

The Realm of Rights, Chapter 6, "Tradeoffs"

Judith Jarvis Thomson

1. As I said at the beginnings of Chapters 3 and 5, it seems right to think that X's having a claim against Y is equivalent to, and perhaps just is, Y's behavior's being constrained in a certain way. We took note (in Chapter 3) of two things included in that constraint on Y, namely the need to seek a release in advance if the claim will otherwise have to be infringed, and the need to compensate later for harms or losses caused by the infringement if a release was unobtainable. But what is central to that constraint on Y is this: other things being equal, Y ought to accord the claim. What are those other things that may or may not be equal? The Tradeoff Idea provides an answer, for it says that it is permissible to infringe a claim--that is, permissible to fail to accord the claim--if and only if sufficiently much more good would come of infringing it than would come of not infringing it. Consider the cases we looked at in the preceding chapter. If the surgeon cuts the young man up and removes his parts, he will thereby save four lives (on balance). Why may he not proceed? I recommend that we say

The Claim Fact: If the surgeon cuts the young man up and removes his parts, the surgeon will thereby infringe some claims of the young man's, and it is not the case that sufficiently much more good would come of infringing them than would come of not infringing them

explains

The Impermissibility Fact: The surgeon may not proceed.

But there are other cases and other claims. Suppose that A has a claim against me that I not kick him in the shin. Suppose also that I would save four lives by kicking A in the shin. I should imagine that it would be permissible for me to do so. It suggests itself that we should say

If I kick A in the shin, I will thereby infringe a claim of A's, but it is the case that sufficiently much more good would come of infringing it than would come of not infringing it

explains

It is permissible for me to proceed.

These ideas need to be looked at more closely.

2. We need in the first place to attend to the question how the notion 'good' in the Tradeoff Idea should be thought to compare with the notion 'value' that we looked at in the preceding chapter. Here is a reminder of the details of one of our cases:

MAFIA: If the surgeon does not cut the young man up and remove his parts, the Mafia will kill five. If the surgeon does cut the young man up and remove his parts, the Mafia will not kill five.

On any reasonable suppositions about value, the act-plus-consequence set of the surgeon's proceeding has more value than the act-plus-consequence sets of his doing any of the other things open to him to do at the time. But I suggested that we should say: it is not the case that sufficiently much more good would come of the surgeon's infringing the claims he would infringe by proceeding than would come of not infringing them. Are we to take that notion 'good' to be the same as the notion 'value'? No.

What is wanted here is a notion 'good for a person'. Suppose you are starving. Then it would be good for you to get some food now, since you will live if you get it and die if you do not.

Moreover, it is equally good for you to get some food now, whatever the source of your need for food, whether your need for food was caused by villains or by negligence on someone's part or by nature. Consider the two communities we looked at in the preceding chapter, one in need of food because of villains, the other in need

of food because of nature. It is equally good for the members of each to get the food they need. That the villains will have killed the members of the one community if they do not get food, and no person (only nature itself) will have killed the members of the other community if they do not get food, makes no difference to how good it is for the members to get food.

Consider TRANSPLANT. It is equally good for the surgeon's five patients to get the parts they need, whether their need for parts is due to nature, as in TRANSPLANT (NATURAL CAUSES), or to villains, as in TRANSPLANT (5 VILLAIN CAUSES).

Or, as we might have put it: it is equally bad for the members of each community to not get the food they need and equally bad for the surgeon's five patients to not get the parts they need, whether the needs are due to villains or to nature or to anything else. These come to the same. That is, I will throughout use "good for a person" and "bad for a person" in such a way that if something that would have been good for a person does not happen, then something happens that is bad for that person, and if something that would have been bad for a person does not happen, then something happens that is good for that person. If a piano falls on my head, then something happens that is bad for me, a piano's falling on my head; if you very kindly intervene, so that the piano does not fall on my head, then something happens that is good for me, a piano's not falling on my head. (Compare the use by legal theorists and economists of the term "opportunity cost".)

And we are not to be looking at the acts that will have taken place if the surgeon does or does not proceed--and if we do or do not give food to those at risk because of villains--and asking what value those acts would have. That is irrelevant to us. We are to be asking only how good it is for those who would be affected by the acts if the acts do or do not take place or, alternatively, how bad it is for them if the acts do or do not take place. 1

So we should take the Tradeoff Idea to say this: it is permissible to infringe a claim if and only if infringing it would be sufficiently much better for those for whom infringing it would be good than not infringing it would be for those for whom not infringing it would be good. Alternatively: it is permissible to infringe a claim if and only if not infringing it would be sufficiently much worse for those for whom not infringing it would be bad than infringing it would be for those for whom infringing it would be bad.

I will not offer an analysis of this notion 'good for a person': it is important for our purposes, and I will lean heavily on it, but I leave it to intuition. It is good for a person who is starving to get food, it is bad for a person that a piano fall on his or her head, and I hope I will not say that such and such is or would be good or bad for a person where this will strike you as implausible. Two things should be stressed, however. In the first place, I do not construe the notion in a subjective sense according to which only pleasant feelings are good for a person and only unpleasant feelings are bad for a person. I so construe it that something that happens to you can be good or bad for you even if you never find out it happened to you, and so are neither pleased nor displeased at its happening to you. Second, I do not construe the notion in a subjective sense according to which what happens to you is good or bad for you if and only if you think it is. I so construe it that something that happens to you can be good or bad for you even if you think it is not. In short, I will everywhere construe the notion 'good for a person' objectively. The Tradeoff Idea in particular is to be understood to say that what matters is not how good those affected by a claim infringement would feel, not how good they think the claim infringement would be for them, but rather how good it really would be for them.

3. Here is the idea again:

The Tradeoff Idea: It is permissible to infringe a claim if and only if infringing it would be sufficiently much better for those for whom infringing it would be good than not infringing it would be for those for whom not infringing it would be good.

We have looked at the question how the notion 'good' in it is to be understood. Let us now take note of an important objection to it: it fails to take the claim holder seriously enough. (It fails to pay the claim holder sufficient respect, we might say.) It tells us that Y may infringe X's claim if and only if a certain balancing comes out satisfactorily, but nothing in the description of the balancing mentions X. For all it says, the claim might be anyone's, and that is surely wrong. If X is the claim holder, then the balancing must give a special

place to X's own good: what is to be traded off is X's good on the one hand and the good of others on the other hand, for it is X's claim that is being traded.

