FAMILY, COMMUNITY, TROLLEY PROBLEMS, AND THE CRISIS IN MORAL PSYCHOLOGY

PAUL BLOOM

A young woman meets a much younger man and takes him into her home. He suffers from terrible limitations. He cannot walk or talk or even sit up; he cannot be left alone and must be carefully fed. He often needs attention at night, and she spends the first years with him in a sleep-deprived fog. Still, this is the most important relationship of her life. She would die for him. She spends many years nursing him as he gradually becomes able to walk, to toilet himself, and to express and understand speech. After they have been together for over a decade, he becomes interested in other women and begins to date, and eventually he leaves her home and marries someone else. The woman continues to love and support him, helping to raise the children that he has with his new wife.

If this younger man were a grown stranger off the street, the woman's actions would be seen as saintly or insane. But, of course, this description summarizes a typical relationship between mother and son. In some regards, this makes her sacrifice all the more impressive, because now we can add additional considerations – if he's not adopted, she keeps him inside her body for nine months, suffering pain, nausea, and exhaustion. Then she gives birth, an act

that is terribly painful and carries certain physical risks. She might then feed him from her own body for months or years afterwards.

Knowing that they are mother and son changes how we think of the woman's actions. The point of this story, told by Alison Gopnik in *The Philosophical Baby*, is that family is special. Knowing that this is her son transforms how we think of the woman's obligations. If she were indifferent towards her child, unwilling to sacrifice for it, treating it just as she would a stranger, this would be judged to be immoral, repellently so. We feel the same, though perhaps not to same degree, when the parent is a father instead of a mother, and when the baby is adopted rather than biological.

These observations illustrate a problem in contemporary moral psychology, which is the field that explores the nature of moral judgment and moral action, including empathetic responses to the pain of others, altruistic behavior, the so-called moral emotions, such as guilt, shame, gratitude, and anger, and considered judgments about what's morally obligatory, permissible, and forbidden. Psychologists in this area explore our moral sense, looking at how it is instantiated in the brain, how it develops in children, and how it evolved.

The problem is that most research in this field, including my own, focuses almost entirely on how people make sense of, judge, and respond to the interactions of unrelated strangers. We have little to say about how people think of interactions that occur between parent and child, brother and sister, and other closely related individuals. We also often ignore moral judgments and moral feelings that concern spouses, close friends, colleagues, allies, and compatriots. I will argue here that these are the interactions that matter the most, and that our failure to explore them leads us to ask the wrong questions, design the wrong studies, and develop the wrong theories.

Psychologists often look toward other disciplines for guidance, and we are right to do so. If one is interested in some domain of human knowledge or action, it makes sense to turn to the most developed theories of the nature of this domain. Put differently, if you are interested in how people think about X, you need to know all you can about what X is. Accordingly, much of the work on the psychology of language is rooted in linguistic theory; much of the

work in human reasoning uses formal logic as a competence model; and much of the work in concepts draws upon developments in epistemology and philosophy of mind.

The relationship between moral psychology and moral philosophy is particularly intimate. Moral philosophers such as Adam Smith and David Hume could be seen as the founders of moral psychology, and many of the contemporary leading figures in the field have had some philosophical training. And, as we shall see, the theories and methods and even the experimental stimuli of moral psychology often come directly from philosophy.

Within the tradition of moral philosophy that has had the most influence on psychologists, the focus is on the morality of actions – what is optional, obligatory, or forbidden. The two main theories fall into the category of *consequentialist* views, which focus on the maximization of some utility (such as happiness) versus *deontological*, or Kantian, views, which propose that certain broader principles should be respected, even if they lead to worse consequences. One can see the difference by thinking about torture and capital punishment. Consequentialists maintain that torturing or killing a person, even an innocent person, would be the right thing to do if it leads to overall better consequences, while some deontologists will insist that torture and killing are always wrong regardless of the consequences.

Philosophers think up examples, often complex and unnatural ones, and use their intuitions about these examples as a tool for theory construction. This is similar to how linguists will use intuitions about the grammaticality and meaning of various sentences (also often complex and unnatural) to develop a theory of grammar. But there is an important difference here; the linguist's task, at least in the Chomskyan tradition, is to explore the mental processes that give rise to these intuitions. Linguistics is seen as a branch of cognitive psychology. For philosophers, in contrast, the goal is to use these intuitions as a springboard to a real theory of morality (whatever, precisely, this is). To see the difference, note that moral intuitions are sometimes contradictory; we might think X is right and Y is wrong, even though X and Y are identical scenarios described in different ways. A psychologist can stop there, accepting these inconsistencies as a fact about the human mind. A philosopher cannot.

