

Chapter 2

Why Traditional Approaches to Ethics Won't Save You

Imagine that you are standing on a footbridge spanning some trolley tracks (see [figure 3](#)). You see that a runaway trolley is threatening to kill five people. Standing next to you, in between the oncoming trolley and the five people, is a railway worker wearing a large backpack. You quickly realize that the only way to save the people is to push this man off the bridge and onto the tracks below. The man will die, but his body will stop the trolley from reaching the others. (You quickly understand that you can't jump yourself because you aren't carrying enough weight to stop the trolley, and there's no time to put on the man's backpack.) Legal concerns aside, would it be ethical for you to save the five people by pushing this stranger to his death? We have just described a very famous philosophy problem known as the "footbridge dilemma."¹ It is often used to contrast two different normative approaches to ethical decision making: a *consequentialist* approach and a *deontological* approach. A *consequentialist* approach is one that determines the morality of an action by its ensuing consequences. Utilitarianism, a common form of consequentialialism, is often described by the phrase "doing the greatest good for the greatest number of people." A very different form of ethical thinking, what Immanuel Kant referred to as *deontological* approach, judges the morality of an action based on the action's adherence to rules or duties.² Kant argued that judgments of whether an act is right or wrong should be determined by a consideration of rights and duties in society. From Kant's point of view, the act of pushing someone off of a bridge would violate his rights and is therefore immoral.

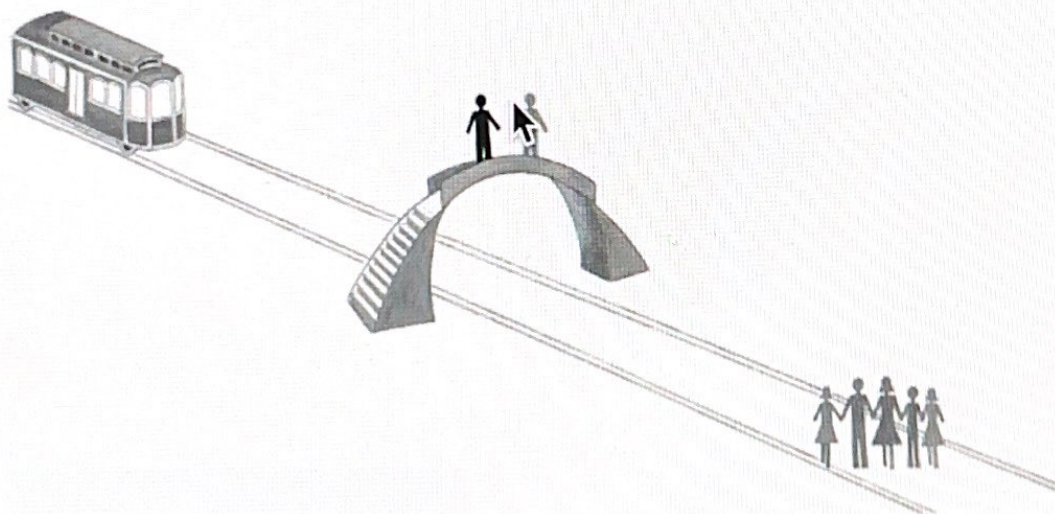


Figure 3. The footbridge dilemma

Indeed, when reading the footbridge dilemma, most people do not believe it is ethically acceptable to push the railway worker off the bridge in order to save five lives. Using a deontological approach, they ask themselves whether they have the right to push someone off of a bridge. If you ask them why they are opposed to the idea of pushing the man off the bridge, common answers include, “That would be murder!” “The ends don’t justify the means!” or “People have rights!”³ By contrast, a utilitarian approach would involve adding up the costs and benefits of each choice and choosing the option that yields the best balance of costs and benefits for all involved—which, in this case, would be to save five lives at the expense of one.

Now let’s look at a problem that was conceived before the footbridge dilemma, the “trolley dilemma”: A runaway trolley is headed for five railway workmen who will be killed if it proceeds on its present course (see figure 4). The only way to save these people is to hit a switch that will turn the trolley onto a side track where it will run over and kill one workman instead of five. Ignoring legal concerns, would it be ethically acceptable for you to turn the trolley by hitting the switch in order to save five people at the expense of one person?⁴

When considering the trolley dilemma, most people (who have not previously been exposed to the footbridge dilemma) say that it is ethically permissible to hit the switch. If you ask them why, their explanations tend to focus on the belief that having five people die would be worse than having one person die.⁵ This is prototypical *utilitarian* thinking because of its focus on the consequences of actions.

