

Penultimate draft of Chapter 1 of [*Conversations on Ethics*](#) by Alex Voorhoeve, Oxford University Press (2009). Contact: a.e.voorhoeve [at]lse.ac.uk

1. Frances Kamm: In Search of the Deep Structure of Morality

In *The Gay Science*, Friedrich Nietzsche argued that only a form of philosophising that sprung from a deep commitment to the subject could ever hope for success. ‘All great problems’, he wrote, ‘demand *great love*’. He continued:

It makes the most telling difference whether a thinker has a personal relationship to his problems and finds in them his destiny, his distress and his greatest happiness, or an ‘impersonal’ one, meaning he is only able to touch them with the antennae of cold, curious thought. In the latter case nothing will come of it, that much can be promised; for even if great problems should let themselves be *grasped* by them, they would not allow frogs and weaklings to *hold on* to them.

Nietzsche went on to complain that, to his knowledge, no one had yet approached moral philosophy in this way:

Why, then, have I never yet encountered anyone, not even in books, who approached morality in this personal way and who knew morality as a problem, and this problem as his own personal distress, torment, voluptuousness, and passion?

No one familiar with Frances Kamm’s work in moral philosophy could share Nietzsche’s complaint. In her two-volume *Morality, Mortality* and in her other work in moral theory and applied ethics, Kamm meticulously and imaginatively analyses moral cases in order to gain insight into our fundamental moral concepts and principles. The tenacity with which she pursues this

aim springs from her personal engagement with the issues she investigates—an engagement reflected in her dedication of the second volume of *Morality, Mortality* to ‘the love of morality’.

At the centre of Kamm’s work lie her development and defence of a nonconsequentialist ethical theory. Consequentialism holds that the rightness or wrongness of our conduct is determined solely by the expected goodness or badness of the consequences of our acts or of the rules to which these acts conform. According to consequentialism, to act rightly is to act in ways that bring about the best possible expected consequences. To act wrongly is to fail to act in such ways.

Nonconsequentialists deny this. In support of this denial, many nonconsequentialists argue that our intuitive judgments of moral cases are inconsistent with consequentialism. The following ‘Bridge Case’ is a standard example: Imagine that you are standing beside a large stranger on a bridge over a trolley track. You see a runaway, driverless trolley heading down the track towards five trapped individuals, whom it will certainly kill upon impact. Before it reaches the five, it must pass under the bridge you are standing on. You realise that you will save the five if and only if you push the large stranger off the bridge and onto the tracks, where the trolley would hit him, kill him, and grind to a halt before it reaches the five. According to consequentialism, since pushing the large stranger into the path of the trolley would ensure the fewest number of deaths, you should push him. Many, however, judge intuitively that it would be impermissible to do so. Of course, intuitive judgments of this kind stand in need of justification. Those who pursue this line of criticism of consequentialism must, therefore, find principles that explain these nonconsequentialist case judgments; they must also show that these principles reflect important moral values.

These are difficult tasks. Consider, for example, the following attempt at formulating a principle that explains why it is impermissible to push the large man: killing someone is worse than letting him die, and indeed so much worse that it is impermissible to kill one (who would otherwise be alright) in the course of saving five others from death. The following ‘Side-track

Case' shows, however, that this principle is problematic. Imagine you see a runaway trolley heading down a main track where it will hit and kill five people. You can divert the trolley onto a dead-end side-track, where it will instead hit and kill one other person. In this case, most people judge it permissible to divert the trolley onto the side-track, even though, like the Bridge Case, it involves killing one and saving five. We need, therefore, to appeal to something other than the distinction between killing and letting die to explain the contrasting judgments in the Bridge and Side-track cases.

Frances Kamm, professor of philosophy at Harvard, has been at the forefront of the attempts to formulate principles that explain and justify our judgments in these and myriad other moral cases. We meet in London in January 2003 to discuss her case-based method, her conclusions about the morality of harming people and saving people from harm, and the view of human beings' moral status that she believes grounds these conclusions.

Alex Voorhoeve: What first drew you to philosophy?

Frances Kamm: I went to a high school for music and art. People there were interested in existentialism, Sartre in particular, so I started reading some philosophy of the Continental variety. I found that I liked thinking about these issues. Then, at Barnard College (the women's college of Columbia University) I took an introductory philosophy course with Robert Wolff. Of course, it had nothing to do with Continental philosophy—it was on Descartes' *Meditations*. What was wonderful about it was that you could read just a page or two and think about it for weeks. I was interested in history and in literature too—I took a wonderful course devoted to Tolstoy and Dostoyevsky. But it involved so much reading! 'Read the Brothers Karamazov in 3 days', they told us. The trouble was that I wanted to read just a few pages and think about it.

I wanted to go to medical school, but these ambitions were dashed because although I enjoyed the theory, I was a complete disaster in the lab. They would bring people over to watch me do experiments so they could have a laugh. So I wasn't going to be a lab person. I considered becoming a psychiatrist, but during a summer internship in psychiatric social work I realised that I simply didn't have enough patience with the patients. So I decided on philosophy.