We can accommodate this point by revising the idea as follows:

The Tradeoff Idea: It is permissible to infringe a claim if and only if infringing it would be sufficiently much better for those for whom infringing it would be good than not infringing it would be for the claim holder.

Very well, but how much better is sufficiently much better? We might allow ourselves to speak of an 'increment of good'; how large does that increment of good have to be?

The answer is that there is no answer. There is no one size such that for any claim you choose, the claim is permissibly infringeable if and only if infringing it would generate an increment of good of that size. One thing (of several) that I have in mind comes out as follows.

We have an intuitive notion of the strictness or, as I will say, the stringency of a claim. The claims the surgeon would be infringing if he proceeded in TRANSPLANT and MAFIA are, we think, markedly more stringent than the claim I would be infringing if I kicked A in the shin. Suppose that if I kick A in the shin, four people will live who would otherwise die; that is very good for them. So infringing A's claim would be very good for those for whom it would be good. Let us suppose that I would cause A a mild, short-lasting pain and a bruise if I kicked him in the shin; that would be in a small way bad for him; thus not infringing the claim would be in a small way good for A. The increment of good in this case seems to be sufficiently large to make it permissible for me to kick A in the shin. But an increment of good of that same size is not sufficiently large to make it permissible for the surgeon to proceed in TRANSPLANT and MAFIA: a net gain of four lives is not enough. Perhaps no increment, however large, would make it permissible for the surgeon to proceed; we will return to that possibility below. At all events, this increment is not sufficiently large.

In short, the size of the required increment of good seems to vary with the stringency of the claim: the more stringent the claim, the greater the required increment of good.²

Let us look more closely at stringency. I suggest that we should take the stringency of a claim itself to vary with how bad its infringement would be for the claim holder. Thus I suggest we accept

The Aggravation Principle: If X has a claim against Y that Y do alpha, then the worse Y makes things for X if Y fails to do alpha, the more stringent X's claim against Y that Y do alpha,

and, what flows from it,

The Comparison Principle: Suppose X₁ has a claim against Y₁ that Y₁ do alpha, and X₂ has a claim against Y₂ that Y₂ do beta. Then X₁'s claim against Y₁ is more stringent than X₂'s claim against Y₂ if and only if Y₁ makes things worse for X₁ if Y₁ fails to do alpha than Y₂ makes things for X₂ if Y₂ fails to do beta.

These principles³ yield in particular that the stringency of a claim does not vary with the source of the claim. But that seems very plausible indeed. Claims generated by promises, for example, are surely no more, or less, stringent than claims generated in any other way; and similarly for any other possible sources of claims.

On the other hand, we should notice that accepting the Aggravation Principle and the Comparison Principle commits us to a conception of the stringency of a claim that might, *prima facie*, have been thought implausible.

We were supposing that A has a claim against me that I not kick him in the shin; and we were supposing that I would do nothing worse to A if I kicked him in the shin than cause him a mild, short-lasting pain and a bruise. Suppose, by contrast, that if I kick B in shin, I will thereby kill him. (B has a new disease, the most interesting symptom of which is that a kick in the shin is fatal.) The Aggravation Principle yields that A's claim that I not kick him in the shin is not very stringent whereas B's claim that I not kick him in the shin is very stringent; and the Comparison Principle yields that B's claim that I not kick him in the shin is markedly more stringent than A's claim that I not kick him in the shin. But I think we feel an inclination to deny both of these things.

I think we feel an inclination to say the following: B's claim against me that I not kill him is a very stringent claim, but B's claim against me that I not kick him in the shin--like A's--is not stringent.

I think we feel an inclination to say, quite generally, that stringency attaches 'directly to the wrong itself'. For many pairs of people X and Y, X has a claim against Y that Y not kick X in the shin, and Y would be acting wrongly in kicking X in the shin. This may plausibly be taken to mean that Y's kicking X in the shin would be Y's doing X a wrong. A wrong no matter what would happen to X in consequence of the kick. Admittedly, if Y kicks X in the shin, Y may also do X some further, other, wrong; admittedly the further wrong may be so very grave as to 'swamp', that is, make seem morally uninteresting, the wrong that Y does X in kicking X on the shin. All the same, whatever else happens, Y does X a wrong in kicking X in the shin. Now it is prima facie very plausible to think that the degree of stringency of the claim just is the degree of gravity of that wrong: that wrong--not some other wrong that Y might be doing to X by kicking X in the shin. We could call the idea that lies behind these remarks the Fine-Grain Conception of Stringency.

But there really is good reason to reject the Fine-Grain Conception of Stringency. It does seem right, prima facie, to attach stringency 'directly to the wrong itself' in the way I indicated, but some wrongs surely are particularly grave precisely because of the wrongs we do to the victim by committing them. Consider, for example, cutting a person's legs off where one ought not do that. Doing it is doing the person a very grave wrong, and what makes it so is presumably the wrongs we do the person by cutting his or her legs off. We do not grow new legs as we grow fingernails and hair, and artificial legs are no very good substitute; so by cutting a person's legs off you permanently cripple the person. I do not say that cutting a person's legs off would be doing the person no wrong at all if we did regrow legs, or if there were happy substitutes for naturally grown legs. Trimming a person's fingernails or shaving off a man's beard, without permission, without any other justification, are wrongs, despite the fact that the person will regrow nails and beard; causing a person to have a squashed nose without justification is a wrong even though there are ways of reconstructing noses. But cutting a person's legs off has got to be counted a more grave wrong by any sane moral theory--and a sane moral theory must surely trace that to the wrong one does to the victim by committing it.

Similarly for the wrong the surgeon does the young man by cutting him up and removing his parts.

So I suggest we instead adopt a large-grain conception of stringency according to which a wrong to a person by which we commit a further wrong to that person inherits gravity from the further wrong. More precisely,

The Large-Grain Conception of Stringency:

If

- (i) X has a claim against Y that Y not do beta, and
- (ii) if Y does alpha then he or she will thereby do beta, and
- (iii) X has a claim against Y that Y not do alpha,

then X's claim against Y that Y not do alpha is at least as stringent as X's claim that Y not do beta.

This conception of stringency does not attach stringency 'directly to the wrong itself'; but it does seem to capture our way of thinking about stringency better than the Fine-Grain Conception of Stringency does.