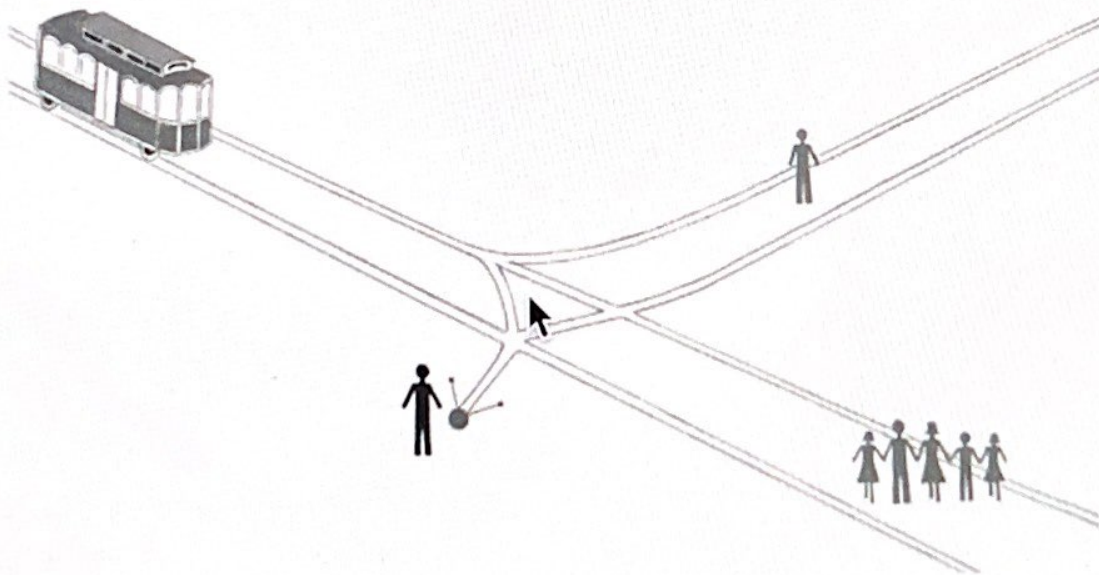


Figure 4. The trolley (switch) problem

When people are exposed to both of these problems, some are bothered by the arguable inconsistency of deciding to flip the switch to turn the trolley (in the trolley dilemma), contrasted with the decision not to push the man over the bridge (in the footbridge dilemma).

Those who are bothered by the inconsistency tend to make the footbridge decision intuitively; later exposure to the trolley dilemma then leads them to greater reflection consistent with utilitarian reasoning. As these two stories illustrate, we sometimes use the implied philosophical principles discussed earlier to make judgments. However, we tend to apply these rules inconsistently, and we sometimes violate what we would do if we gave the question more thought.

We have no vested stake in whether you are more of a utilitarian or a deontologist, or if you decide to become one or the other upon finishing this book. You are welcome to your own opinion about what to do in the footbridge and trolley problems. Our aim is simply to alert readers to potential inconsistencies in their decisions and actions—and, in particular, to the gap that exists between their behavior and their perceptions of their behavior, a gap that traditional ethical approaches tend to ignore.

Can Ethicists Improve Our Ethics?

After the collapse of Enron and other organizations during the same period, professional schools and corporations have been called on to give ethical matters more serious deliberation. As society demands more ethical behavior from organizations, it is useful to examine whether traditional ethical analysis offers a promising solution. The first logical step in this process is to examine what ethicists have to offer to the issue. Within the field of philosophy, ethics has historically been studied from a normative perspective—that is, an approach that seeks to determine the morally correct course of action. This type of perspective focuses on asking the question, “How should people behave?” Philosophers have considered, for example, whether a utilitarian or deontological approach to the footbridge dilemma is more appropriate.