I went to graduate school at MIT because I thought I would do philosophy of psychology, and its graduate programme connected philosophy, psychology, and linguistics. But when I got there, I didn't feel strongly motivated in that area. Then I took an ethics course with Bob Nozick at Harvard—and that's what did it, because I found what he was doing *really* interesting. You know, about one-and-a-half years before he died, Nozick gave a talk in my ethics colloquium at NYU, where I was teaching at the time. And when I introduced him, I said that for the last twenty years I had been finishing the term paper for his class. Most of the topics I have worked on were there: the distinction between harming and not aiding, the distinction between intentionally harming someone and harming them as a foreseen, but unintended side-effect of one's actions, the question of abortion, and others.

Alex Voorhoeve: What in that class captured your imagination?

Frances Kamm: First of all, I should say that we were all captured by Judith Thomson's imagination. She was a professor at MIT, but at that time she wasn't teaching what she was writing on, which was on the topics in moral theory we discussed in Nozick's class. In her class, we spent about half a semester on three or four pages of G. E. Moore's *Principia Ethica*. I was going *out of my mind!* At the time, this was not for me. Still, it was an eye-opening experience to see the level of care, detail, and rigour with which she

approached Moore's writing. I wasn't used to that. Nozick, however, *was* dealing with topics that interested me, and his classes introduced me to this case-based method of moral analysis. So what captured my imagination was the combination of my discovery of these new standards of rigour, of imaginativeness, and the use of cases or thought-experiments in ethics. I seemed to be able to do this, and after finding that I wasn't capable of doing other things, I thought, 'Well, I might not be a Renaissance woman, but I have one little ability, I enjoy doing it, and it seems to be something that the world would like to see done. So why don't I do this?'

Alex Voorhoeve: 'It seems to be something the world would like to see done'—what did you mean by that?

Frances Kamm: I meant I thought I could get a Ph.D. by doing it! You must realise that I wasn't the person voted most likely to succeed in philosophy by a long shot. I was hopeless for a while. It took me a long time to find my way. I went back to New York, hung around Columbia... My supervisor, Barbara Herman, was very patient; she saved my life.

I do think, though, that moral philosophy is very important, even when it doesn't relate to public policy issues, though that is probably what the world is interested in. I think that, nowadays, people realise that philosophy is not just the explication of erudite texts or incomprehensible discourses, but that it emphasises rigour and clarity—well, at least Anglo-American philosophy does. So people come to philosophy for practical guidance, certainly in bioethics. We now even have philosophers on medical wards, wearing beepers, who get asked by doctors, 'Come advise me whether I should pull the plug on this person's ventilator'. I have never done that. I tried teaching at the medical

school at NYU, but they were more interested in answers to specific real-life cases and I was more interested in theoretical issues that related to questions of practical import.

Alex Voorhoeve: Why do philosophers have authority in these matters?

Frances Kamm: Well, philosophers are different from religious thinkers in that they obviously are not referring to sacred texts; nor are they claiming that they have answers that have been revealed to them by a higher power. I mean, they are simply trying to provide universally accessible reasons for certain judgments. (Of course, there are individuals within religious traditions that do the same. Aquinas, for example, says that natural reason can give us answers to many things.) I think that philosophers very often do it better than others. When I look at arguments on war, abortion, or stem-cell cloning that are put out by people who have some interest in ethics, but who are not philosophers, but rather, say, literature professors, or even scholars who are well-known in the Catholic or Jewish traditions for doing this kind of work, I find that people who are trained in philosophy generally present better papers and are much better able to judge the validity of positions other than their own. I think people with analytic training simply have better argumentative skills.

Alex Voorhoeve: Is part of what sets philosophers apart a willingness to explore a very wide range of possible judgments, to imagine and consider cases or arguments that others might find too wacky to consider?

Frances Kamm: It is interesting that I would agree with that, since Peter Unger, who was a colleague of mine when I taught at NYU, calls me a 'preservationist' in philosophy because I very rarely reach conclusions that differ radically from our everyday morality.

Consequentialists are the ones who are prepared to accept anything if it will maximise good consequences, no matter how much it is at odds with our everyday moral thinking—like chopping up one healthy person to get organs to save five people who are in need of a life-saving transplant. My approach is generally to stick with our common moral judgments, which I share and take seriously.

I also think that some of the consequentialist philosophers who argue that we have very demanding duties to aid and that not aiding is morally equivalent to harming, but who don't live up to these demands, do not really believe their own arguments. You can't seriously believe that you have a duty to give almost all your money away to help others who would die without your aid and that refusing to do so is just as bad as killing them, and then, when we ask why you don't live up to that, say, 'Well, I'm weak'. Because if you found yourself killing someone on the street to save \$1000, you wouldn't just say, 'Well, I'm weak!' You would realise that you had done something terribly wrong. You would go to great lengths to avoid becoming a person who would do that. Now *that* is a sign that you believe you have a moral obligation not to kill someone. But when somebody says, 'My theory implies that you should be giving \$1000 to save someone's life and that failing to do so is just as bad as killing someone' and he also says, 'I don't give the \$1000 because I'm weak!', then I can't believe he really thinks he has that obligation to aid and that his not aiding is equivalent to killing. Imagine him saying, 'I just killed someone for \$1000 because I'm weak'. Gimme a break! This is ridiculous. Either there is something wrong with that theory, or there is something wrong with its proponents.