If we accept the Large-Grain Conception of Stringency--and I suggest we now do--then B's claim that I not kick him in the shin is very stringent, for it is at least as stringent as B's claim against me that I not kill him.

I suggest we should also accept something even stronger than the Large-Grain Conception of Stringency. The Large-Grain Conception of Stringency thesis tells us (roughly) that if you have a claim that a person not use a certain means, and a claim that the person not carry out the end, then your claim that the person not use the means is at least as stringent as your claim that the person not carry out the end. It is very plausible to think, more strongly, that if you have a claim that a person not carry out the end, then it follows from that by itself that you have a claim that the person not use any of the means by which the person would carry out the end, a claim that is at least as stringent as your claim that the person not carry out the end. In other words, it is very plausible to think we should accept

The Means Principle for Claims:

If

- (i) X has a claim against Y that Y not do beta, and
- (ii) if Y does alpha then he or she will thereby do beta, *then* X has a claim against Y that Y not do alpha, that claim being at least as stringent as X's claim that Y not do beta.

This thesis says that more than stringency is inherited: it says that the status of claim itself is inherited--it says that claims extend down from ends to means.

I think it a very attractive thesis, and in two ways at that. Here is my doorbell. So far so good. Does anybody have a claim against me that I not press my doorbell? Nicholas, for example? What on earth could be reason to think he does? Then we learn that if I press the doorbell, I will thereby kill him. (He and a battery are wired to the doorbell.) Nicholas has a claim against me that I not kill him. If he also has a claim that I not press the doorbell, then the Large-Grain Conception of Stringency tells us that that claim of his is at least as stringent as his claim that I not kill him. The Means Principle for Claims tells us that he does in fact have a claim that I not press the doorbell. Admittedly I would be doing him no wrong if it were not the case that by pressing the doorbell I would do him a harm. But that is the case. And don't we think he therefore has a claim--a very stringent claim--that I not press the doorbell?

A second reason for thinking the Means Principle for Claims an attractive thesis is its link with a thesis we met in Chapter 4: the Means Principle for Permissibility, which we can re-express in a way similar to that in which we just expressed the Means Principle for Claims:

The Means Principle for Permissibility: If

- (i) if Y does alpha then he or she will thereby do beta, and
- (ii) it is permissible for Y to do alpha, *then* it is permissible for Y to do beta.

This thesis says that permissibility too is inherited, but in the reverse direction--it says that permissibility extends up from means to ends. And it would be no surprise if the two theses were jointly true. If claims extend down from ends to means, then if there are claims infringed in carrying out the end, there are equally stringent claims infringed in choosing any means to it. No wonder it is permissible to carry out the end if it is permissible to choose a means to it.

So I suggest we accept not merely the Large-Grain Conception of Stringency but also the Means Principle for Claims.

Given we accept the Large-Grain Conception of Stringency, the Aggravation Principle and the Comparison Principle are safe. Consider the young man in TRANSPLANT and MAFIA. The young man has a claim against the surgeon that the surgeon not kill him. In and by killing him, the surgeon makes things very bad for him, since the young man dies. The young man also has a claim against the surgeon that the surgeon not cut him up and remove his parts. In and by doing this, the surgeon makes things equally bad for him, since the young man dies. The Aggravation Principle yields that both claims are very stringent; the Comparison Principle yields that they are equally stringent. That is surely as it should be.

Let us now go back. We were attending to

The Tradeoff Idea: It is permissible to infringe a claim if and only if infringing it would be sufficiently much better for those for whom infringing it would be good than not infringing it would be for the claim holder.

How much better is sufficiently much better? I said that the size of the required increment of good seems to vary with the stringency of the claim: the more stringent the claim, the larger the required increment of good. In light of our account of stringency, that means: the worse the claim infringer would make things for the claim holder by infringing the claim, the larger the required increment of good.

4. The Tradeoff Idea offers us necessary and sufficient conditions for permissibly infringing a claim, But in fact it is grossly oversimple.

Consider first the possibility that the prospective claim infringer is confronted with a conflict of claims. There are Lewis and Michael. Here am I, confronted, alas, with only two alternatives: kill Lewis or kill Michael. (How so? I am driving, and my brakes have just failed. If I continue straight, I drive into and kill Lewis, My steering wheel is constricted, so my only alternative to continuing straight is turning to the right. But if I turn to the right, I drive into and kill Michael.) We may well suppose that Lewis and Michael each has a very stringent claim against me that I not kill him, and that their claims are equally stringent: I would make things bad for each in the same degree if I killed him. If we suppose also that no one other than Lewis and Michael will be affected whichever alternative I choose--and for simplicity, let us suppose this--then there is no increment of good to be got by killing Lewis and no increment of good to be got by killing Michael. Yet it is surely permissible for me to toss a coin (heads I kill Lewis, tails I kill Michael), and then kill the one who loses the coin toss. In this case, then, no increment of good is required for permissibility, though the claim I will be infringing is very stringent.

That is a very simple case of conflict of claims, because of our suppositions that the claims are equally stringent and that no one else is involved. What am I to do in a case in which only two people are involved, but one's claim is more stringent than the other's? Then I presumably ought to infringe the weaker of the claims. Again, what if many people are involved? If their claims are equally stringent, I presumably ought to choose the course of action by which I infringe fewest. What if, to pile on complications, the claims of the many are not equally stringent? Thus what if I have to choose between doing alpha, by which I would cause a number of people to suffer in a variety of different ways, and doing beta, by which I would cause a number of other people (the same number or a different number) to suffer in a variety of still other ways? Then I presumably ought to make things as little bad for those potentially affected as possible.

But I say only "presumably" throughout. For example, it might be that I can compensate Y for the harm I would cause Y by infringing Y's claim, whereas I cannot compensate X for the harm I would cause X for infringing X's; then even if Y's claim is the more stringent, it might be that what I ought to do is infringe- Y's-claim-and-later-compensate-Y-for-the-harm-I-cause-him-or-her. (Notice that if I could get a release from X or Y in advance, then I am not in fact confronted with a conflict of claims, for I have a non-claim-infringing alternative, namely get a release from X or Y in advance.) For another example, it might be that I ought to infringe Y's more stringent claim today if I (or perhaps some other person) had had to infringe X's more stringent claim yesterday. There are any number of other possibilities here too. I think that a theory of rights cannot be expected to supply a non-vague general formula by means of which it can be decided what it is permissible to do in cases of conflict of claims. At all events, this one will not so much as try to supply a decision procedure for such cases.