Contemporary philosophers have argued that philosophical thinking is central to a moral education, that it will make us better citizens, and that it will give us the courage needed to stand up for justice.⁶ Yet legal scholar and judge Richard Posner argues that no empirical evidence exists to support these claims.⁷ In fact, ethicists themselves provide the perfect sample to test whether traditional, normative training in ethics leads to more ethical behavior, notes Eric Schwitzgebel, a philosophy professor at the University of California at Riverside. Because ethicists devote their careers to studying and teaching morality, we might expect ethicists to behave more ethically than the rest of us.⁸

Yet in his research, Schwitzgebel finds that if morality is equated with “not stealing,” ethicists do not score very well, at least by certain measures. Surveying thirty-one leading academic libraries in the United States and the United Kingdom, Schwitzgebel compared the rate at which ethics books were missing from the shelves to the rate at which nonethics books in philosophy, comparable in age and popularity, were missing. He found that ethics books were more likely to be missing than nonethics books. Next, he examined the presence of fairly obscure philosophical texts that would likely be borrowed only by advanced students and

professors. Among those texts, he found that a philosophy book was 50 to 150 percent more likely to be missing if it was an ethics book than if it was a nonethics book.

Schwitzgebel conducted related research on whether ethicists are more likely to engage in the arguably prosocial behaviors of voting and not eating meat. Comparing philosophical ethicists with both other philosophers and with professors in other departments, he found that while ethicists were more likely to condemn meat-eating than were other nonethicist philosophers and other academics, they were no less likely to eat meat themselves. Across other contexts, the researcher found little support for the notion that traditional ethics training creates more ethical citizens.⁹ Schwitzgebel concluded that his research undercuts the widespread assumption that enrollment in ethics courses will improve students' future ethical behavior.¹⁰

Even professional philosophers appear to be divided regarding the ethical behavior of ethicists. A poll of philosophers at an American Philosophical Association meeting in April 2007 found that although a substantial minority (especially ethicists) expressed the view that ethicists do behave morally better, on average, than nonethicists of a similar social background, a majority of respondents said that ethicists do *not* behave better than nonethicists.¹¹

Surprised? You might not be if you thought about the focus and underlying assumptions of a philosophical approach to ethics. Normative ethicists from a philosophical tradition have focused on exploring how we *should* behave and have made great strides toward answering these types of questions. However, little empirical attention has been devoted to examining how people actually *do* behave and how their ethical behavior can be improved—knowledge that is needed to understand and improve not just how philosophers behave, but also how the ethical and economic crises of the past decade emerged. As we will discuss later in this book, how we think we *should* behave is very different from how we *want* to behave. We may predict we will behave in a manner consistent with our expectations for ourselves. But when the time comes to make a decision, we often behave the way we want to behave.

The Limits of Traditional Approaches to Ethics

Another barrier that has kept scholars of ethics from fully dealing with ethical issues concerns the central role they give to decision makers' ethical intentions. Most approaches to ethics assume that people recognize an ethical dilemma for what it is and respond to it intentionally. By contrast, research on *bounded ethicality* examines unethical behavior that arises without intentionality. Consider J. R. Rest's influential descriptive model of ethical decision making. Rest claims that individuals faced with ethical decisions go through the following four phases:¹²

Moral Awareness → Moral Judgment → Moral Intention → Moral Action

Moral awareness, judgment, intention, and action certainly are important factors in understanding many ethical decisions. Yet this model is incomplete and potentially misleading. The model presumes that (1) awareness is needed for a decision to have moral implications, (2) an individual's reasoning determines judgment, and (3) moral intention is required for her to understand her moral action. Each of these assumptions, which are implicit in traditional approaches to ethics and many ethical training programs, ignores evidence to the contrary. In doing so, the model directs our attention away from critical elements of decision making and judgment that lead to unethical behavior. As we explain in the sections that follow, those who teach us to behave more ethically neglect many of the situations in which we actually find ourselves, including those where we lack moral awareness, judge before reasoning, and misjudge moral intention.

When We Lack Moral Awareness

Imagine that you are a salesperson who works on full commission. All of your income depends on how much you sell. You have been given aggressive sales quotas, and you focus on how to meet these goals. At the end of the year, you accomplish these goals and are rewarded generously by your company.