But I admit that when you put together some intuitive judgments, sometimes you get surprising conclusions. For example, it may be that many people are opposed to euthanasia or physician-assisted suicide and assume that only philosophers would regard it

as permissible. But in my article ‘A Right to Choose Death?’ I take ideas that people would ordinarily accept and show that they have surprising implications.¹

So sometimes you can shock yourself. Sometimes you are amazed, you know. You think, ‘Look at that!’ Even *I* come up with unusual conclusions, starting off in perfectly ordinary places, just by trying to think carefully. In a way, it is like what an artist does, creating something surprising from ordinary colours and shapes.

Alex Voorhoeve: Is there an analogy between what an art critic does when discussing an image and what an ethicist does when discussing a moral case? Faced with an observer who does not find a masterpiece beautiful, an art critic might direct the observer’s attention to aspects of the image in order to explain to him why it is beautiful; she might do the same for an observer who finds the painting beautiful while being incapable of articulating his reasons for this judgment. Analogously, in a moral case about which the observer’s initial judgment may be uncertain, or about which he may have a strongly held intuitive judgment that he is unable to support with reasons, an ethicist’s role may be to draw attention to various aspects of the case that provide reasons for coming to a moral judgment about it.

Frances Kamm: Yes. This may be just an autobiographical fact, but I don’t really have a considered judgment about a case until I have a visual experience of it. I have to deeply imagine myself in a certain situation, with an open mind. It is almost as if you are looking at something with no preconceptions. You have to attend to it, and then things will pop out at you. First you may get the intuitive judgment of what you really should do in the circumstance you are imagining. Then you wonder, ‘Why am I reaching this conclusion?’

¹ In ‘A Right to Choose Death?’, Kamm argues that when a patient wants to die and when his death is a lesser evil and his pain relief is a greater good, it is sometimes permissible to intend the patient’s death for the sake of relieving his pain.

And your inner eye focuses on one factor as driving this judgment. I suppose that it is the same sort of thing when you look at a painting. Make sure you are attending to it and aren't having stray thoughts. You start to focus on what is so fascinating about it. And it can take a while. You can develop a whole theory about what is causing you to have an aesthetic judgment, and the same can be said about judgments in moral cases.

Alex Voorhoeve: The idea that we just 'see' or 'intuit' the right response to a moral case suggests that such judgments are simply personal emotional responses or that we use a mysterious faculty in order to make them. You distance yourself from this interpretation of our case judgments. Why do you do so and what alternative view do you have of them?

Frances Kamm: The term 'intuition' has a long history, and when you use it people tend to think that you are talking about immediate access to atomic entities of some sort. I am not talking about intuitive judgments about the atomic structure of some substance—I am just talking about a judgment about a case. These judgments may be wrong, and we need to be able to give reasons for them, reasons which are not simply our emotional responses. What I am saying is that, in order to have a judgment about a case, you really have to situate yourself in the case. For example, imagine a case in which a runaway trolley will kill five people if a bystander does not divert it onto another track where she foresees that it will kill one different person. You have to imagine this case in detail. For example, you have to ask yourself, 'Which way is the side-track going? Is there a loop in the side-track, so that if you diverted the trolley away from the five, it would rush around and run into the five from the other side were it not that the trolley will hit the one person (thereby killing him) and grind to a halt?' You have to sink into these details, not just say immediately, 'Oh, it is one versus five, so of course you have to sacrifice the one', or, 'Oh, the track is going around, so...' Just *situate* yourself in the case.

I think this is what Judith Thomson is trying to get us to do in her famous discussion of the case where there is a loop in the track. In this Loop Case, some would immediately draw on a general principle like, ‘Don’t treat someone as a mere means!’² People who would immediately invoke this principle would say, ‘In the Loop Case, the hitting of the one is causally necessary to save the five, so if we turned the trolley onto the one, we would be using him as a mere means to the greater good. So it is impermissible.’ But I take Judith Thomson to be saying: ‘Forget about the principle! What do you think when you really have in place all the facts?’ And, surprisingly, people often say, ‘Well, I *do* think it is impermissible to use people as mere means, but I also think it is all right to turn the trolley in this case’. You can be very surprised about your own responses to a case like this one.

Alex Voorhoeve: But the principle that one should not treat someone as a mere means for other people’s ends is an important one that appears to capture an idea of moral significance, since it prohibits involuntary harmful exploitation. Why should we abandon it merely because of our—possibly tentative—response to an imaginary Loop Case of which we have had no experience?

Frances Kamm: Well, the Loop Case may show not that you should abandon the principle, but rather that turning the trolley away from the five in this case does not actually violate it. Instead, the Loop Case may indicate that there is a moral difference between acting to turn the trolley away from the five *only because* you will hit someone and *treating someone as a mere means*. In the Loop Case, you are turning the trolley away from the

² The principle ‘Don’t treat someone as a mere means!’ is attractive because it is consistent with the common judgments in the cases mentioned above: in the Bridge Case, it would prohibit pushing the large man in front of the trolley because this would involve using him harmfully as a means to the greater good. By contrast, in the Side-track Case, the killing of the one, though a foreseen (and unintended) consequence of turning the trolley away from the five, would not be a means to saving them. The principle would therefore not rule out turning the trolley onto the side-track in this case.

five for some goal (final or intermediate) other than hitting the one, since your intermediate goal is to stop the threat the five are facing now—being hit from the front. This is not true in the Bridge Case, where your intermediate goal is to push the man onto the tracks.³

