What does it mean? How do you know?

The Stanford Observer, November 1981

What we think of ourselves, of others, and the interrelationships of all of us has been a puzzlement since the beginning of time. This fact of life has provided a field day for philosophers, one of whom is **Michael Bratman**, *associate professor of philosophy at Stanford. He has been on the faculty since 1974. Bratman is a graduate of Haverford College, and earned his Ph.D. at Rockefeller University with a thesis entitled "Thought, Action, and Acting Against One's Better Judgment." That topic provided a starting point for a talk he gave on campus to alumni earlier this year. Following is an edited version of that talk, taken from the taped transcript.*

PHILOSOPHICAL disputes frequently articulate underlying tensions in the general intellectual culture. This is especially true of the disputes in that area of philosophy commonly known as ethics, or moral philosophy. A striking example of this is the dispute which has enlivened much philosophical reflection up to the present day between the 18th century German philosopher Immanuel Kant and the 19th century philosopher John Stuart Mill concerning the fundamental moral principles by which we are to lead our lives.

This dispute between Mill and Kant is the mere tip of a cultural iceberg. It's the reflection of a deep-seated tension in our common moral culture, a kind of moral schizophrenia to which all of us are, to some degree, subject.

In philosophy, contrary to certain popular misconceptions, serious understanding frequently requires reflection on a wide variety of fairly concrete examples. So I'll proceed in a rather example-oriented way, making some general remarks When they seem appropriate. But first a caveat.

Philosophers interested in ethics, myself included, sometimes strike others as rather odd for a variety of reasons. I'm going to tell you about one. Many in our culture assume that once we have stumbled, in our thinking or conversation, into a question of what a person's moral obligations are, or how one should lead one's life, there ceases to be any room for actual reflection or debate. It is as if such matters were as much a matter of taste as questions about what flavor of ice cream to order.

For philosophers, in contrast, it is at this point that things get really interesting, that the need for reasoned reflection and discussion is seen as most pressing. Just how a philosopher might proceed in a reasoned fashion with such matters is something of which I hope the rest of my talk may give you some idea.

SUPPOSE you're the only doctor in an emergency room in a small, rural hospital. There's been a terrible accident. They bring in six people, five of them seriously injured, one absolutely critically injured.

It becomes clear very quickly that you have a choice. You can devote full attention to the critically injured person and save that person's life; but if you do, the five very seriously injured people will surely die. Or you can devote your attention to the five very seriously injured people, saving their lives, but if you do that the very critically injured one will die. You cannot do both; there's no time to get additional aid. You're too far from Stanford Hospital. What do you do?

Well, many of the students I've talked to about this kind of problem say, save the five. Their principle seems to be that if you must choose between five people surviving and one person surviving, you choose the former.

That sort of principle is in the spirit of that perspective on morality which John Stuart Mill was a principal advocate, and is frequently called utilitarian. Thus the title of Mill's book, *Utilitarianism.*

The utilitarian's concern is with

increasing, as much as possible, human welfare. Since saving the five in this case clearly does this, he would see our clear preference for saving the five, even at the cost of the one person's life, as an indication that we are all deep down utilitarians.

Utilitarianism is a moral conception. It is a conception of how one should act, what one ought to do. A utilitarian supposes, roughly, that what we should do is what would have the best consequences for human welfare, what would, as it's sometimes put in the jargon of the trade, maximize human welfare. And this is surely a plausible view.

But problems emerge when we change the case slightly.

Suppose, now, that the night before, in this small town hospital, a hitchhiker whom nobody knows, and for all anyone can tell, has no family or friends, has stopped in for a routine checkup. Let's call him Sam. Sam is upstairs sleeping, awaiting his checkup, a perfectly healthy guy.

There now has been a terrible accident, and in come five very seriously injured people. A quick look at these people makes it quite clear to you, the doctor, that each of them has a very bad injury to one of their organs. In each case it's a different organ—one person's kidney is severely injured; one person's... You're catching on. You can usually tell when an audience gets the point of the example. You're a quick one.

So, what are you going to do? You need to transplant an organ for each person, to save that person's life.

Well, the thought of Sam comes to mind. He's up there, right?. He's got five very good organs. Pretty nasty thing to do to Sam—still you have to choose between five surviving and one surviving, and the principle used in the first case says choose the five. So you go upstairs and use Sam's organs to save the five.