There are other sources of trouble for the Tradeoff Idea beyond the possibility of a conflict of claims, but there is more that comes out of consideration of conflicts of claims that is important for our purposes, and we should stop to take note of it.

5. Consider TRANSPLANT and MAFIA. The young man has a claim against the surgeon that the surgeon not kill him. But what about the surgeon's five patients in TRANSPLANT? What about the five whom the Mafia will kill if the surgeon does not proceed in MAFIA? Don't they have claims against the surgeon that the surgeon save them? If so, then these cases too are cases of conflict of claims---just as the case of Lewis and Michael is--for the surgeon cannot save the five without killing the young man.

The idea that people have a claim to be saved, or perhaps more generally to be provided with what they need, is one that many people have found attractive. There are difficulties. For example, against whom is the claim had? Anyone who is available and can do the saving? Everyone who is available and can do the saving? Again, suppose Y can, and Y alone can, save X but only at a considerable sacrifice; does X all the same have the claim against Y?⁴

But we can bypass those difficulties, for we are now in possession of a strong argument to the effect that there is no such claim. Of course an argument is no stronger than its premises are, and I have no doubt that some who dislike the conclusion of this argument will wish to go back and reject one or more of its premises. But if

you have been coming along with me so far, you will not find that easy to do. (Indeed, this is one reason why I showed you the premises first.)

Suppose that Alfred is one of the surgeon's five patients in TRANSPLANT: Alfred will die if he does not get a new heart. Does Alfred have a claim against the surgeon that the surgeon save his life? The surgeon can in fact save Alfred's life--for there is an available heart, the young man's, and the surgeon can install it in Alfred--and to bypass the difficulties I mentioned, I stress that no one else can save Alfred's life and, moreover, the surgeon can do so at no great sacrifice to himself. So let us suppose that Alfred does have a claim against the surgeon that the surgeon save his life. The young man, of course, has a claim against the surgeon that the surgeon not kill him. So what confronts the surgeon is a conflict of claims.

Can it be thought that the two claims in conflict are equally stringent? Hardly. Lewis' claim against me that I not kill him is as stringent as Michael's claim against me that I not kill him; so flipping a coin was permissible--heads I kill Lewis, tails I kill Michael. But it is a datum that the surgeon must not kill the young man. Flipping a coin between Alfred and the young man is morally ruled out.

So Alfred's claim that the surgeon save him had better be less stringent than the young man's claim that the surgeon not kill him. But how can that be so? If the surgeon fails to save Alfred, that is very bad indeed for Alfred: Alfred dies. So

The Aggravation Principle: If X has a claim against Y that Y do alpha, then the worse Y makes things for X if Y fails to do alpha, the more stringent X's claim against Y that Y do alpha

tells us that Alfred's claim against the surgeon is very stringent. Of course if the surgeon kills the young man, that is very bad indeed for him: he dies. Is it worse for the young man that he die than it is for Alfred that Alfred die? It's pretty terrible to die in consequence of being cut up and having one's parts distributed for use by others; so add such details as you like (as it might be, a lot of pain) to fix that Alfred's death would be equally terrible. Then we are entitled to conclude, by

The Comparison Principle: Suppose X1 has a claim against Y1 that Y1 do alpha, and X2 has a claim against Y2 that Y2 do beta. Then X1's claim against Y1 is more stringent than X2's claim against Y2 if and only if Y1 makes things worse for X1 if Y1 fails to do alpha than Y2 makes things for X2 if Y2 fails to do beta,

that the young man's claim that the surgeon not kill him is no more stringent than Alfred's claim that the surgeon save him. So if Alfred has a claim that the surgeon save him, it is fully as stringent as the young man's claim that the surgeon not kill him. So it should be permissible for the surgeon to flip a coin. But it is not permissible for the surgeon to do that.

Consider also Bert, who is another of the surgeon's five patients but who only needs a kidney. Does Bert have a claim against the surgeon that the surgeon save him? Suppose he does. Bert will die if he does not get the kidney he needs, so the Aggravation Principle tells us the claim is very stringent indeed. The young man has a claim against the surgeon that the surgeon not cut him up and remove a kidney, (That is battery and theft,) How stringent is that claim? It's not nearly so bad to lose a kidney as to die: a healthy person can live with only one kidney, The Comparison Principle therefore tells us that Bert's claim against the surgeon that the surgeon save him is more stringent than the young man's claim against the surgeon that the surgeon not cut him up and remove a kidney. Faced with this conflict of claims, surely it is permissible for (perhaps even morally required of) the surgeon to cut the young man up and remove a kidney. But that won't do at all. The fact that I need a kidney for life, and that you can live with only one, certainly does not make it permissible for (much less required of) a surgeon to remove one of yours against your will.

I conclude that neither Alfred nor Bert, nor any of the surgeon's other three patients, has a claim against the surgeon to save his life.

What this argument does is to take the (putative) claim to be saved seriously. If there really is such a claim, then it should have the same *kind* of moral weight as any other claim does (it should have some or other

degree of stringency), and then it should be assessible for *amount* of moral weight (for degree of stringency) in the same way as any other claim. But supposing it does and is has unacceptable consequences.

Why do people want to have it that there is a claim to be saved? The answer seems plain. (i) There are situations in which it would be very easy for us to save and thus in which we ought to save, indeed in which it would be monstrous in us to fail to save. Here is a child drowning; I have a life preserver and can easily throw it to the child; it would be monstrous in me not to do so. Moreover, (ii) there are situations in which it would not be very easy to save--saving might require a considerable effort on our part, perhaps even some sacrifice--but in which we all the same ought to save, though perhaps "monstrous" would be too strong a word for a failure to save. There is a child drowning; I can swim out to save it, though that would require effort and is accompanied by some risk to me (drowning people tend to climb up those who try to save them), but perhaps the case is one in which I ought to expend the effort and run the risk.

But why should it be thought that those in need of saving in situations of kinds (i) and (ii) have a claim to be saved? Our old friend (from Chapter 3)

(T2) If Y ought not let it fail to be the case that p, then there is an X such that Cx, yp

says that wherever one ought to do something there is someone who has a claim against us that we do it; as we saw, however, that thesis will not do.