So, I think we don't really understand what believing in a principle amounts to until we consider various cases to which it might apply. But I also think that our responses to other cases may undermine the validity of the mere means principle. In the first volume of *Morality, Mortality*, I consider a case in which we have to choose between saving either a doctor's or a janitor's life with an organ transplant. Assume that the doctor could do a great deal more good than the janitor: he would, let's imagine, go on to save another five people. In this case, I actually thought, 'Well no, we shouldn't give the doctor priority, even if he will go on to save five lives'. This appears to support the principle 'Don't treat people only as a means!' But then I thought about a case where I have several organs to distribute. One of these organs can help either person A or person B, both of whom are on my side of a river that I am too weak to cross. The rest of these organs can help five people on the other side of the river. Suppose that if I give the first transplant organ to person B, then she will be strong enough to cross the river; also suppose that A does not have this ability. Now, in this case, I thought, 'Ah, now that's interesting. I think it *is* permissible to give a transplant organ to B in this case even though I think it is *not* permissible to give it to the doctor rather than to the janitor—if A's ability is *connected to* distributing the resource I have, then it *is* okay. An instrumental difference between A and B is then a permissible ground for distribution. By contrast, the doctor is useful to people who need *him*, but not useful to the people who need what *I* have to distribute.' Right

³ In the Loop Case, once you turn the trolley away from the five, thereby ending the first threat to the five, they face another threat—being hit after the trolley loops around—which is defused only by the presence of the one on the side-track. And you turn the trolley away from the five only because you know that this second threat will be defused by the trolley grinding into and killing the one. Nonetheless, your initial aim does not involve the death of the one, and this might make a moral difference, Kamm believes.

away, this showed me that it wasn't just the purely instrumental use of people that was the issue.

The lesson I draw from such cases is that these variations save you from overgeneralisation, from an immediate attachment to a moral principle. There is so much variation, there are so many delicate distinctions to be drawn, that I don't think these very general principles we immediately jump to can possibly be correct.

Alex Voorhoeve: I don't understand what these cases have to do with the injunction not to use people merely as a means. For when we give B priority, we do not treat her merely as a means—we are also counting the fact that the transplant would save B's life as a reason for giving her the transplant. Moreover, B would presumably consent to being 'used' in this way. So the preference for saving B would not violate the injunction. The same would apply to the doctor, should we decide to give him priority over the janitor.

Frances Kamm: You have to consider it from the perspective of the person who is being *denied* the transplant because he is *not* useful. Take A: he is being treated as a mere means because we consider whether he could serve our goal and then refuse him the transplant because he will not. So I think the judgment that we should give the transplant to B violates the injunction. But it is true that, in saying this, I am relying on my sense that one can treat someone as a mere means even when he is not being made into something causally useful.

Alex Voorhoeve: What do you think about the following explanation for your judgments about these cases that, if correct, would undermine their purported authority. In actual cases our limited knowledge and limited ability to think through all the consequences of our acts, together with the advantages of specialisation, lead to a moral division of labour

and a concomitant division of responsibility. This division of responsibility is, roughly, one in which you are required to think only about *your* resource, about how *you* could put it to best use. So when what someone could go on to do with his life is relevant to the use of your resource, you take that into account. But you don't think in general about *everything* that the people who you are helping will do with their lives. The moral intuitions you have in these cases are a product of being educated into practices that embody this division of responsibility. But since the grounds for this division of responsibility are absent in your imaginary cases (where we know and can take into account everything that people will do with their lives) these moral intuitions may not be trustworthy.

Frances Kamm: Your explanation seems to be this: 'Well, ordinarily, we don't know what doctors will do with the life we give them, so we just tend to ignore this question—we'll make it an irrelevant good. Whereas here we *do* know.' But I don't think this explanation is sensitive enough to what is going on. In this case, I *can* suppose that I know that the doctor is going straight from his treatment to his surgery to treat five people only he can save. When I imagine this, I still see a difference between these cases. But then you might say: 'You have these ingrained tendencies from ordinary life; you are reasoning in a habitual mode. I know you think you've convinced yourself that the doctor is going to do this, but your habitual mode of thinking is: "Who knows what's going to happen, I am going to block out what the doctor is going to do".' I simply don't believe our capacity to reason about cases is so limited by habitual responses.

Nonetheless, I must admit that I was worried when my colleague Marc Hauser, a psychologist, reported to me in a research seminar that the moral distinctions people make in the various trolley cases are respected cross-culturally, in the sense that people from different cultures would turn the trolley onto the one to save the five in the Side-track Case, but would not throw an innocent bystander in front of the trolley to save the

five in the Bridge Case. I was also worried when a famous developmental psychologist told me that children had the same response. Because even though I have taken the view that many intuitive judgments reflect some deep structure not completely accessible to the person who has it and that this structure may be universal, I thought: ‘That’s terrible! I am just reporting some ingrained response! If even babies share this view...’

Alex Voorhoeve: *Babies?*

Frances Kamm: Well, you know, young children.

Alex Voorhoeve: Why does that worry you?