Most would find this morally objectionable, and, if they thought doctors proceeded in this way, would not use hospitals too often. But it appears you were applying the same principle as in the first case; if you have to choose between five people surviving and one person surviving, choose the five.

WHY DO we discriminate between these two cases? You might be inclined to say the difference here is this:

In the first case you didn't kill the critically injured person by saving the five. You merely let that person die. Whereas in the second case, you didn't just let Sam die, you *killed* him. Or would kill him if you cut him up. So what distinguishes these two cases, in which you're reasonably sure you ought to save the five in the first case, but not in the second, is that you think there's a morally significant distinction to be made between letting someone die and kiling him.

The problem is that, on reflection, it's not obvious that you believe this. How can you find out if you really believe it? You do a thought experiment. You construct a pair of cases in which the only significant difference between the cases is that in one case you kill someone, and in the other, you let him die.

That is, the general issue is whether there is an important difference between causing some harm to someone and merely letting some harm happen. In this case the harm is death. Our question is whether the difference between the two emergency room cases really is the difference between killing and letting die.

One way of seeing whether we really believe that is to construct a pair of cases in which that's the only difference, and to see whether we really think there's some morally important difference in the two cases.

Let me do that, proceeding in my example-oriented way.

YOU HAVE an old grandmother from whom you stand to inherit a substantial amount, which you desperately need. You visit her on Wednesdays. You come in one Wednesday and she's taking a bath. She's all alone, and it occurs to you that you could substantially improve the chances of getting this inheritance, so you dunk her under the water,

killing her. This is a very bad thing to do. That's killing your grandmother out of greed. That's one case.

We'll make the second case exactly like the first, and change just one thing. It's not going to be a case of killing, it's going to be a case of allowing to die.

In this case, you visit your aging grandmother who's taking a bath, and all of a sudden, you hear some funny sounds coming from the bathroom. You walk in and she's slipped and fallen, and her head has fallen under the water. You could just reach down and pick her up. But it occurs to you that if you don't reach down and pick her up, you will inherit the money. So you don't. You allow her to die, out of greed.

I would claim that in this case, the only real difference is between killing and letting die. In the first case, you killed her and in the second case, you allowed her to die. I think most of you will find, upon reflection, that there's not really any morally important difference between the cases. In each case, what the person did was equally morally reprehensible, as wrong as it could be, and there just doesn't seem to be a substantial difference, morally speaking, between killing and letting die. So you can't believe that the difference between killing and letting die makes that much of a difference.

But if it doesn't, what accounts for our discrimination between the two emergency room cases, if it's not just a distinction between killing and letting die? How can we explain the fact that you ought to, or at least it seems you ought to, save the five in the first case, but in the second case—cutting up Sam—it seems inappropriate, to say the least.

LET'S change the examples to one that's a little harder.

You're driving along a country road with a friend and have a terrible accident. It's late at night. Your friend is pinned under the car, in severe pain, though it's not a life-threatening situation. There's nothing you can do without help. You're stunned, but in pretty good shape. It's quite clear that you should find the nearest help, to pull him out and relieve him of this terrible suffering.

So you begin hoofing it down the road. Unfortunately, you're in the middle of nowhere. You finally come across an old house, knock on the door and go in, and there's a little old lady who's taking care of her grandchild, a three year old. He's asleep in the living room.

The grandmother, however, has been reading newspaper accounts of muggings and is quite petrified when she sees you. She grabs her pocketbook which contains her keys to her car, runs into the bathroom, and locks the door.

You say, "I'm not here to hurt you. I need your car so I can go get help and save my friend." This of course only convinces her even more that you're a sly bugger. Everything you say persuades her even more that you're really out to harm her, and you think of your poor friend who's screaming, stuck under the car, and you think, what am I going to do? I need this lady's car to try to get someone else to help me get my friend out, but she won't give me the key. What do I do?

Aha, it occurs to you that what you could do is twist the child's arm severely so it really hurts. The lady has a peephole in the bathroom door so she can see what you're doing, and the kid will scream bloody murder. You're going to torture the young child until the lady give you the keys, and you're really sure that this is the way of getting the keys to the car. You then will get help in the town and save your friend in all of his pain.