At the same time, though, an adequate moral philosophy can't depart too far from intuitions. One wouldn't take seriously a moral theory that said that torturing innocent babies for fun is a good thing to do. Such a conclusion would be so unrelated to what we naturally think of as right and wrong that it wouldn't be a moral theory at all.

The working moral philosopher resolves this tension by engaging in what John Rawls described as "reflective equilibrium" — going back and forth between general principles and specific cases, ultimately coming to a point where a theory captures certain intuitions but rejects others. As a result, moral theories will make counter-intuitive claims. There are deontologists such as Kant, who tell us that lying is always wrong (always wrong? even if the Nazis are at the door, asking if there are Jews in the attic? Yes!) and utilitarians such as Bentham, who say that it's perfectly fine to torture and kill a baby if this action increases the sum total of the world's happiness by even a smidgen (a baby? an innocent little baby? Yes!).

Some of the most influential examples in modern philosophy concern runaway trains and runaway trolleys. Consider this one by Peter Unger, which has been expanded and elaborated by Peter Singer.

> One day when Bob is out for a drive, he parks his Bugatti near the end of a railway siding and goes for a walk up the track. As he does so, he sees that a runaway train, with no one aboard, is running down the railway track. Looking farther down the track, he sees the small figure of a child very likely to be killed by the runaway train. He can't stop the train and the child is too far away to warn of the danger, but he can throw a switch that will divert the train down the siding where his Bugatti is parked. Then nobody will be killed — but the train will destroy his Bugatti. Thinking of his joy in owning the car and the financial security it represents, Bob decided not to throw the switch. The child is killed. For many years to come, Bob enjoys owning his Bugatti and the financial security it represents.

Singer gives his own similar example: Imagine that Bob is walking by a lake and sees a child drowning in shallow water. Bob

could easily wade in and pull the child out, but this will ruin his shoes, which are quite expensive. So Bob walks on, letting the child drown.

These scenarios are constructed so that it's clear that Bob did something wrong through his failure to act. But now consider other failures to act. There are many dying children in the world, and Bob can save some of them by giving to charity. He can save a life for far less than the price of a Bugatti or even of Italian loafers. Singer's argument, following Unger, is that there is no relevant distinction here. Bob's choice not to sacrifice his car or his fancy shoes isn't relevantly different from Bob's choice to buy the car and the fancy shoes in the first place, instead of going to www .oxfam.org and sending the money to needy children.

Obviously, there are all sorts of differences here; for instance, when Bob fails to throw the switch, he is condemning an individual child to death, while when Bob fails to send money to charity, the effects are far less discrete. But Singer and Unger argue that such differences are morally irrelevant – this is a case where we have different intuitions about X and Y, but where X and Y are, in relevant regards, identical.

Consider now a similar runaway trolley case, one that has been hugely influential in moral psychology. This was first introduced by Philippa Foot in 1978, and subsequently expanded upon by Judith Jarvis Thomson. There are two related scenarios:

> SWITCH: A trolley is running out of control down a track. In its path are five people who have been tied to the track. You could throw a switch, which will lead the trolley down a different track. Unfortunately, there is a single person tied to that track and this will kill him. Should you throw the switch or do nothing?
>
> BRIDGE: A trolley is running out of control down a track. In its path are five people who have been tied to the track. You are standing on a bridge above the track, next to a very large stranger. The only way to stop the trolley is to shove the man off the bridge and into the trolley's path, killing him but saving the five. (It won't help to jump yourself; you're too small to stop the trolley.) Should you push the man or do nothing?

Y

Many people believe that you would be right to throw the switch, yet wrong to push the man, even though the outcome is identical in both situations − both throwing the switch and pushing the man would save five people and kill one. This suggests, among other things, that we are not natural consequentialists.