Frances Kamm: It doesn’t *really* worry me. I just meant that it doesn’t necessarily help us with the principles we are trying to explain and justify. That a lot of people agree doesn’t show that something is correct. It doesn’t help to support the normative authority of the judgment. Now, Marc Hauser himself recognises that. But there was a colleague of his at the research seminar who said upon hearing the empirical data, ‘Well, that’s great. What more do you need?’ Now, what worries me about that view is, of course, that *everyone could be wrong*. For example, the research by Daniel Kahneman and Amos Tversky on the psychology of decision-making under risk reveals that everyone makes certain mistakes when making decisions based on the probabilities of things happening. That everyone in fact decides as though they have a certain theory of probability ingrained in them doesn’t mean that the theory they are using is correct!

In addition, I was worried that if even very young children have these responses, then what is at work here is some remnant of the primitive part of the organism rather than a rational capacity—although I suppose someone might dismiss my worries, saying

‘That’s a good thing! Maybe morality is more widespread than we thought. Even *reptiles* would have this morality; they would flick their tails to turn the trolley onto one lizard rather than five, but they wouldn’t push another lizard in front of the trolley.’

Alex Voorhoeve: I don’t see why you should be worried simply because a particular response is ingrained. For example, vampire bats tend to rest next to the same neighbours day after day, and a bat who returns from his hunt to the bats’ resting place well-fed will often regurgitate blood into a hungry neighbour’s mouth. This pattern of interaction is repeated over time, and a bat who has fed his neighbour will, on other occasions, be fed by his neighbour in return. If a bat’s future access to such feeding is dependent on him feeding his neighbour when the occasion calls for it, then a disposition to engage in this kind of reciprocal altruism can be favoured by evolutionary pressures. Suppose that we have a similarly deep-rooted disposition to engage in acts of reciprocal altruism. That would not show that this tendency would be something that we would disapprove of on reflection.

Frances Kamm: I guess you could give an explanation of this ingrained tendency that is consistent with morality being something that you would have reason to follow. But we want to know if morality is something that rational beings could choose to develop; we don’t just want to know that they did develop it. The point is that the mere fact that we have a certain response to various trolley cases doesn’t indicate to us that there is a good reason to act on that response. It doesn’t solve the problem of normativity.

Alex Voorhoeve: You claim that this problem of normativity would be solved if we managed to find principles that explain our case judgments and that express morally

significant ideas. But what if the principles that best explain our case judgments fail to embody morally significant ideas?

Frances Kamm: Well, I guess that ultimately, you *could* reject all our moral judgments once you find out what is driving them. But I must admit that when I find a principle that appears to explain my judgment in a particular case and my reaction to it is, ‘How could *that* be of any importance?’, then my next response is that I should think more deeply about it because I must be missing something. It *could* be the case that when I see everything that is there, I will think, ‘This can’t be right’. But I am much less willing to give up my intuitions in particular cases than some proponents of the method of reflective equilibrium.⁴

Nonetheless, for me, these intuitions don’t provide sufficient justification. I am interested in figuring out our everyday moral consciousness, but I am prepared to accept as a possibility that once we see what the fundamental principles are, we will conclude that the intuitive judgments in particular cases are not worth adhering to.

Alex Voorhoeve: Can you give an example of this process of examining the moral significance of a principle that captures some of our intuitions?

Frances Kamm: Suppose that in order to explain my judgments in various cases I come up with a principle that says that the greater good itself can cause lesser evil, but the means that lead up to the greater good cannot.⁵ Then I need to ask, ‘Why should that be

⁴ The method of reflective equilibrium consists in working back and forth among our judgments about particular cases, the principles that we believe govern them, and the considerations that we believe bear on accepting these judgments or principles, revising any of these elements wherever necessary in order to achieve an acceptable coherence among them. We achieve reflective equilibrium when we arrive at an acceptable coherence among these beliefs.

⁵ This is a rough version of Kamm’s proposed Principle of Permissible Harm. At a first approximation, this principle states that it is permissible to harm some people when it is necessary to save a greater number of other people just in case the harm is an effect or aspect of the greater good that one produces. One may do greater good for some even

true? The person is going to die as a result of the greater good happening in both cases, so why should it make a difference whether the means that bring it about involve harming someone rather than that the harm is produced by the greater good itself? I also need to ask, ‘What might this constraint express about what it means to be a person?’ In the second volume of *Morality, Mortality*, I propose an answer based on the idea that it is a different matter when the fate of a person, who is of worth as an end-in-himself, is confronted *directly* with the good of other ends-in-themselves than when he is confronted with a chain of events which in themselves have no worth but which are means to a good end. Now, when I look back on that stuff I ask myself, ‘What does this amount to?’ But my point is that if you can’t find something of this sort, then there is something very unsatisfying about the whole thing.

I also want to say that when I look at these cases, maybe I can come up with something entirely new, something that may have nothing to do with treating people as ends or with the worth of a person. I might be awakened to some new aspect of the universe never before seen. When you read Kant, there is something that happens to students (it certainly happened to me!)—it is like a whole new dimension of the moral universe is opening up to you, but it comes to be known *a priori* as opposed to via empirical methods. That is why it is so difficult to understand Kant. Sometimes I think I am such a primitive form of humanity compared to the kind of being that he is on about. As Kant says, ‘We have reason not for the purpose of producing happiness, for instinct would have been better at that’—now whether that is true I don’t know—‘but for the purpose of producing a will that is good in itself’. You know, a will that is good in itself? *That* is the point of my life?! That is a standard I had never thought of measuring my life against! I taught a course on Kant once, and I thought, ‘What is going on here?’ But you feel there is gold to be mined. It involves a completely different conception of what your life is

though lesser harm to others may come of it. One may not, however, do harm to some in order that good to others will come of it. It is easy to verify that this principle fits our intuitions in the Side-track and Bridge cases.

about. We should always be open to the idea that life has dimensions that we haven't recognised yet.