Now, it's pretty clear to you that the child will, of course, suffer a substantial amount of pain, but nothing like the pain your friend is suffering. For example, it will only last 10 or 15 minutes, whereas if you don't get the car, you've got several more hours of hoofing it down the road before you can help your friend out.

It seems pretty clear that by torturing the child, you save your friend from even more suffering and pain. Do you twist the child's

arm and torture it to get the keys to the car to save your friend? Many people would say no, you're not permitted to torture a child. You're not permitted to treat an innocent person that way, even as a means to a worthy end—in this case, the end of saving your friend from enormous pain.

To the extent that you're moved by such intuitions, to say no, you simply cannot torture the child, your intuitions go against the utilitarian grain. They seem rooted in a very different moral conception, one that imposes strict side constraints on the permissible means we can use in pursuing worthy ends.

THE historical philosopher who most systematically explored this moral conception was Immanuel Kant. His views are complex, but for our purposes, his main formulation of this conception may be put in one of the following two ways (which he supposed were equivalent).

First, never treat a person, oneself or another, merely as a means.

Second, never treat a person, oneself or another, in a way that fails to respect that person as a rational being.

Now, this second moral conception deals naturally with some of the intuitions you may have had about some of the cases I've talked about so far. For example, it deals naturally with the view that you are not permitted to cut Sam up to save the five, since cutting someone up and using his organs is as clear a case of using someone merely as a means as anyone could think of.

Secondly, it deals naturally with the view that you're not permitted to torture the child as a means of saving your friend from even worse suffering. So those of you who would torture the child, but not cut Sam up, are probably quite ambivalent.

Let's explore this moral conception further.

J.L. Austin, a famous English philosopher in the early 1950s, once said you discover most of philosophy by asking two questions—what do you mean? and how do you know? I think he was doubtless right. He probably didn't even need both of them, We'll start with one of them.

WHAT does it mean to treat a person in a way that expresses respect of that person as a rational being? What does it mean not to treat a person merely as a means? Let's explore some cases.

Starting with sexual morality, it seems pretty clear that rape would be as clear a case as we're going to get of violating the requirement never to use a person merely as a means. On the other hand, sex with love between consenting adults seems as clear a case as we're going to get of acting in a way that expresses respect for a person.

Consider another kind of case: What about self-defense against an aggressor, when someone comes running at you with a machete or Colt pistol? Are you permitted to defend yourself even to the extent of using deadly force? I think it's clear the answer is yes.

This person has voluntarily chosen to attack you, and so you do not fail to respect him as a rational being, by defending yourself—though there are limits. If he's throwing a pillow at you, you're not allowed to shoot him to protect yourself. So certainly there are clear cases of self-defense which do not violate the Kantian requirement.

There are some hard cases here.

Suppose you are climbing a mountain, and there's a person behind you having his lunch. Along comes a big boulder, rolling downhill toward you. You see it coming, but the person behind you doesn't, and there's no chance to warn him. All that's open to you is to duck or not to duck. If you don't duck, you've had it; if you do duck, he's had it.

Are you permitted to duck to avoid the falling boulder although you know that if you do the boulder will smash the person below you to smithereens, if you also know that if you don't, it'll smash you to smithereens?

Or is ducking a case of failing to respect him as a rational being, or using him as a means? I think Kant would say no. I think

Kant would say you're permitted to duck.

Let's contrast that with a different story.

You're going up the mountain. The boulder is coming at you. You can't get out of the way of the boulder, but you can grab the person behind you and put him in front. Now, that's a case of using someone as a means. And that isn't permitted.

Well, things get pretty subtle here, but I'd like not to pursue this any further and turn my attention now to cases in which the issue is whether you're treating *yourself* with respect.

ONE of the most interesting cases that raises this problem is the case of what you might call the servile housewife. This is the traditional model of a woman who gives up her interest in pursuing some career outside the home for the sake of the husband's pursuit of his career. Is she failing to act in ways that show *self* respect?

The Kantian concept says you must act in ways that show respect for all persons as rational beings. You're not allowed to act in ways which fail to show respect for yourself any more than you're allowed to act in ways which fail to show respect for others.

The housewife case is a hard case. It raises a general problem: is self-sacrifice ever permissible? I think Kant would want to allow certain kinds of self-sacrifice, but he's going to make some important distinctions here. Let me explain what they are.