Perhaps the thought here is that anyway those in need of saving in situations of kind (i) have a claim to be saved, in light not of the overstrong (T2), but of the weaker

If Y not only ought not let it fail to be the case that p, but would be gravely at fault if he or she let it fail to be the case that p, then there is an X such that Cx, yp.

But why does there have to be a claim where we would be gravely at fault if we failed to do a thing? Is positing the claim supposed to explain its being the case that we would be gravely at fault if we failed to do the thing? But why do we need that explanation? If someone is in dire need, and I can easily supply what is needed, then why does that not by itself yield that I would be gravely at fault if I failed to do so? Why does the explanation have to run through a claim? I have been pointing as we go to a variety of ways in which people try to simplify the phenomena of morality, and here is yet another--and it is one for which, so far as I can see, there is no need.

6. We were looking at The Tradeoff Idea:

It is permissible to infringe a claim if and only if infringing it would be sufficiently much better for those for whom infringing it would be good than not infringing it would be for the claim holder.

How much better is sufficiently much better? I had said in section 3 that the size of the required increment of good seems to vary with the stringency of the claim: the more stringent the claim, the larger the required increment of good. But I then said in section 4, that the Tradeoff Idea is all the same grossly oversimple. Cases of conflict of claims make trouble for it. Where a prospective claim infringer is confronted by a conflict of claims, it may be permissible for him or her to infringe a very stringent claim, despite not producing any increment of good at all.

In a case of conflict of claims, there are three parties: A has a claim against me that I not do alpha, but a third party B has a claim against me that I do alpha. But third parties can make trouble for the Tradeoff Idea in yet another way. Suppose that A has a claim against me that I not kick A and that no one has a claim against me that I kick A; so I am not confronted by a conflict of claims. Suppose that if I kick A, I thereby saves four lives. Fine, we think: large increment of good, weak claim, so it is permissible for me to proceed. But Suppose that the saving of four lives is only net, Villains have said they will kill five unless I kick A; but A is standing next to B near the edge of a cliff, and if I kick A, A will be jostled against B, and B will fall off the cliff, so if I kick A, I will thereby kill one. Five minus one is four, so I save four if I kick A. But surely I may not kick A. If I kick A, I infringe only a weak claim of A's; but by kicking A I infringe a stringent claim of B's. I may not infringe that

stringent claim of B's for a saving of four lives. It follows (by the Means Principle for Permissibility) that I may not infringe that weak claim of A's.

In a case of conflict of claims, as I said, A has a claim against me that I not do alpha, but a third party B has a claim against me that I do alpha; in the case we have just looked at, A has a claim against me that I not do alpha, and the third party B also has a claim against me that I not do alpha. B's claim generates interference here, not by conflicting with A's claim but rather by reinforcing it.

And there are other sources of trouble for the Tradeoff Idea than cases in which the claims of third parties generate interference. Suppose that the only claims bearing on what I may do now are A's. Suppose that I can produce a large increment of good by infringing A's claims. So far so good. But what if I have an alternative, non-claim-infringing, way of producing that same increment of good? (We might notice here again that if I have a way of getting a release from A in advance, then I do have an alternative, non-claim-infringing, way of producing that same increment of good, though of course if I have to buy the release from A then the distribution of the increment of good will differ from what it would otherwise be.) However weak a claim, a large increment of good to be got by infringing it does not make infringing the claim permissible if one has a non-claim-infringing way of producing that same increment of good, or even a comparably large increment of good.

The Tradeoff Idea is obviously oversimple. It offers us necessary and sufficient conditions for permissibly infringing a claim, but what makes a claim infringement permissible or impermissible is an extremely complex affair, turning not only on the stringency of the claim, and the size of the increment of good to be got by infringing it, but on other things as well. What we have noticed in particular is that it is also relevant what the potential claim infringer's circumstances are. Do third parties have claims bearing on whether a potential claim infringer may infringe the claim? Does the potential claim infringer have alternative ways of producing the increment of good? Its being permissible or impermissible to infringe the claim turns on these matters too. I said earlier that I think a theory of rights cannot be expected to supply a nonvague general formula by means of which it can be decided what it is permissible to do in cases of conflict of claims; a fortiori I think that a theory of rights cannot be expected to supply a nonvague general formula by means of which it can be decided, quite generally, when it is permissible to infringe a claim. Supplying such a formula requires having in hand an exhaustive account of all of the other things besides the fact that you would here and now be infringing a claim in doing alpha that are relevant to the question whether you may here and now permissibly do alpha, and thus seems to me to require a full-scale moral theory. Claims do of course have moral significance, they do of course bear on what one may permissibly do; but they are not the only things that do.

Still, a theory of rights must tell us what the moral significance of a claim is, and that means it must tell us how claims bear on what one may permissibly do. I suggested that X's having a claim against Y is equivalent to Y's behavior's being constrained in a certain way; I suggested also that central to that constraint on Y is that, other things being equal, Y ought to accord the claim. How are we to cash that "other things being equal"?

Let us go back to the Tradeoff Idea. Perhaps it is possible to emend it? Notice, after all, that what we have pointed to by way of objection to it are considerations having to do with the potential claim infringer's circumstances; thus they are considerations that, while relevant to the question whether he or she may infringe the claim, are irrelevant to the question how stringent the claim is. What if we were to stipulate them away? Suppose X has a claim against Y that Y not do alpha, but would produce an increment of good by doing alpha. May Y do alpha? Perhaps we can say: if it is or were the case that no third parties have claims against Y either that Y do alpha or that Y not do alpha, and if it also is or were the case that Y has no non-claim-infringing way of producing that increment of good, then it is or would be permissible for Y to do alpha if and only if the increment of good would be sufficiently large, the required size varying with the stringency of X's claim.

This is an attractive way of proceeding. What we do in this way is to zero in on X's claim itself--we stipulate away what is irrelevant to *its* moral weight and attend only to that. Doing so eliminates the blur produced in actual circumstances by Y's alternatives, the claims of third parties, and so on; it lets us see the moral significance of X's claim itself. That is exactly what we wanted to see.

But there are two further objections to the Tradeoff Idea, so that two further emendations would be called for.