Now, not everyone agrees with me on this. Baruch Brody [a professor of philosophy at Rice University] once said to me: 'I'll go along with your intuitive analysis, but why are you always searching for deeper principles? It is important that we realise that our judgments don't display the superficial irrationalities that consequentialists claim they display when we don't believe it is permissible to chop up one (who would otherwise live) in order to get transplant organs with which we can save five or to push the large bystander under the trolley, but we do believe it is permissible to turn the trolley even though it will hit a single person on a dead-end side-track. But all you have to do is show, as you do, that in the latter case the greater good leads to the lesser evil and in the former cases the lesser evil leads to the greater good. These intuitions wear it on their face that they are correct. Why do you have less confidence in them and more confidence in some deeper underlying theory of the person, or of interpersonal relations, or something like that? I am less confident about the correctness of the latter than I am that these intuitive judgments are correct.' I don't see it this way, even though I admit that we might never have the confidence in the theory that we have in our intuitions.

The reason we can't simply stop at the considered case judgments in which we have greatest confidence or at the relatively superficial principles that explain them is that there is a further question to ask: Why should we act as our moral judgment tells us we should? After all, adhering to moral requirements may be costly—we may have to sacrifice our aims, our lives, or the lives of others in order to respect moral requirements. A critical understanding of morality therefore requires that we understand which values are expressed in the moral principles that explain our case judgments and that we can judge whether these values are worth respecting or promoting.

The need to uncover a deeper rationale for our everyday judgments is particularly apparent in the case of nonconsequentialist rights against being harmed. (These rights constrain others' permissible actions towards a person and are therefore often referred to as 'constraints'.) These rights are notoriously hard to justify, because it appears irrational that one is prohibited from violating a right when such a violation will prevent more of the same type of right from being violated. By way of illustration, imagine a murderous twist on the Bridge Case: someone has maliciously sent a trolley hurtling towards five people with the intention of killing them. It is impermissible, nonconsequentialists hold, to kill an innocent bystander by pushing him in front of the trolley, even when this is the only way to prevent the murder of the five. But if it is so important that rights are not violated, how can the violation of the *single* bystander's right outweigh *five* equally harmful violations of the very same right?

In response to this question, Kamm attempts to formulate an idea of human status that such rights express and to show that this status would be undermined by the permissibility of acting to minimise the number of rights violations.

Alex Voorhoeve: Could you explain your idea of status and how it relates to constraints on harming?

Frances Kamm: When I developed this idea, I was grappling with the so-called 'paradox of deontology': if you care about rights, about people not being harmed in certain ways, then if you can stop five from being harmed in a certain way by harming one (who would otherwise be alright) in the same way, why wouldn't you do that? This would be a kind of consequentialism of rights-violations. Now, most people working in this area, like Bernard Williams and Thomas Nagel, attempted to resolve this paradox by bringing in an agent-centred perspective—they said, 'Well if *I* violated the rights of the one in order to stop the five from having their rights violated by *someone else* it would be *me* doing it,

whereas if *I* did nothing it would be *someone else* violating their rights'. Then I thought, 'If *I* had produced the threat, if *I* had, for example, set a bomb that would kill five, and the only way that *I* could now prevent this bomb from going off would be to kill a single, different person, would it then be permissible for *me* to kill this person in order to minimise the number of rights that *I* was violating?' And I realised that the answer was no.

So I decided to look in a different direction: not to focus on the *agent* who would violate the constraint, but instead to consider the potential *victims* of violations of the constraint. Because I thought, 'Suppose it *was* the case that I could kill the one to stop five from being killed? That would imply something about *all of us*. It would imply that *all of us* were useable in a certain way.' You see, there is a kind of status that is defined in terms of what it is permissible to do to people. One measure of people's worth is what we consider permissible to do to them. It is true that the five will be mistreated if I don't harm the one in order to stop their mistreatment. But they will still be the kind of beings *who should not be treated in that way*—who are *inviolable* insofar as it is *wrong* to harm them in a certain way, even in order to realise the greater good of minimising that type of harm. If it were permissible to kill the one to save the others, then *no one*—neither the one nor the five nor anyone else—would have the status of a highly inviolable being. So it is this value that is being expressed by this constraint. True, the constraint stands in the way of saving more lives and preventing more wrongs. But it expresses the fact that the nature of the individual is such that we are required to treat her in certain ways and are barred from treating her in other ways.

Alex Voorhoeve: What is it about our nature that gives rise to these requirements?

Frances Kamm: I don't know, although I suspect it has something to do with the fact that we are rational beings. But that was not what I was trying to establish. What I cared about was this paradox of deontology. And I realised that the answer lies in the fact that a moral system expresses the worth of a person, and that the worth of a person increases when people are less violable—as there are fewer constraints that it is permissible to violate. It is not only *what happens* to people that matters, but also *how their nature requires us to treat them*. The latter determines their worth.