LET'S consider two different cases, similar in certain respects. Here you are in a wartime situation. You are out in a field. Five of your comrades are over on one side of a boulder. You're behind the boulder, and a grenade rolls in behind your comrades. There's no time for them to escape even if you should yell and tell them about the grenade.

It's pretty clear you have a choice: You could duck behind the boulder, in which case you will be protected, but your comrades will be destroyed; or you could jump on the grenade. Your comrades would be saved, but you, of course, would be killed.

Let's suppose that out of love for your comrades you jump on the grenade. It's not that you see your own life as any less worthy of love and respect, or yourself as not worthy of having been helped out in similar ways by other people in other circumstances. You have a great deal of self-respect but still, out of love for your comrades, you give up your life for them.

This form of self-sacrifice seems clearly compatible with Kant's conception, because there is no way in which you are failing to show respect for yourself, though it certainly is a case of self-sacrifice.

Now, contrast that with a different case. Imagine a person who, as we say in 20th century California terminology, has a negative self-image. He thinks he's worth something like negative 2 on the scale of human worth. We all, sadly, know people like this.

Out of his sense of his own lack of worth, he figures, well, what the hell, I'll jump on the grenade. I'm not worth anything, anyway. Now, that person saves his comrades, to be sure, but really does fail to show his respect for himself. He fails to satisfy the Kantian demand for self-respect.

THE LESSON here is that a lot depends on the *attitude* that you have, and that is expressed in your conduct. The action of jumping on the grenade will have a very different moral status, depending on the attitude expressed by that action, whether it's because of love for your friends or because of a bad self-image. This contrasts with Mill's conception, which focuses solely on the consequences—on the actual effects on human welfare.

The actual consequences are pretty much the same: you die, they live. But for Kant, that's not all that matters; it matters a great deal how you conceive of what you are doing.

Such cases of potential self-sacrifice help us see yet another important difference between these perspectives, which I'll put somewhat paradoxically in this way. Kant's

perspective is in a way more, but in a way less, demanding than that of utilitarianism. (I promised you it would be paradoxical.)

To see the sense in which Kant's perspective is less demanding, think about the grenade case and ask yourself what a utilitarian would say. He'd say that you'd have to jump on the grenade.

For after all, if you don't jump, five die, and if you do jump, one dies. And it doesn't seem, from the point of view of the utilitarian, that it's much different from the first emergency room case. In each case you have to choose between one dying and five dying.

Of course, in this case the one dying is you, which somehow seems important.

But if you're concerned about human welfare, if you have this general picture that each person counts for one—no more than one—it won't seem quite as important as it might seem to you hiding behind the rock. So the utilitarian, lacking some special story about the case, will say you have to jump on the grenade.

Now, this seems awfully demanding. That is, we frequently make, in our moral thinking, a distinction between what you are morally *obliged* to do and what, if you did it, you'd be a hero for doing.

We give medals to people who jump on grenades and call them heroes. But people who don't jump on grenades we don't consider to be, somehow, immoral; they're just not heroes.

The utilitarian doesn't make that distinction. So it would seem that what he requires is awfully demanding.

What would a Kantian say? He'd probably say that you'd be heroic if you jumped on the grenade out of love for your friends, but if you jumped on it out of lack of self-respect, it's a whole other story. But in any case you're not *required* to jump on the grenade, for in not jumping on the grenade you need not have failed to respect your friends as rational beings. Respect for your friend doesn't require that you give up your life for them—that's what I meant by saying the demand for Kant can be somewhat less demanding than the requirements of utilitarianism.

But the Kantian perspective can also be more demanding. To see why let's consider another kind of case, again one involving potential loss of life.

SUPPOSE you're one of a group of five spelunkers exploring a cave. There's a small earthquake, the entrance to the cave collapses and the earthquake unleashes an underground stream. The water in the cave is rising quickly and the five of you are in mortal danger.

Thank goodness you brought with you one stick of dynamite. It's clear that you need to blow up the sealed entrance to the cave. The problem is that a little peephole is left in the entrance and you can see, not far away, an innocent picnicker sitting around having her lunch. You yell, you scream, she doesn't hear you.

The water rises quickly and it's getting very wet in here. What do you do? Can you blow up the blockage, knowing that it will kill the innocent picnicker but save the five of you? Or must you allow yourselves to drown?