What's the difference between the two cases? One proposal rests on what is known as the Doctrine of Double Effect (the DDE), an idea often attributed to Thomas Aquinas. The DDE posits a critical moral difference between killing someone as an unintended by-product of causing a greater good to occur (which can be permissible) versus killing someone in order to bring about a greater good (which is not permissible). To see the DDE at work, consider the act of bombing an enemy military base knowing that the bombs will cause the death of some innocents who work at the base. This is done with the goal of destroying the base, perhaps ending the war quickly, saving millions of lives. The innocents are "collateral damage," like the dead man in the switch case. Now compare this to the act of bombing an enemy military base knowing that the bombs will cause the death of some innocents who live near the base. This is done with the goal of killing them and thereby intimidating the population into surrendering, perhaps ending the war quickly, saving millions of lives. The innocents die to bring about a greater good, like the dead man in the bridge case. Even though the overall goal is the same (to win the war), and even though the same number of people die in each case, many see the second act as worse than the first. This intuitive difference is elegantly captured by the DDE − in the second case, the deaths of innocents is a means to an end, while in the first case it is a regrettable by-product.

In John Mikhail's doctoral research, completed in 2000, he used trolley problems as a tool through which to investigate the structure of moral intuitions, exploring the idea that humans possess a moral faculty akin to a language faculty, one which includes subtle unconscious principles, including the DDE.

A year later, in an important paper published in *Science,* Joshua Greene and his colleagues used brain-imaging techniques to explore which parts of the brain are active when people reason about trolley and trolley-like situations. Since then, psychologists have

conducted Web-based surveys to assess the intuitions of hundreds of thousands of people from different countries and cultures, and variants of trolley problems have been tested with young children, with people living in small-scale societies, with psychopaths, and with patients suffering from various sorts of brain damage.

One clear finding is that just about every non-brain-damaged person shares the intuition that the switch and the bridge cases are different – most people, even young children, think that it is right to throw the switch; but wrong to push the man. Other studies, using different methods, support the claim that the DDE is a psychologically real principle – we draw a moral distinction between harmful acts that are unintended by-products versus those that are means to an end – and hence it might explain, at least in part, the intuitive difference between the switch and the bridge cases.

Other studies find that more corporeal factors, such as the presence of bodily contact, also affect our moral judgments about trolley cases. For instance, Greene and his colleagues find that people are more willing to accept the morality of using the man as an instrument to stop the runaway train if, instead of shoving him, they can throw a switch that opens a trapdoor, which makes him fall onto the track. This cannot be explained by the DDE. Greene argues, based on his fMRI data, that those intuitions that depart from consequentialism (as in the bridge case, where we are usually unwilling to kill one to save five) are best explained in terms of certain emotional responses. More generally, he proposes – in a paper titled "The Secret Joke of Kant's soul" – that deontological theories such as Kant's are merely rationalizations of these instinctive gut responses.

It is clear that all sorts of factors affect trolley intuitions. One clever study looked at effects of implicit cues as to the race of the characters. Will you choose to sacrifice an individual named Tyrone Payton to save a hundred members of the New York Philharmonic? Will you choose to sacrifice Chip Ellsworth III to save a hundred members of the Harlem Jazz Orchestra? Liberals were sensitive to this race manipulation. They were more likely to sacrifice a white person to save a black group than vice-versa – even though, when asked, they explicitly claimed that race shouldn't be a factor. In another study, people were given trolley

Y

problems after seeing a humorous clip from *Saturday Night Live*. This made them more consequentialist – more likely to endorse pushing the large man in front of the train.

There are many scholars who are uncomfortable with trolley problems, and with how bizarre and contrived they often are; in his *Experiments in Ethics*, Kwame Anthony Appiah observes that the dense trolley literature "makes the Talmud look like Cliff Notes." But there is little doubt that they have proven to be powerful tools for exploring the structure of our intuitions. They might be the fruit flies of moral cognition.

What about moral judgments concerning situations that involve family or friends? Real-world moral dilemmas often concern people we know, frequently family members. What about the sort of example we started with, about moral issues such as a mother's obligation to her son, or, more generally, about the rights and responsibilities that one family member has toward another?

Moral philosophers have used trolley and trolley-like problems to address precisely such issues. The idea here is that we can think more clearly about these controversial and emotionally fraught cases if we translate them into simplified dilemmas involving strangers. Indeed, when the trolley problem was introduced by Foot, it was intended to explore issues revolving around *abortion*, looking at the case in which the death of the baby is related to saving the life of the mother.