Alex Voorhoeve: What about the following, different response to the paradox of deontology. The question, 'If you care about rights, why not aim to minimise rights violations?' might derive its force from the idea that caring about something (e.g. that people's rights be respected) always involves taking oneself to have reason to promote its occurrence, other things being equal. But it is not true of every valuable thing that if I care about it, I should promote its occurrence. Caring about friendship, for example, doesn't commit me to 'promote its occurrence' by maximising the number or quality of friendships I have. Breaking off one relatively time-consuming friendship so I can have two new friends would not show that I cared deeply about friendship. Perhaps what is true of friendship is true of rights.

Frances Kamm: Though this may be right about friendship, it doesn't work in the case of rights. Take a case where a single person and five different people are having their rights violated. Suppose we have to decide whether to save the single person or the five from this rights violation (we cannot save all six). In this case, we should save the five from having their rights violated. So it is clear to me that when I don't have to mistreat a person, I should maximise the number of people whose rights are respected. So maximising the number of people whose rights are respected *is* important.

Alex Voorhoeve: That's very interesting...

Frances Kamm: [Laughs.] I've just said something interesting! I can't believe it!

I mean, the other response is to say that one life is as precious as any number of lives. That faced with a choice between saving one life and a million other lives, you should flip a coin. But I think the view that I am proposing—the view about the importance of constraints in determining the status of a human being—has nothing to do with a refusal to count numbers.

You know, I am glad you are looking like you are interested, because that is what this stuff is all about. I mean, we are trying to do something important here. I am always surprised when people say, 'Oh, that was a nice discussion. That was fun.' I think, 'Fun? *Fun?* This is a serious matter!' You try and try to get the right account of the moral phenomena in moral cases, and getting it right is just as important as when you are doing an experiment in natural science or any other difficult intellectual undertaking. If we had worked on a NASA rocket and it launched well, we wouldn't say, 'Well that was fun!' It was awe-inspiring—that would be the right way of putting it!

Alex Voorhoeve: What about Shelly Kagan's challenge to your argument for the status-enhancing role of constraints? Kagan pointed out that as we *increase* our inviolability, thereby increasing our moral status because less may be done to us for the sake of maximising the balance of good over evil, we also *decrease* 'saveability', or what may or must be done in order to save us. And saveability, too, is a mark of moral status...

Frances Kamm: I don't always dwell on the objections that critics raise, because... well, because I don't always know how to answer them, but also because it is important for

someone who believes they are on the right track to keep going. But let me say something about it. Shelly is deriving a person's degree of saveability from facts about how much has to be done to save him when he is a member of a *group* of people who can be saved. But saveability seems to be a mark of someone's status only if it is a mark of how much you should do to save him *as an individual*. The status of a person *qua* individual is a function of what is true of any individual person. So Shelly's case isn't an indication of the status of a person as an individual. I guess the test for saveability is more like: if someone were drowning, how much of a loss would you impose on someone else, less than death, to save this person? That would be a way to argue that a degree of saveability shows someone's value. That is different from the question, 'How many would you kill to save a thousand people?' For the answer to *that* question is a function of the total size of the group. Look, you might have a very valuable jug in your house, and a lot of *tshatsbkes* [knick-knacks]. And given a lot of them, you might be willing to sacrifice the jug for them. But that wouldn't be a mark of their high status as individual items.

Alex Voorhoeve: Still, how many of the *tshatsbkes* it would take before you sacrifice the jug might reveal something about the value of each *tshatsbke*. Likewise, it seems that I am more inviolable when my right not to be killed can be overridden only in order to save the lives of more than twenty others than I am when it can be overridden to save just two.

Frances Kamm: It isn't straightforward that the number of people whose lives must hang in the balance before we can consider harming you in order to save these others bears on the determination of your inviolability. For example, Judith Thomson thinks that when it is a question of the violation of rights in order to help others, you have to engage in pairwise comparison of the loss that the right-holder would sustain if we violated his right and the harm that would thereby be prevented for *each* of the other people. You have to

ask about each of these people how much worse off he or she would be if the right-holder's right were respected. The degree of violability of a person's rights is then determined by how bad the fate of *one* other person would have to be before you transgress the right. The analogous idea in the saveability case would be that you have to think about it *one person at a time*; it would be to see, for example, how much I could take from one person to save another person.

Alex Voorhoeve: If inviolability is a mark of status, why shouldn't we make ourselves more inviolable—thereby giving ourselves higher status—by making it impermissible to turn the trolley away from the five and onto the one in the Side-track Case?

Frances Kamm: I don't believe that this idea of our 'making' something is correct. I don't believe that we construct morality—that we can make it the case that we have a certain status. Given what we are, we either have it or we don't. I believe that given a certain conception of the person, the rest of it follows.

Alex Voorhoeve: Even though you don't have this conception of the person?

Frances Kamm: I don't know what it is yet. I have it. I have it. There is no doubt about that, because I must have it, since I have the intuitions that express it.

Alex Voorhoeve: But you don't have it articulated...

