action are alone the province of humans. Canine happiness may come from answering the perennial question that dogs face: "Do I eat it, pee on it, or roll in it?" But human happiness, he thought, consists in the use of reason to guide our lives. Happiness is not the result of rational decision-making, but the process of it in accordance with virtue. Happiness is a journey, not a destination.

2.84 Whatever the merits of Aristotle's function argument, it is rather nonspecific about exactly which virtues contribute to the good human life. You can agree with him that the rational life is the proper function of a human being, but still aver that this does not yield much in the way of contentful virtues. So we are still stuck with the relativism problem.

2.85 The present objection is similar to the problem of clashing virtues, but with a twist—one might simply relativize virtues and vices to cultures. Thus humility and modesty are vices in the culture of WNBR cyclists while being virtues for the Amish, and vice versa. The problem with endorsing relativism about virtue is that virtue ethics is supposed to tell us those qualities of personality that are valuable for *everyone* to have. There may be different sporting virtues for offensive linemen than there are for riders in the Tour de France, but the traits that make for the good life, for being an honorable human being, are the same. At least, that's the promise of virtue ethics. Yet allowing relativism about virtue reneges on that promise. What's more, a relativized virtue ethics would then have to deal with all the objections to cultural relativism that were discussed earlier in the last chapter.

2.86 *Objection 4: There is no such thing as character* The final objection to virtue ethics comes from experimental psychology, and is perhaps the most powerful complaint. Essential to virtue ethics is the notion of character, the idea that each person has a tendency to act in certain sorts of ways given the appropriate situation. Thus an honest person will refrain from cheating or stealing, even in a situation when they could clearly get away with it. A gentle person won't harm others, even when they are told to do so, or when annoyed or made angry by another. Tendencies to behave in certain ways are *explained* by the fact that a person is honest, gentle, or has other char-

acter traits. Furthermore, those personality characteristics are *predictive*—knowing that a person is caring allows us to predict that she will help others who need a hand.

According to the experimental psychologists, "character" is just a story 2.87
that we tell about people after they act in certain ways. If someone demonstrates bravery in battle, we say that it was due to his or her courageous character. When participants in a prison simulation descend into barbarity we tell ourselves that their true natures were revealed. But the experimental evidence from psychology is that these categories are no more than convenient pigeonholes. People aren't "jerks" or "losers" or "alpha males" or "nature's nobility." Appeal to character is a tidy plotline to understand the behavior of others (and ourselves), but isn't nearly as predictive as the general facts of human nature that the experimentalists are uncovering.

What are these experimental results? John Doris, a prominent contem- 2.88
porary critic of virtue ethics, summarizes the argument as follows.

> Numerous studies have demonstrated that minor situational variations have powerful effects on helping **behavior**: hurried passersby step over a stricken person in their path, while unhurried passersby stop to help;[16] passersby who find a bit of change stop to help a woman who has dropped her papers, while passersby who are not similarly fortunate do not. Situations have also been shown to have a potent influence on harming: ordinary people are willing to torture a screaming victim at the polite request of an experimenter, or perpetuate all manner of imaginative cruelties while serving as guards in a **prison** simulation.[17] The experimental record suggests that situational factors are often better predictors of behavior than personal factors, and this impression is reinforced by careful examination of behavior outside the confines of the laboratory. In very many situations, it looks as though personality is less than robustly determinative of behavior. To put things crudely, people typically lack character. (Doris, 2002, p. 2)

The present objection is that there is no such thing as character. The concept of character does not adequately explain people's actions, nor can it be effectively used to predict what people will do. But if "acting in character" or "acting out of character" is just a kind of fictional narrative we tell about our actions, then there is no character for you to improve by trying to make virtues habitual. There is nothing for virtue ethics to be *about*.

Conclusion

2.89 Some moral theories enjoy currency among the general public, especially religious moralism, ethical egoism, and moral relativism. These were discussed in Chapter 1. Others are taken more seriously by professional philosophers, particularly utilitarianism, Kantianism, and virtue ethics. These were addressed in the present chapter. Contemporary philosophers have developed sophisticated versions of those theories in an attempt to address the criticisms of the basic models canvassed above. The two ethics chapters in this book have not discussed every ethical theory, but hopefully you see how moral theories are crafted and debated, and what the pros and cons are of the theories that have been presented. Moral reasoning is quite a bit more than just registering one's opinion on the hot-button topic of the day.

2.90 Many issues in moral philosophy are unexplored here, and are beyond the scope of a general introduction to philosophy. For example, what is the meaning of moral propositions? Is there some kind of objective moral reality to which they correspond? Or do such propositions do no more than express the approval or disapproval of the speaker towards certain sorts of actions? How are human ethical intuitions related to the moral or proto-moral instinctive actions of nonhuman animals? Can evolutionary psychology provide a unified account of these intuitive responses, and what would this mean for moral theory, or for the idea that there is a moral reality?

2.91 Finally, and this may be the hardest question to answer satisfactorily, why be moral? Clearly the response "because you should" misses the point. Such an answer presupposes that paying attention to morality matters and you ought to care about it, which is the very issue at hand. Perhaps you should only care about *appearing* to be a morally upright person, displaying public virtue yet preserving private, secret vice. The real question here is "Why should I adopt the moral point of view, why should I enter into the game of morality at all?" That question is a difficult one indeed.