Years earlier, Judith Jarvis Thomson also addressed abortion, developing an example quite similar to the one by Gopnik that began this essay. She noted that debates about abortion often turn around the status of the fetus, over whether or not it is a *person* – the assumption being that if it is, then abortion is immoral. Thomson argues this is mistaken. To make this point, she presents the following scenario.

> You wake up in the morning and find yourself back to back in bed with an unconscious violinist. A famous unconscious violinist. He has been found to have a fatal kidney ailment, and the Society of Music Lovers has canvassed all the available medical records and found that you alone have the right

blood type to save him. They have therefore kidnapped you, and last night the violinist's circulatory system was plugged into yours, so that your kidneys can be used to extract poisons from his blood as well as your own. [If he were unplugged from you now, he would die; but] in nine months he will have recovered from his ailment, and can safely be unplugged from you.

Thomson's intuition here is that, while it would be nice for you to stick with him, there is no moral obligation: "If you do allow him to go on using your kidneys, this is a *kindness* on your part, and not something he can claim from you as his due." Similarly, then, a woman who is tethered to a fetus is not obligated to stay tethered to it, even if one assumes it is as much of a person as you or I.

This is a controversial argument, to say the least. Not everyone shares Thomson's intuitions about what should be done about the violinist. Also, pregnancy is usually the result of a voluntary act. To the extent that there is a clear parallel here between the violinist and the fetus, it is in cases of rape, and, indeed, rape is often recognized as a case in which abortion is permissible. (Thomson addresses this concern with a modified example later in the paper.) Note further that Thomson's argument, if correct, holds not just for abortion but also for *breast feeding*. It might be nice for the mother to feed her child (particularly if there is nobody else around to do so, and the child would die otherwise), but if the case is truly analogous to that of the violinist, there is no obligation for her to do so. One can take this as a surprising and valuable moral insight, or as a reason to question the merits of the analogy.

The trolley scenario with Bob's Bugatti relates to family matters in a different way. Consequentialists conceptualize right and wrong in terms of how they relate to overall happiness, fulfillment, or other forms of utility. Some family and group bonds might well be defendable on such grounds – it might be better if people fed themselves first, then their children, and then, after that, worried about others. This selfish priority ordering might be the best for all, in the same way that the instruction for using oxygen masks in airplane emergencies (you first, then your children) is the best system for ensuring that everyone survives. Any

rational consequentialist will also keep in mind the psychology of human attachment. For instance, parents often love their children, and get pleasure and satisfaction (and hence utility) from knowing that their children are safe and happy.

But Singer argues that the resources we selfishly give to ourselves and our families are far too great. It is a moral mistake, he argues, to lavish our children with luxuries in order to raise their happiness a little when the same resources could be used to save the lives of strangers. The Bugatti trolley example is intended as a stark illustration of our misplaced priorities.

This is one way to do moral philosophy. One develops general and abstract principles – perhaps *very* simple ones, as in consequentialism – and then extends them to particular cases, such as the relationship between mother and child. This is similar to how a linguist might propose general and abstract syntactic principles and then apply them to specific conversational contexts.

There are alternatives. A concern with issues of caring and trust is a central focus in certain strands of feminist philosophy; an emphasis on non-universal norms and standards is central to the philosophical school sometimes known as communitarianism. Then there are theories of "virtue ethics" that are less concerned with principles that apply to particular dilemmas and more focused on questions of moral character, often explicitly addressing family and community.

But it is the abstract philosophical approach, identified with scholars such as Mill and Kant, that influences the empirical study of moral thought and action. And so the method of moral psychology typically involves scenarios with anonymous strangers. It is not just trolley problems. Consider the large literature in experimental economics that explores our generosity and sense of fairness with others, using methods such as the dictator game, in which subjects are asked how much of their money they want to hand over to another. These games are usually played with strangers, under conditions of anonymity. Developmental psychologists from Lawrence Kohlberg to Elliot Turiel have given young children hundreds of different moral dilemmas but, again, these almost always involve strangers. (One exception to this is a famous

dilemma by Kohlberg, that asks whether a man has a right to steal on overpriced drug to save his dying wife.) The just published *Moral Psychology Handbook* has index entries for "Aristotle", "Hume", "Kant", and "Rawls", but none for "mother", "son", "family", or "friend".