7. Suppose that A has a claim against me that I not kick him. I would cause him some pain if I did but not much, so the claim is not very stringent. Suppose that no third parties have claims against me that I kick him or that I not kick him. Suppose, last, that I would produce a large increment of good by kicking him and that I have no non-claim-infringing way of producing it. Fine, we think: it is permissible for me to kick A. But what if that large increment of good would be the sum of very tiny increments of good (small satisfactions, perhaps) for a large number of people? Surely that won't do. The objection we turn to here issues not from a difficulty about the potential claim infringer's circumstances, but from a difficulty about the distribution of the increment of good.

For my own part, I think we should agree to a very strong distributive constraint on claim infringement. Let us for the remainder of this section ignore the difficulties made by the potential claim infringer's circumstances and attend only to distribution. What I suggest by way of preliminary (the need for a revision will be pointed to in the following section) is this: I may not kick A unless there is at least one person, B, for whom I would thereby provide an increment of good of a size such that, if only B gained only that amount of good, his gaining it would by itself make it permissible for me to kick A. Where the increment of good I would produce by kicking A is the sum of tiny increments of good for many, there is no such B, for there is no person whose gaining a tiny increment of good would by itself make it permissible for me to kick A. Not so if the increment of good I would produce by kicking A is or includes B's being saved from death or blindness or losing his legs, for that increment of good for B would by itself make it permissible for me to kick A. In short, where claims are concerned, the sum of goods across people does not count. I do not mean that goods across people have no sum; I mean only that where claims are concerned, their sum across people does not matter. In still shorter form, where claims are concerned, the numbers do not count.⁵ Let us call this the High-Threshold Thesis.

Though I think we should accept the High-Threshold Thesis, there is no denying that it is very strong, and in two ways at that. In the first place, I take it to be plain that one may not (other things being equal) infringe a more or less stringent claim on the mere ground that one would thereby save one other person from a worsening comparable in gravity to the worsening that fixes the stringency of the claim. I may not cause you pain just to save Bloggs from a comparable pain; I may not kill you just to save Bloggs from death. If we agree that the numbers do not count, then we are committed to agreeing that I may not cause you pain to save five or twenty-five or hundreds from a comparable pain, and that I may not kill you to save five or twenty-five or hundreds from death. I said that the surgeon may not proceed in TRANSPLANT in that the increment of good he would thereby produce is insufficiently large; if the High-Threshold Thesis is true, we might have said instead that the surgeon may not proceed in that the increment of good he would thereby produce does not satisfy the distributive constraint.

Where the numbers get very large, however, some people start to feel nervous. Hundreds! Billions! The whole population of Asia! I suspect that those who start to feel nervous as the numbers get larger are moved to do so by the feeling that where billions (or perhaps even just hundreds) could be saved from pain, a person ought to volunteer to suffer it, and that where the lives of billions (or perhaps even just hundreds) could be saved from death, a person ought to volunteer his or her life. Perhaps so. (But then why isn't it also the case that he or she ought to volunteer for the sake of five? Or two? Or one?) I leave open whether he or she ought to. What is in question for us is only what may be done if the claim holders refuse to volunteer. The mere fact (if it is a fact) that the claim holders ought to volunteer really cannot be thought to make it permissible to infringe their claims. For example, many people do not volunteer to participate in Red Cross blood drives; the mere fact that they ought to (I believe they ought to) can hardly be thought to make it permissible to extract blood from them against their will.

A second consequence of accepting the High-Threshold Thesis is that accepting it commits us to there being what we might call maximally stringent claims. Let us say that a claim is within the range of maximally stringent claims if and only if no increment of good to be got by infringing it, however large, would make it permissible to infringe it--if the potential claim infringer had a non-claim-infringing alternative to infringing it. (Why I speak of a range of maximally stringent claims will come out shortly.)

Suppose that if I kill you I thereby produce an increment of good for Bloggs and only Bloggs. Is there any possible increment of good I could produce for him by killing you that would be sufficiently large to justify my

killing you? I should think it plain that there is not. If we accept the High-Threshold Thesis, then, the numbers do not count, and no matter how many I would save by killing you, from no matter what horrors, that does not make it permissible for me to kill you. No matter how large an increment of good I would produce, it is impermissible for me to kill you, and your claim against me that I not do so is maximally stringent. For my own part, I think it is, and thus that the fact that the High-Threshold Thesis has this consequence is no objection to it.

I said that the surgeon may not proceed in TRANSPLANT in that the increment of good he would thereby produce is insufficiently large; if the High-Threshold Thesis is true, we might have said instead that the surgeon may not proceed in that the young man's claim against him is maximally stringent.

I think that many other claims are also maximally stringent. Is there my increment of good I could produce by blinding you or cutting off your legs that would justify my blinding you or cutting off your legs? I think it plainly impermissible to do these things to you to save the life of one person; the High-Threshold Thesis yields that it therefore is impermissible to do these things to you to save the lives of any number at all. That consequence of the thesis also seems to me no objection to it.

It should be stressed, however, that the High-Threshold Thesis does not yield that just any claim is maximally stringent. The thesis does yield that I may not infringe your claim that I not cause you pain to save billions from comparable pain; the thesis is entirely consistent with its being permissible for me to infringe that claim of yours to save the lives of billions, or even to save the life of one.

Moreover, it is worth noticing that one can consistently suppose that some maximally stringent claims are more stringent than others. That is why I spoke of a range of maximally stringent claims. X's dying is markedly worse for X than X's going blind, or losing his or her legs, is for X, so it follows by

The Comparison Principle: Suppose X1 has a claim against Y1 that Y1 do alpha, and X2 has a claim against Y1 that Y1 do beta. Then X1's claim against Y1 is more stringent than X2's claim against Y1 if and only if Y1 makes things worse for X1 if Y1 fails to do alpha than Y1 makes things for X2 if Y1 fails to do beta

that X's claim to not be killed is markedly more stringent than X's claim to not be blinded or maimed. One can nevertheless suppose that these are maximally stringent claims.

Perhaps even clearer is this. It is surely worse for a person to be tortured to death than just to die. The Comparison Principle tells us therefore that the claim to not be tortured to death is more stringent than the claim to not be killed. I am sure that more people could be got to agree that the claim to not be tortured to death is maximally stringent (many people believe that torture is impermissible, no matter what the circumstances) than could be got to agree that the claim to not be killed is maximally stringent. But one can consistently believe that one claim is more stringent than the other and that both are maximally stringent.