Frances Kamm: I don't have it articulated. It is there, but I don't know what it is. If I ever manage to articulate it, it might be like one of those revelations when you read Kant. The thing is, I have beliefs that don't fit into a consequentialist model. I believe that I

shouldn't throw a fat bystander in front of the trolley when that is the only way to stop it from hitting the five if that will paralyze him, but that I can turn the trolley away from the five and onto the one and kill him when the one is on a dead-end side-track. There are certain ways in which you can bring about a result—it can't just be legislated. You might ask, 'Why don't we just eliminate that restriction and introduce another?' It isn't like that. It's a package deal. It's like a theory of grammar. Once you have the core, everything else follows. About a million billion sentences. And you can't say, 'Let's take five sentences and make the adverb work differently here'. You can't. The structure is hard, rigid. It is not something we put together, or can fool around with. And you have to figure out why it is rigid in exactly this way.

Alex Voorhoeve: Why did you dedicate the second volume of *Morality, Mortality* to the love of morality?

Frances Kamm: When I was writing that book, I was so immersed in it. Going over it again and again, you become completely detached from other people. And there is this structure that looms in your presence. At certain points, I just had to sit in my bedroom—the most secure place I could find—with a big bag of potato chips and a big bag of popcorn, and try to read the whole thing in one go. I had to encompass the whole thing in my mind all at once to see whether it would all fit together. Remember that when I think of these examples I have to think, to feel myself into them. And when I was doing this with the whole structure, I was trying to think as deeply as I possibly could—and I am no Saul Kripke!—and I just felt there wasn't anything else. This morality structure that I thought I had uncovered just seemed to me to be life. And it seemed to me to be so completely different to care about these things—most people I know didn't seem to grasp these principles. It was like a new world that I was having insight into and that other

people weren't aware of. And I dedicated my book to it, because I felt like I had been granted admission into this new world.

References and further reading

Much of Kamm's work is collected in her two-volume *Morality, Mortality* (Oxford: Oxford University Press, 1993 and 1996) and *Intricate Ethics* (Oxford: Oxford University Press, 2006). She is also the author of *Creation and Abortion* (Oxford: Oxford University Press, 1992). The article on euthanasia referred to is F. M. Kamm, 'A Right to Choose Death?' *Boston Review* (Summer 1997) (available online). Michael Otsuka, 'Kamm on the Morality of Killing', *Ethics* 108 (1997): 197-207 provides an accessible critical summary of Kamm's central ideas. Peter Unger vigorously criticizes what he calls Kamm's 'preservationist' approach to our moral intuitions in his *Living High and Letting Die* (Oxford: Oxford University Press, 1996). Kamm replies in *Intricate Ethics*.

The opening quotations are from F. Nietzsche, *The Gay Science*, ed. B. Williams, translated by J. Nauckhoff (Cambridge: Cambridge University Press, 2001), remark 345. Philippa Foot first formulated various trolley cases in her 'The Problem of Abortion and the Doctrine of the Double Effect', in *Virtues and Vices* (Oxford: Oxford University Press, 1978), pp. 19-33 and Judith Thomson developed other versions in her 'Killing, Letting Die, and the Trolley Problem', *The Monist* 59 (1976): 204-17, 'The Trolley Problem', *Yale Law Journal* 94 (1985): 1395-415, and 'Turning the Trolley', *Philosophy & Public Affairs* 36 (2008): 359-74. Thomson discusses rights in her *The Realm of Rights* (Cambridge: Cambridge University Press, 1990). How people respond to trolley cases is reported in Marc Hauser, *Moral Minds* (New York: HarperCollins, 2006) and in Joshua Greene, 'The secret joke of Kant's soul' (with discussion by John Mikhail and Mark Timmons), in *Moral Psychology, Vol. 3: The Neuroscience of Morality: Emotion, Disease, and Development*, edited by W. Sinnott-Armstrong, (MIT Press: Cambridge, 2007). Kazuo Ishiguro's novel *Never Let Me Go* (Knopf: New York, 2005) imagines a world in which healthy people are routinely killed

to harvest organs for transplantation and therefore offers a look at a world in which it is common practice to treat people merely as means to others' ends.

Some of the research on the psychology of decision-making under risk to which Kamm refers is collected in *Judgment under Uncertainty: Heuristics and Biases*, edited by Daniel Kahneman, Paul Slovic, and Amos Tversky (Cambridge: Cambridge University Press, 1982). Reciprocity among vampire bats is discussed in G. S. Wilkinson, 'Food Sharing in Vampire Bats', *Scientific American* 262 (1991): 64-70.

The method of reflective equilibrium in moral theory was first proposed by John Rawls in *A Theory of Justice* (Oxford: Oxford University Press, 2nd ed. 1999), section 4. See also Norman Daniels, 'Reflective Equilibrium', in *The Stanford Encyclopaedia of Philosophy Online* (2003).

The passage about Kant and the good will to which Kamm refers can be found in *Grounding for the Metaphysics of Morals*, translated by J. W. Ellington (Indianapolis: Hackett, 1994), Akademie page numbers 395-6.

Robert Nozick first drew attention to the paradox of deontology in *Anarchy, State, and Utopia* (New York: Basic Books, 1974), p. 30. Kamm refers to Bernard Williams's contribution to Jack Smart and Bernard Williams, *Utilitarianism: For and Against* (Cambridge: Cambridge University Press, 1973), pp. 93-100 and to Thomas Nagel, *The View from Nowhere* (New York: Oxford University Press, 1986), pp. 175-80. Shelly Kagan discusses the trade-off between inviolability and saveability in 'Replies to My Critics', *Philosophy and Phenomenological Research* 51 (1991): 919-28, on pp. 919-20.