Does this approach make sense? I am not concerned, here, with moral philosophy. But I am intensely interested in moral psychology: Does it make sense for psychologists to start with strangers and view family and community as a special case?

To answer this, consider first one core aspect of our moral natures. Certain basic moral impulses — both inner-directed, as reflected in feelings such as guilt, and outer-directed, as reflected in feelings such as righteous anger — are universal and show up very early in child development. They are likely to be at least in part the product of evolution. As such, they would have emerged in a world very different from the one we now live in. The natural history of morality began with small groups of people in families and tribes. Think of summer camp, not midtown Manhattan.

Morality, then, did not evolve to guide our actions and judgment in dealings with anonymous strangers. It evolved to help us deal with those with whom we are in continued interaction, in part to surmount situations in which selfish short-term desires must be suppressed for maximum gain. Those ancestors who were inclined to help others, gratified by the help of others, and motivated to punish those who defected, would have out-reproduced those without these sentiments, and this is why they exist today. But the same logic of natural selection dictates that these altruistic and moralizing impulses should not be indiscriminate — there is a strong reproductive benefit to being biased to favor friends and family, and one would expect this to be incorporated as part of an innate moral sense.

I should stress that there is no consensus as to precisely how morality evolved. Some scholars endorse an approach, developed by scholars such as W. D. Hamilton and Robert Trivers, in which our moral sense follows directly from the forces of kin selection and reciprocal altruism. But others, such as Robert Boyd and Peter Richerson, argue for a two-stage account, in which initial moral instincts get established, and then, as human society becomes

somewhat larger, a dedicated system for the acquisition of norms, including moral norms, emerges. There are interesting controversies here. But what's clear is that human morality did not evolve for the sake of dealing with strangers.

Indeed, much of the original context for morality is smaller than the community; it's the family. Humans are the Mama's boys of the biosphere. We have the longest childhoods of any creature; a long period of serious vulnerability. Even a strong empiricist, inclined to see all morality as learned and to dismiss the evolutionary claims of the previous two paragraphs, would have to agree that these early years are critical to the development of a moral sense, and hence that this sense will be calibrated to the family situation. Note further that the long period of dependence is a two-way street: our uniquely long period of fragility entails a correspondingly long period of serious parental investment. Babies depend on their parents; parents love their babies.

Now, none of this might matter from the perspective of developing a normative theory of morality. A philosopher might concede the importance that people give to family and community but argue that this is best addressed by a broader systematic theory. A philosopher might also believe that the interactions between strangers are the interesting ones. We need, after all, to know how to deal with the billions of strangers who share the world with us. Indeed, if it's true (as I shall argue below) that our natural moral sensibilities are nonexistent or blunted when it comes to faraway people, this is precisely where philosophy might need to step in. Family and community can take care of themselves.

(I should add that there's a less generous, but probably equally accurate, reason for the dismissal of family by philosophers and other scholars, summarized by Alison Gopnik on edge.org: "Child-rearing has been women's work, philosophy, psychology, theology and politics have belonged to men.")

What about moral psychology? It is possible that, although our moral understanding has emerged in a narrow context, it is nonetheless a general abstract system, one that draws no systematic distinction between friends and strangers and kin. But this is unlikely. The evolutionary history of morality has left its mark. Most people care more about their neighbors than about strangers,

more about their country than about other nations. Certainly most people care more about their children and their siblings and their parents than about their hypothetical strangers at the end of a track.

More than that, some moral intuitions turn around the importance of this caring. That is, we believe that certain preferences and attitudes towards kin and group are morally important. This is a point explored by Jonathan Haidt, one of the few contemporary psychologists interested in the moral attitudes we hold concerning friends and family. Haidt has developed an important research program exploring a "thick conception" of morality, which he explicitly contrasts with the narrow moral psychology perspective that focuses only on harm and fairness. In an article published in *Science*, Haidt presents a series of examples that pertain to family matters and matters of community. How much would someone have to pay you to do each of the following four actions?

Say something slightly bad about your nation (which you don't believe to be true) while calling in, anonymously, to a talk-radio show in your nation.

Say something slightly bad about your nation (which you don't believe to be true) while calling in, anonymously, to a talk-radio show in a foreign nation.