This scenario describes the situation faced by millions of employees working for organizations around the world. At face value, the situation appears completely acceptable. Now let's add more detail. The year is 2006, the salesperson is a mortgage lender, and his quotas require him to lend money to homeowners independent of their ability to pay. Or imagine that the salesperson works at Enron in 1999, selling a new concept: the firm's "special purpose entities," or shell firms—ways to hide its debt—to private equity investors such as JPMorgan Chase, Citigroup, Credit Suisse First Boston, and Wachovia.

While some of the salespeople at these companies were probably aware of the ethical consequences of their decisions, many more were probably unaware. They may have viewed them as "business decisions" and believed they were following accepted practices to achieve the ultimate business goal of making money. Or they may have seen them as legal decisions and asked themselves merely whether the sales strategies they followed were technically legal. It's likely that most didn't, however, view these decisions as ethical ones.

Training in business ethics tends to be largely based on the approaches to ethics described above: that is, emphasizing the moral components of decisions with the goal of encouraging executives to choose the moral path. But the common assumption that this training is based on—that executives make explicit trade-offs between behaving ethically and earning profits for their organizations—is too narrow. It ignores the fact that decision makers often fail to see the "ethics" in a given ethical dilemma. In many situations, decision makers do not recognize the need to apply the type of ethical judgment they may have learned in ethics training courses to their decision-making process.

As we described in [Chapter 1](#), our minds are subject to *bounded ethicality*, or cognitive

limitations that can make us unaware of the moral implications of our decisions. The outside world also limits our ability to see the ethical dimensions of particular decisions. For example, aspects of everyday work life—including goals, rewards, compliance systems, and informal pressures—can contribute to *ethical fading*, a process by which ethical dimensions are eliminated from a decision.¹³ Ann and her colleague Dave Messick have argued that these common features of organizations can blind us to the ethical implications of a decision, leading us, for example, to classify a decision as a “business decision” rather than an “ethical decision” and thus increasing the likelihood that we will behave unethically.

The organizational practices that contribute to ethical fading may be as subtle as differences in the language used to describe the decision. A case in point is Albert Speer, one of Adolf Hitler’s government ministers and most trusted advisers. After the war, Speer admitted that by labeling himself an “administrator” of Hitler’s plan, he convinced himself that issues relating to human beings were not part of his job.¹⁴ This labeling allowed Speer to reclassify ethical decisions as business decisions, such that the ethical dimensions faded from the decision.

Why does our classification of a decision matter? Because classification often affects the decisions that follow. When we fail to recognize a decision as an ethical one, whether due to our own cognitive limitations or because external forces cause ethical fading, this failure may very well affect how we analyze the decision and steer us toward unintended, unethical behavior.

When We Judge before Reasoning

Consider the following two stories:

- A woman is cleaning out her closet, and she finds her old U.S. flag. She doesn’t want the flag anymore, so she cuts it up into pieces and uses the rags to clean her bathroom.
- A family’s dog is killed by a car in front of their house. They have heard that dog meat is delicious, so they cut up the dog’s body and cook it and eat it for dinner.

When psychologist Jonathan Haidt and his colleagues presented these stories to study participants, most of them immediately decried the behaviors depicted to be wrong, though they couldn’t immediately offer informative explanations for their opinions.¹⁵ Instead, they responded with statements such as, “I don’t know, I can’t explain it, I just know.”¹⁶

Intuitionist psychologists such as Haidt argue that such emotional reactions precede moral judgment. In other words, moral reasoning doesn’t influence moral judgment. Rather, it’s the other way around: moral judgment influences moral reasoning. According to this view, quick, emotional reactions drive our judgments, and it is only after making such judgments that we engage in deliberate moral reasoning to justify our initial reaction.¹⁷ These emotive processes generate our initial verdicts about ethical issues, such as the use of the flag and consumption of

dogs. In direct contradiction of the Rest model of ethical decision making, *only after reaching these verdicts do we come up with reasons to explain them.*

The strong influence of emotional reactions on moral judgment is supported by research showing that individuals with neurological damage to the regions of the brain responsible for emotion have a reduced capacity for moral judgment and behavior.¹⁸ These findings cast doubt on the notion that judgment always precedes action, a premise that has dominated traditional approaches to ethics.