That's a hard question but I think it's clear what Kant would say even if it may not be clear on reflection what you would want to say. The answer is no, clearly no.

You're not permitted to blow up the cave since in doing so you will kill this innocent person, therefore failing to show respect for her as a human being. The utilitarian on the other hand would probably say yes because you can save five lives rather than one. In this case the Kantian is more demanding than the utilitarian.

This last case, in which the Kantian view is you're not permitted to save yourselves by blowing up the cave, together with some lingering suspicions you may have had about the Kantian prohibition on torture of the child, may make some of you increasingly uneasy. I hope it does.

You might ask, "Might not the

prohibition on treating someone merely as a means be overridden if enough can be gained by it?" Let's explore that, together with some real-life cases.

THERE is a famous 19th century case in British law about a cabin boy named Dudley.

There was a shipwreck and four or five of the sailors together with a cabin boy survived and floated around for weeks in a rowboat in the middle of the ocean. It was reasonable for them to expect that there would be a search party but they quickly ran out of food and fresh water and were in terrible shape. They got weaker and weaker, until it was clear the cabin boy was clearly the weakest and would probably die in a day or so. The sailors were starving; what were they to do? The sailors thought: this is a terrible situation, but if we don't do anything we will die. If we cut up the cabin boy and eat him we can live another week, and in another week we may get saved. So that's what they did. While the cabin boy was asleep they killed him and ate him. A few days later they were rescued.

They were taken back to the mainland and tried for murder. They were convicted and given minimal sentences. (That's the way the common law deals with hard problems.)

What shall we say? I think it's clear Kant would say: what they did was wrong. They were not permitted to treat Dudley, the cabin boy, that way. Cutting up and eating the cabin boy is clearly a case of using him merely as a means and so what they did was wrong. To be sure, the temptation was probably overwhelming because the fear of death has such a grip on most of us. So it probably was reasonable for the court to give them minimal sentences. But the court was right in ruling that they had committed murder. They were not permitted to treat the cabin boy as a means even if that's a necessary means to survival.

So says Kant. Of course the utilitarian would have a very different perspective, since it was clear that if they didn't do that they would die.

A utilitarian would ask, "What action would have the best consequences from the point of view of human welfare? If they don't cut up the cabin boy they will likely die; if they do, they'll survive," So the utilitarian would be much more sympathetic to cutting up the cabin boy.

Let's consider two more cases.

IN 1944 the French partisans were organized into the Free French forces of the Interior, and they were fighting openly and on a large scale with the occupying German forces. The Germans treated captured Free French partisans as traitors subject to summary execution. The French protested that they should be treated as prisoners of war, with immunity against such treatment.

By August of that year, 1944, many German soldiers began surrendering to partisan troops. It also became known that the Germans had recently executed 80 French prisoners. In reprisal the French forces executed 80 German prisoners of war and as far as anyone could tell there were no further executions of French prisoners by the Germans. The cases get hard, don't they?

Now, the Kantian must say that the French forces not only violated the existing policies of war which quite clearly forbid the execution of prisoners of war, but also acted in ways that were clearly not morally permissible, for there was no reason to think that these 80 Germans had done anything that would justify their execution. For example, there was no trial to show that these *particular* 80 Germans had themselves violated the rules of war and executed some civilians.

As prisoners, these 80 Germans had ceased to pose any threat to the French forces, so there was no justification for killing them. The 80 Germans were treated merely as a means to the end of deterring further German executions of the French forces. Kant would say the worthy goal of stopping further German executions cannot be carried out by means which treat other persons—the 80 German prisoners of war—

merely as a means and without respect.

One more example and then I'm going to draw some lessons from the whole thing. Think now of probably the hardest example, the British terror bombing of German cities in 1942.

IN 1940 TO '42 the British saw their bombing of German cities as their single effective weapon to prevent what they feared would be a Nazi victory and the destruction of all they held dear. But the bombing was clearly a terror bombing aimed at a civilian population, including many young children and old people who were clearly innocent of any involvement in Nazi aggression. There's no plausible argument that can connect a one-year-old kid with Nazi aggression or for that matter an 80-year-old invalid.

The British, in effect, used the killing and maiming of hundreds of thousands of such innocent civilians as a means of preventing what they saw to be an absolute disaster, a Nazi victory.