Slap a friend in the face (with his/her permission) as part of a comedy skit.

Slap your father in the face (with his permission) as part of a comedy skit.

People typically feel that the options on the right are more aversive than those on the left. This reflects feelings that are special to groups and families: loyalty towards one's country and respect for

Y

one's father. These feelings extend to moral judgment; we would be harsh towards someone who lacked these feelings of loyalty and respect.

There are other moral issues that are special to groups and to families. The strong moral response to incest arises (by definition) only when one thinks about sex among family members. The feelings of moral outrage that are sparked by sexual purity violations also seem strongest when the violations are committed by family members. This is reflected in those gruesome instances in which young women who engage in sexual activity before marriage are murdered by their brothers or fathers.

Note finally that, outside of the seminar room, the moral issues that most engage us have to do with families. This is appreciated by the authors of religious texts. If you want to read about moral dilemmas that speak to one's gut, put aside Hume and Kant and read instead the Hebrew Bible, which has familial morality at its core: brothers pitted against brothers, daughters seducing fathers, and so on. Think about Abraham commanded to kill his only son, a tale later mirrored in the Gospels, with God's sacrifice of his beloved only son, Jesus.

What about now, in a world in which many adults willingly live hundreds or thousands of miles from blood relatives? Readers of this article are most likely citizens of what Joseph Henrich and his colleagues have dubbed WEIRD societies – Western Educated Industrialized Rich Democracies. WEIRD societies are weird; they are not typical for our species, and one of their unusual features is the relative unimportance of family. But even in these societies, family matters. Like the rest of the world, WEIRD people are obsessed with the moral problems that arise concerning parents and children, brothers and sisters, and husbands and wives. If you doubt this, read the letters that people send to advice columnists, look at the titles of popular self-help books, or watch popular movies and television programs.

If we entertain the view that there are different moralities for different categories of individuals, it might help us better understand certain puzzles in moral psychology.

Consider again the switch case of the trolley problem. Should you throw the switch to save five strangers at the expense of one

stranger? People tend to say yes. One interpretation of this response is that we are moral consequentialists along the lines of Bentham and Mill. In the absence of emotional distractions, our judgments about right and wrong are based on how the world will be if one acts or doesn't act. Since five deaths are worse than one death, one should throw the switch.

A different possibility, though, is suggested by the fact that the individuals in the switch case are anonymous, faceless strangers. There is a large body of evidence from psychology and anthropology that our default response to such individuals is, at best, indifference. These individuals become morally relevant only if we are somehow induced to think of them as real people – for instance, if we see their faces, or if we know their names. But the added personal information isn't present in the switch case, and this raises the possibility that our intuition here isn't moral at all. As Richard Shweder has argued, we treat the dilemma as little more than a math problem: Which is less: 1 or 5? The majority who answer that the right act is to switch the trolley are doing precisely the same sort of reasoning that they would employ if they were asked whether to destroy one shoelace or five shoelaces.

This proposal differs from the standard consequentialist analysis, and does so in a way that's testable. The key difference is that the consequentialist analysis presumes that the switch case initiates *moral* judgments. If this is right, it should connect to certain other beliefs and inclinations—this is what it is to be moral. As an example, I don't like raisins. This is a preference, not a moral attitude, and so I don't care whether other people like raisins and don't think raisin-eaters should be punished. I also don't like killing babies. This is different, though; this is a moral attitude, and so certain other things follow: I believe other people *shouldn't* kill babies. I believe that baby-killers should be punished. And I believe that those who punish baby-killers are doing the right thing and those who reward them are doing the wrong thing. (These are the sorts of criteria my colleagues and I explore when testing whether infants and toddlers can make moral judgments, although, of course, we use milder moral infractions than baby-killing in our experiments.)

My prediction, then, is that our intuition about the switch case is more like raisin-eating than baby-killing. People might agree

that it's the "right thing" to throw the switch, but this is an abstract intellectual decision, not a moral one, and hence there will be little disapproval of those who fail to throw the switch, little desire to punish them, and so on. In order to turn the situation into one with moral import, the individuals have to cease to be strangers; they need to become morally evaluable beings.