When We Misjudge Moral Intention

Traditional philosophical approaches to ethics, particularly certain segments of deontological ethics, place intention as a central consideration in judgments of unethical behavior. That is, when judging an individual's ethicality, these approaches consider whether or not the person intended to behave ethically. But consider that judgments of intentions can be based on erroneous factors, as this example developed by Yale University philosophy professor Joshua Knobe illustrates:¹⁹

The chairman of a company has to decide whether to adopt a new program. It would increase profits and help the environment too. "I don't care at all about helping the environment," the chairman says. "I just want to make as much profit as I can. Let's start the new program." Would you say that the chairman intended to help the environment?

Now consider a variation on the situation.

The chairman has decided to adopt a new program. The program would increase profits but harm the environment. "I don't care at all about helping the environment," the chairman says. "I just want to make as much profit as I can. Let's start the new program." Would you say that the chairman intended to harm the environment?

Despite the fact that, in both scenarios, the chairman's only goal is to make money, people's judgments of the chairman's intention seem to depend on the "side effect" of the chairman's decision. After study respondents read the first scenario, in which the program led to environmental improvements, only 23 percent said the chairman had intentionally helped the environment. By contrast, after respondents read the second scenario, 82 percent believed the chairman had intentionally harmed the environment. This was true despite the fact that the chairman expressed identical intentions in both scenarios.

Such inconsistencies are driven by factors irrelevant to a decision maker's intentions. As such, they cast doubt on approaches that make intentionality a defining characteristic of ethical versus unethical behavior. It's true that intentionality can drive responses to ethical behavior, but that is not true for all situations and all decisions. As Ann and her colleague Kristin Smith-

Crowe have argued, “‘good’ and ‘bad’ people make ‘good’ and ‘bad’ decisions.” Therefore it’s important to be able to identify and understand intentional *and* unintentional ethical decisions. Traditional approaches neglect the latter.

The variables that the Rest model of ethical decision making encapsulates are important. But some of the model’s elements—moral awareness, a set order of stages, and intentionality—obscure key factors that lead to unethical behaviors in organizational life. By ignoring ethical decisions that occur without moral awareness, the model leaves a substantial portion of unethical decisions, and the reasons behind them, unexamined.

A subset of a new school of philosophy that examines actual behavior, known as *experimental philosophy*, responds to some of our criticisms. These philosophical rebels, in hopes of shedding light on traditional philosophical issues, run experiments to gather information about the judgments that people actually make when faced with moral dilemmas.²⁰ This type of research should provide valuable information on how people actually behave. At this stage, however, it represents a small and somewhat isolated subset of philosophers whose work has yet to affect traditional approaches to ethics training.

If philosophical approaches don’t provide all of the keys needed to reduce unethical behavior in organizations, what will? Unlocking the door to unethical behavior requires an insightful understanding of the subtle influences on our behavior—influences of which we are often unaware—and their impact on how we think about ethical dilemmas. In the next section and chapters that follow, we describe how psychologists are applying those insights to the burgeoning field of behavioral ethics.

Two Cognitive Systems, Two Modes of Decision Making

The field of behavioral ethics emphasizes the need to consider how individuals actually make decisions rather than how they would make decisions in an ideal world. Research reveals that our minds have two distinct modes of decision making. By understanding these modes, we can reach key insights to help improve the ethicality of our decisions.

Not surprisingly, decision making tends to be most ethically compromised when our minds are overloaded. The busier you are at work, for example, the less likely you will be to notice when a colleague cuts ethical corners or when you yourself go over the line. An important psychological concept sheds light on why this tends to be the case: the distinction between “System 1” and “System 2” thinking.²¹ According to this view, System 1 thinking is our intuitive system of processing information: fast, automatic, effortless, implicit, and emotional. System 1 is also efficient, and thus serves as an appropriate tool for the vast majority of decisions we make on a daily basis. By comparison, System 2 thinking is slower, conscious, effortful, explicit, and more logical.²² When you weigh the costs and benefits of alternative courses of action in a systematic and organized manner, you are engaging in System 2 thinking.