But, as I said, the High-Threshold Thesis is very strong. Some people would reject it. On their view, the numbers do count, or anyway very large numbers do. So while I hope you will join me in accepting the High-Threshold Thesis, I will therefore not assume it true or rest any weight on it in later argument.

On the other hand, the distributive constraint cannot just be dropped. Those who reject the High-Threshold Thesis must find some other account of that constraint, for surely it is on no view permissible to kill a person to save billions from a minor headache. I leave it open what move those who reject the High-Threshold Thesis should make here.

8. A quite different kind of objection to the Tradeoff Idea comes out as follows. Suppose that if I kick A, I will thereby save B's life. (B will die if I do not kick A.) Then on any view I would produce an acceptably well-distributed and sufficiently large increment of good by kicking A to make it permissible for me to proceed. (On the assumption that no other parties have claims that are relevant. Let us once again ignore the difficulties made by the potential claim infringer's circumstances.)

Suppose instead that if I kick A, I will thereby very probably save B's life. (B will certainly die if I do not kick A.) Then I would very probably produce an acceptably well-distributed and sufficiently large increment of good by kicking A to make it permissible for me to proceed. Does the fact that I would very probably produce an acceptably well-distributed and sufficiently large increment of good by kicking A to make it permissible for me to proceed itself make it permissible for me to proceed?--thus whether or not B will in fact live if I proceed?

Another way to put this question is this. Suppose that if I kick A, I will thereby very probably save B's life. So I do kick A. As things turn out, B dies anyway. Was it all the same permissible for me to proceed?

I think we want to say yes. If B died anyway, then I did not in fact produce any increment of good by infringing A's claim. (Quite to the contrary: B gained nothing, and A suffered some pain.) But I think we want to say it was enough to justify my kicking A that I would thereby very probably save B's life.

We do welcome getting chances of getting things it would be good for us to get; we are even glad that we did get those chances, when it later turns out that we did not get the things it would have been good for us to get. I am grateful to you for making me a gift of a ticket in the state lottery; I am still grateful to you, when it later turns out that the ticket you gave me is not the winning ticket. Let us introduce a new piece of terminology. Suppose it would be good for X to get a thing Z. Let us say that if Y makes it probable that X will get Z, Y gives X an advantage; and let us measure the size of the advantage in a familiar way, namely by taking it to be the product of the amount by which Y increases the probability that X will get Z and the measure of how good it would be for X to get Z. When Y makes it probable that X will get Z, it may be that Y does more than give X an advantage, for it may be that Y will produce an actual increment of good for X. But Y need not do that in order to have given X an advantage.

Then if we want to say that it was permissible for me to kick A even though B died anyway, we could say that what made it permissible for me to do this is that I thereby produced not a large enough increment of good for B, but anyway a large enough advantage for B.

Should we say that it was permissible for me to kick A? Well, what is the alternative? If it was not permissible for me to kick A, then it was the case that I ought not kick A. Is it at all plausible to say that the mere fact that B died anyway makes it be the case that it was wrong for me to kick A--despite the fact that by kicking A I was very probably going to save B's life?

There is a bad argument to the effect that it was not wrong for me to kick A, that it was not the case that I ought not, that it was not impermissible for me to do so, and we should look at it.

What I have in mind is the two-step argument that goes as follows: "Since I was very probably going to save B's life by kicking A, it follows that I was not at fault, or to blame, for kicking A; and it follows from my not having been at fault, or to blame, for kicking A that it was permissible for me to kick A." But both steps are unacceptable.

First, its having been the case that I was very probably going to save B's life by kicking A just does not entail that I was not at fault for kicking A. What relieves of fault is not probabilities, but an agent's estimates of probabilities. (These may of course differ. I might think it highly probable that a six will turn up on the next throw of a die--perhaps because no six has turned up in recent throws--and be mistaken in thinking so.) If Y kicks X out of the belief that doing so will very probably kill Z, then (other things being equal) Y is at fault for kicking X, even if in fact Y's kicking X will very probably save Z's life.

The second step is equally unacceptable but in some ways more interesting. The second step presupposes the general thesis that (1) X was not at fault for doing alpha

entails

(2) It was permissible for X to do alpha.

In light of our decision on usage (Chapter 3, section 3), saying that (1) entails (2) comes to the same as saying that (1) entails

(2') It was not the case that X ought not do alpha.

Let us call the thesis that (1) entails (2) and (2') the Requirement-of-Fault Thesis for Ought-Not. Very often when we would not be at fault for doing a thing, it is also true that it would be permissible for us to do it, and thus that it is not the case that we ought not do it; The Requirement-of-Fault Thesis for Ought-Not says something stronger, namely that there is entailment here.

Why might a person accept the Requirement-of-Fault Thesis for Ought-Not? Here is Bloggs, trying to figure out whether or not to do alpha. His best estimate, on very careful consideration, is that it is likely to be better that he do alpha, and conscience tells him to do alpha. So he does alpha. As things turn out, his estimate was wildly off, and his doing alpha was disastrous. But isn't one's best estimate, on very careful consideration, good enough? Isn't Bloggs fault-free? It seems reasonable to think he is. How can it be right to say he ought to have acted against conscience? But this really won't do. Suppose Bloggs' baby has a fever. Bloggs' best estimate, on very careful consideration, is that it would be best to starve the baby. (Feed a cold, starve a fever.) So he does, and the baby dies. If that really was Bloggs' best estimate, on very careful consideration--if Bloggs really was not at fault for his mistaken belief--then we do not blame Bloggs for the baby's death, we do not say he is at fault for it. Do we conclude that he did not fail to act as he ought? The doctor will later say "I quite understand, it's not your fault"; but he or she will go on to add "but you really ought not have starved the baby, you ought to have fed it aspirin dissolved in apple juice".⁶ Not just that it would have been a good idea, rather that it is what Bloggs ought to have done.

In light of such possibilities, most people who take this line say we need to distinguish. The word "ought", they say, has two senses. It has a subjective sense according to which what a person ought to do is only what he or she would be at fault for not doing, thus a sense for which the Requirement-of-Fault Thesis for Ought-Not is true; and it also has an objective sense according to which what a person ought to do is not fixed by what his or her beliefs are, not fixed by those beliefs even if the person was not at fault for having them, thus a sense for which the Requirement-of-Fault Thesis for Ought-Not is not true.