When does someone cease to be a stranger? This question is addressed in the parable of the Good Samaritan. It begins with a series of questions by a lawyer to Jesus, asking what he should do to inherit eternal life. Christ responds by talking about love; you should love God and you should love "thy neighbor as thyself." The lawyer then challenges him, asking, "And who is my neighbor?" This is the answer:

> A certain man went down from Jerusalem to Jericho, and fell among thieves, which stripped him of his raiment, and wounded him, and departed, leaving him half dead. And by chance there came down a certain priest that way: and when he saw him, he passed by on the other side. And likewise a Levite, when he was at the place, came and looked on him, and passed by on the other side. But a certain Samaritan, as he journeyed, came where he was: and when he saw him, he had compassion on him, and went to him, and bound up his wounds, pouring in oil and wine, and set him on his own beast, and brought him to an inn, and took care of him. And on the morrow when he departed, he took out two pence, and gave them to the host, and said unto him, Take care of him; and whatsoever thou spendest more, when I come again, I will repay thee.

Jesus asks the lawyer which of these three men was the neighbor of the victim, and he responds, "He that shewed mercy on him." And Jesus then says, "Go, and do thou likewise."

The Samaritans were despised by the Jews, treated with open contempt. And so this story plainly has a moral message about transcending traditional ethnic boundaries. It is often taken as more than this, though: as a defense of a universal morality that encompasses all people. Your neighbor is everyone. There are no strangers.

But as Jeremy Waldron points out in an article in *The Monist,*

this is a misreading. Clearly, the Samaritan wasn't looking to help all humanity. Rather, he helped because he came across someone in his vicinity who was in need. The priest and the Levite are the villains of the story because they saw the man and did nothing. This is not a story about how everyone is your neighbor; rather it is a moral tale of the moral significance of *physical proximity*. As Waldron says, "Never mind ethnicity, community, or traditional categories of neighbor-ness. They are there and that makes them his neighbors."

This tale is a normative claim about how we should act, not a description of our moral psychology. But it does capture the conditions under which someone becomes morally relevant, both in our own actions and in how we judge the actions of others. To return to the Singer example, if I'm walking my dog in the woods near the house, and I see a child drowning in the river, I will surely be compelled to wade in and rescue her – even at the cost of my nice new shoes. This is true even if her skin color is different from mine and she is screaming for help in a foreign language. It would be true even if I were far from home. If I told you that I let her drown because of my nice shoes, and because I care only about people from my own neighborhood, you'd think that I was a monster and you would be right. Common sense tells us that proximity has moral significance. It may not trump family and may not trump community, but, still, it matters.

This brings us back to the bridge case. A scenario in which one person shoves another person in the path of a train puts the individuals into actual physical contact. Perhaps this is close enough that it transforms the potential victim into someone of value – it is no longer 1 stranger versus 5 strangers. This is a hard moral problem, and perhaps pushing the man is the right thing to do, but, in any case, this is a genuinely *moral* issue, eliciting emotions like guilt and anger, and notions of punishment, reward, and justice.

I'll end with a remark about kin. Like the categories of "stranger" and "neighbor", the question of who is and isn't "kin" is a surprisingly complicated one. One particularly interesting complexity is that societies have a notion of "fictive kin" – non-genetic relatives who are nonetheless thought of as blood. This is true even in WEIRD cultures. Where I was raised in Montreal, my neighbors and other

friends of my parents were described as aunts and uncles – this was presumably to signal that I should treat them, and trust them, and value them, as family.

In *How the Mind Works,* Steven Pinker points out that the most familiar example of genetically unrelated people being treated as kin is that of spouses, because their relation can be functionally identical to one of kinship. That is, "if spouses are faithful, if each acts on behalf of the union's children rather than other blood relatives, and if the marriage lasts the lifetime of both, the genetic interests of a couple are identical. Their genes are tied up in the same package, their children, and what's good for one spouse is good for the other." Although these conditions are rarely met in reality, still, it's no surprise that husband and wife form a tight family bond.

It is over such bonds that psychology departs from philosophy. Appiah tells a story about William Godwin, a contemporary of Bentham's, and a devout utilitarian. Godwin asks us to imagine that we could rescue only one person from a fire – an Archbishop of great moral distinction or someone more typical, like our father. As a devoted utilitarian, Godwin says the right answer is to leave Dad behind. This might be good philosophy, but it is not a decision that many of us would make, and it's not one that we view as morally right. Our theories of psychology should be able to explain why this is the case.

R