Annotated Bibliography

Aristotle (350 BCE) *Nicomachean Ethics*, full text available at http://classics.mit.edu/Aristotle/nicomachaen.html, accessed May 8, 2012. Here Aristotle presents his seminal defense of virtue ethics. It remains the touchstone for virtue ethicists after nearly two and a half millennia.

Bentham, Jeremy (1789) *An Introduction to the Principles of Morals and Legislation* (Oxford: Clarendon Press, 1907), full text available at www.econlib.org/library/Bentham/bnthPML.html, accessed May 8, 2012. Bentham's classic presentation and defense of utilitarianism, including the felicific calculus.

Bentham, Jeremy (1825) *The Rationale of Reward* (London: John and H.L. Hunt), full text available at www.archive.org/stream/rationaleofrewar00bent/rationaleofrewar00bent_djvu.txt, accessed May 8, 2012. Contains Bentham's defense of the view that only the quantity of pleasure matters. He derides the self-appointed arbiters of good taste in defense of his claim that push-pin is as good as poetry.

Doris, John (2002) *Lack of Character: Personality and Moral Behavior* (Cambridge: Cambridge University Press). Doris argues against virtue ethics on the basis of "situationist" evidence from experimental psychology.

Kant, Immanuel (1785) *Groundwork for the Metaphysics of Morals* (Cambridge: Cambridge University Press, 1998), full text available at www.earlymoderntexts.com/pdf/kantgrou.pdf. 415 editions of this book were published in 20 languages between 1785 and 2008. It contains Kant's most famous presentation of deontology, including the categorical imperative.

Klein, Robert (2000) *The Beer Lover's Rating Guide*, 2nd edn (New York: Workman Publishing). Somewhat dated and eccentric, nevertheless a useful guide to good beer.

Mill, John Stuart (1863) *Utilitarianism* (London: Parker, Son, and Bourn), full text available at www.earlymoderntexts.com/pdf/millutil.pdf, accessed May 8, 2012. Like Bentham, Mill was interested in penal reform, and this is his presentation of utilitarian theory, including the idea of higher and lower pleasures.

Sidgwick, Henry (1874) *Methods of Ethics*, full text available at http://archive.org/details/methodsethics03sidggoog, accessed May 8, 2012. Sidgwick was, along with Bentham and Mill, the third great classical utilitarian. *Methods of Ethics* is the most systematic treatment of the theory prior to the twentieth century.

Suikkanen, Jussi (forthcoming) *This is Ethics: An Introduction* (Oxford: Wiley-Blackwell). A comprehensive introduction to ethics. It covers such topics as pleasure and happiness, well-being and the meaning of life, egoism and altruism, ethical theories and principles, moral motivation and reasoning, and how to apply ethics in ordinary life.

Online Resources

1 The definition and causes of friendly fire, along with historical examples: http://en.wikipedia.org/wiki/Friendly_fire

2 A dictionary entry on "summum bonum": http://en.wiktionary.org/wiki/summum_bonum
3 Beer Advocate's description and assessment of Brooklyn Brewery's Black Ops imperial stout: http://beeradvocate.com/beer/profile/45/40149
4 A sophisticated discussion of the mathematics underlying Simpson's Paradox: http://plato.stanford.edu/entries/paradox-simpson/
5 Highlights from the 2009 Wimbledon men's final: www.youtube.com/watch?v=rAu_woONfUs
6 A substantive overview of Kant's moral philosophy: http://plato.stanford.edu/entries/kant-moral/
7 The entry on Oskar Schindler at the United States Holocaust Memorial Museum: www.ushmm.org/museum/exhibit/focus/schindler/
8 A review of the concept of the golden mean in classical philosophy: http://en.wikipedia.org/wiki/Golden_mean_(philosophy)
9 The life, times, and philosophy of Cicero, the greatest orator of ancient Rome: http://plato.stanford.edu/entries/ancient-political/#CicRomRep
10 A news report on the arrest of Sicilian mob capo Salvatore Lo Piccolo and the Mafia's "ten commandments": http://news.bbc.co.uk/1/hi/world/europe/7086716.stm
11 A discussion of the meaning and origins of omertà, the honor code of silence: http://en.wikipedia.org/wiki/Omert%C3%A0
12 A substantial biography of famed mob informant Joseph Valachi: www.trutv.com/library/crime/gangsters_outlaws/mob_bosses/valachi//index_1.html
13 An overview of The Who's 1967 album *The Who Sell Out*: www.thewho.com/discography/index/album/albumId/62
14 The home page for The World Naked Bike Ride: http://worldnakedbikeride.org/
15 A more in-depth look at Aristotle's function argument and conception of the human good: http://plato.stanford.edu/entries/aristotle-ethics/#HumGooFunArg
16 An overview of the behavioral psychology's good Samaritan experiment: www.experiment-resources.com/helping-behavior.html
17 The Stanford prison experiment is presented and discussed in a video: www.youtube.com/watch?v=sZwfNs1pqG0&feature=related