It is quite common for people to have emotional, System 1 responses to ethical problems. However, such responses are sometimes at odds with what we would have decided with more

deliberation. Moreover, the importance of real-world decisions does not necessarily protect us from the limits of the human mind. In fact, the frantic pace of modern life can lead us to rely on System 1 thinking even when System 2 thinking is warranted.²³ In one study, researchers found that “cognitively busy” study participants were more likely to cheat on a task than were less overloaded study participants.²⁴ Why? Because it takes cognitive energy to be reflective enough to stop one’s impulse to cheat. Kern and Chugh found that the impact of outside influences on our ethical choices—such as whether the same outcome is framed as a loss or a gain—depends on how much time we have to make the decision.²⁵ They asked people to imagine themselves in the following situation:

You are trying to sell your stereo to raise money for an upcoming trip overseas. The stereo works great, and an audiophile friend tells you that if he were in the market for stereo equipment (which he isn’t), he’d give you \$500 for it. You don’t have a lot of time before you leave for your trip. Your friend advises that you have a 25% chance of getting the sale before you leave for your trip. [A separate group was told that they would have a 75% chance of losing the sale.] A few days later, the first potential buyer comes to see the stereo and seems interested. The potential buyer asks if you have any other offers. How likely are you to respond by saying that you do have another offer?

As in other research by Kern and Chugh, study participants were more willing to cheat to avoid losses (“losing the sale”) than to accrue gains (“getting the sale”). However, the framing as a loss or a gain only affected decision making when individuals were under time pressure and told to respond as quickly as they could. Individuals who were under no such pressure—those who were told to take their time in responding and to think carefully about the question—were not affected by the irrelevant framing (irrelevant to the choice as an ethical one) of a potential gain versus a potential loss.

Which way of thinking is better, System 1 or System 2? Many people were pleased when author Malcolm Gladwell made a case for trusting our intuition in his book *Blink*. We like to go with our gut (System 1). Moreover, System 1 thinking is sufficient for most decisions; it would be a waste of time to logically think through every choice we make while shopping for groceries, for instance. However, trusting our gut instinctively, without ever employing System 2 thinking, can widen the gap between how we want to behave and how we actually behave. System 2 logic should be part of our most important decisions, including those with ethical import.

If your gut reaction is different from the decision you reach after more deliberative processing, it is important to reconcile this inconsistency. If you let your gut rule, something as simple as whether a choice is framed as a gain or a loss might influence a decision. But if you ignore your gut and completely base your decision only on a cold calculation of the costs and benefits, you may be ignoring internal warning signs that “something isn’t right,” such as the omission of the decision’s ethical implications from the calculation—the type of signs to which those who contributed to the financial crisis of 2008 should have listened. It’s important to get the two systems to talk to each other.

Essentially, when the two systems disagree, that is your hint to have each system “audit” the other system. Your gut can help you figure out what feelings you may have left out of your careful calculation, and rational analysis may help you determine whether irrelevant factors are influencing your gut response.

The Importance of Ethical Self-Awareness

As evidenced by the research findings presented in this chapter, people generally fail to recognize that their ethical judgments are biased in ways they would condemn with greater awareness. Unfortunately, informing us about our biases doesn’t seem to help us make better choices. We tend to believe that while others may fall prey to such inconsistencies, we ourselves are immune to them. For example, when participants in one study were asked to predict whether financial incentives would influence their own and others’ decisions to donate blood, most overestimated the influence of self-interest on others; at the same time, they denied it would affect their own decision.²⁶ Most of us dramatically underestimate the degree to which our behavior is affected by incentives and other situational factors.

The decisions we make on behalf of ourselves, our organizations, and society at large can create great harm. To improve our ethical judgment, we need to understand and accept the limitations of the human mind. Yet the solutions that have been offered to reduce the undesirable outcomes of these decisions—including laws and ethics remediation and training—don’t take these limitations into account. Without an awareness of blind spots, traditional approaches to ethics won’t be particularly useful in improving behavior. If, like most people, you routinely fail to recognize the ethical components of decisions, succumb to common cognitive biases, and think you behave more unethically than you actually do, then being taught which ethical judgment you should make is unlikely to improve your ethicality. By contrast, the lessons of behavioral ethics should prove useful for those who wish to be more ethical human beings but whose judgments don’t always live up to their ideals or expectations.