Now, a Kantian need not forbear self defense. He can allow, in appropriate circumstances, the bombing of military installations, munitions factories and the like, but he cannot allow deliberate terror bombing of children and old people who are clearly not the aggressors. That's a clear violation of his requirement that we not treat others merely as means.

Having said all this, many of you will be very uneasy, perhaps even impatient with the Kantian. Surely, you might say, we are morally required to treat some merely as a means so that many others can survive and important cultural values remain intact. (Though notice that if one of these important cultural values is respect for persons, as it surely was for the British, one's actions threaten to be self-defeating.)

To the extent you want to say there are circumstances in which you're required to treat some merely as a means in order to protect many more others and their important cultural values, I think you are being gripped by something like the utilitarian perspective.

But I think it's clear from other examples we've looked at, for example, our original case of Sam in the emergency room, that few of you are gripped solely by this perspective. Would you really cut up Sam?

WHAT EMERGES is a deep ambivalence. We're pulled by both perspectives, even when they are clearly in conflict. I began this talk with the remark that philosophical controversies frequently articulate deep underlying tensions in our cultures. I think you can now see why I said that the Mill-Kant controversy —and its 20th century analogues—is an example of this phenomenon.

I make so bold as to conjecture that there is some of both Mill and Kant in almost all of us, because both perspectives are deeply imbedded in our moral culture. But you may be getting impatient. How, you may want to ask, are we to reflect on this and reach some reasonably thoughtful resolution that is at least persuasive to ourselves?

This issue has been the focus of much recent philosophical research (where, by the way, many of the examples I've used in this talk have come from). But let me try to give you the flavor of one thing you might say.

Let's ask how these different moral conceptions see people and their relations to each other. Let me start by asking what may seem to be a very different question: What is it to be prudent?

TAKE A very mundane example, going to the dentist. Many of us know this will involve a certain amount of pain. Why do we do it?

Because we think it's worth incurring these present pains in order to avoid even more severe future pains. We try to minimize the overall amount of pain in our lives even if this means suffering some present pains to do it. In this trade-off we're willing to incur the present pain of the drill to avoid the future pain of the abscessed tooth. This is what we call prudence.

Now, suppose someone were to object: "Look here, that's not fair to your present

self, you're using your present self merely as a means to avoiding pain for your future self." We'd think that person bonkers. (The dentist can feel better now.) Why?

Well, we see reasonable tradeoffs between present and future pain of our own as the very heart of rational prudence. We don't feel we're being unfair to our present self by using it to avoid pain to our future self.

Now, the way the utilitarian sees morality, it is as a kind of generalization of this conception of prudence. If it's reasonable for your present self to make sacrifices for your future self, when your present self incurs the pain of the drill so that your future self can avoid the pain of the abscessed tooth, shouldn't it equally be reasonable for you to make sacrifices for me, or for him to make sacrifices for her?

The utilitarian says if it's reasonable for me to go to a dentist and incur present pain to avoid even more severe future pain, shouldn't it be reasonable to inflict pains on the child to avoid even more severe pains for the accident victim? After all, the pains in each case are all equally pains. Just as we try to minimize the overall amount of pain in our lives, in morality shouldn't we try to minimize the overall pain in human society and twist the arm of the young child? The utilitarian answers yes, and in this sense sees morality as a generalization of prudence.

Now, a Kantian will reply that this is a grave mistake, that in thinking this way the utilitarian fails to take seriously the distinctness of persons. The Kantian's going to say the relation between me and you is fundamentally different from the relation between my present self and my future self.

The utilitarian, says the Kantian, is inclined to see the welfare of each person as part of a larger whole, the "social welfare", in much the same way as prudence sees the welfare of my present self and of my different future selves as parts of a larger whole, which is my total individual welfare.

The Kantian would say this is a fundamental error, for it does not do justice to the distinctness of persons.

Well, we've gotten into deep water now. Our ambivalences in ethics, it turns out, go even deeper than we might at first have suspected. Among other things, they go toward even deeper ambivalences about what it is to be an individual person with a single life that extends over time and over significant changes in character, goals, values, and the like.

But having ventured into such waters, I'm afraid we'll have to stop. Philosophy always stops in the middle and today is no exception.