For my own part, I see no such ambiguity in the word "ought". Suppose Bloggs was thinking out loud as he engaged in his very careful consideration of what to do for the baby. "Yes", he concluded, "I ought to starve it". I am inclined to think that there just is no sense of the word "ought" according to which he spoke truly. No sense of the word "ought" according to which anyone would have been speaking truly if he or she, or they, said that Bloggs ought to starve that baby.

In any case, I will throughout use "ought", and its cognates such as "permissible", in the objective sense. ⁷ If this practice makes you feel uncomfortable, you are invited to make emendations in the text. If you think that "ought" has two senses, a subjective as well as an objective, then you may wish to write "objective" before every occurrence of "ought". If you think that "ought" has only one sense, the subjective sense for which the Requirement-of-Fault Thesis for Ought-Not is true, then you may wish to insert clauses attributing appropriate beliefs to the agents in question.

To return now to probabilities. It was the case that if I kicked A, I would very probably save B's life. So I kicked A, but B died anyway. We were asking whether it was all the same permissible for me to kick A. What we have just done is to look at a bad argument for the conclusion that it was permissible for me to kick A, a two-step argument that proceeds via fault. But both steps are unacceptable. I know of no good argument for the conclusion that it was permissible for me to kick A. Some people say "Well, isn't the point of saying 'ought' to guide action? And wouldn't it have been a terrible thing to say to you at the time--given the probabilities--that you ought not kick A then?" But the fact that it would have been terrible to say a thing does not mean it was false.

On the other hand, I know of no good argument for the conclusion that it was not permissible for me to kick A. And we do very naturally think it was permissible. Moreover, I will suggest in the following chapter that a

certain hard problem in moral theory is best dealt with if we allow probabilities, or more precisely advantage, to have a direct bearing on what it is permissible to do.

But if we do allow this, if we do say that it was permissible for me to kick A, then much of what has preceded needs revision. We have to say that what made it permissible for me to proceed was not that I would produce an acceptably well-distributed and sufficiently large increment of good by kicking A, for I did not in fact do this, but rather that I would produce an acceptably well-distributed and sufficiently large advantage by kicking A. As is plain, any requirement on tradeoffs would have to allow for the justification of claim infringements by appeal to advantages produced. And the High-Threshold Thesis would have to be revised: presumably it should be so revised as to require that the advantage to at least one for whom infringing the claim would give advantage must be sufficiently much larger than the advantage to the claim holder if the claim is not infringed for that advantage to the one, by itself, to justify infringing the claim.

9. We were looking at

The Tradeoff Idea: It is permissible to infringe a claim if and only if infringing it would be sufficiently much better for those for whom infringing it would be good than not infringing it would be for the claim holder.

As we saw, it is grossly oversimple. Still, tradeoffs of claims on the one hand for increments of good, or anyway increments of advantage, on the other, are permissible. And the thought that lies behind the Tradeoff Idea is surely right. Moreover, we can extract from the thought that lies behind it an account of what is central to that constraint on Y's behavior that X's having a claim against Y is equivalent to. Suppose X has a claim against Y that Y not do alpha. If we stipulate away matters having to do with Y's circumstances (are there third parties who have claims against Y that Y do alpha or that Y not do alpha? and so on), these being matters relevant to whether Y may infringe X's claim, but irrelevant to the stringency of X's claim, then we thereby zero in on the moral significance of X's claim itself. And we can say that the moral significance of X's claim lies in the very fact that--if matters relevant to whether Y may infringe X's claim, but irrelevant to the stringency of X's claim, are or were absent--then it is or would be permissible for Y to infringe the claim if and only if Y would thereby produce a sufficiently large and appropriately distributed increment of good, or advantage, the size of the required increment, and the appropriateness of its distribution, turning entirely on the stringency of the claim.

In light of

The Aggravation Principle: If X has a claim against Y that Y do alpha, then the worse Y makes things for X if Y fails to do alpha, the more stringent X's claim against Y that Y do alpha

we can rephrase the last clause as follows: the size of the required increment, and the appropriateness of its distribution, turning entirely on how bad things are for the claim holder if the claim is infringed.

That is what is central to the constraint on Y's behavior that X's having a claim against Y is equivalent to.

Some things that happened along the way might be worth mentioning just briefly again by way of summary.

I argued that we should accept the Means Principle for Claims, which says that claims, and their stringency, extend down from ends to means. And I argued also that the fact that Y can save X's life does not give X a claim against Y that Y do so.

I offered the High-Threshold Thesis as an account of the distributive constraint on increments of good, or advantage, to be produced by infringing a claim: it says that where claims are concerned, the numbers do not count. And I said that it seems to me no objection to the thesis that it yields that some claims--such as the young man's against the surgeon in TRANSPLANT--are maximally stringent. But I will not assume you join me in accepting the High-Threshold Thesis.

I also said that I would so use "ought" and its cognates that the Requirement-of-Fault Thesis for Ought-Not is not true: I throughout use them objectively, that is, in such a way that its being the case that X would not be at fault for doing a thing is compatible with its being the case that he ought not do it.

NOTES

1. It will be clear, then, that the notion 'good for a person' is the same as the result of constraining the notion 'value' by what I called the narrower idea in section 8 of Chapter 5.

2. Notice that according to the thesis that all claims are absolute, all claims are equally stringent--yet another count against the thesis, if another were needed. Ronald Dworkin says that rights are "trumps", and Robert Nozick calls them "side-constraints". See Dworkin, *Taking Rights Seriously* (Cambridge: Harvard University Press, 1977), and Nozick, *Anarchy, State, and Utopia* (New York: Basic Books, 1974). You might conclude that Dworkin and Nozick think that all rights (and a fortiori all claims) are absolute, thus equally stringent. I doubt that they think this. (I cannot believe they think it impermissible to kick a in the shin to save four lives. I also cannot believe they think it no infringe of a claim to kick a man in the shin. But they are alas not very generous with detail.) In any case, since claims vary in stringency, those attractive metaphors have to go. The most we can say is that rights are more or less high cards or that they are more or less spongy side-constraints, which I grant lack the charm of the originals.

3. They will need qualification, for reasons that will come out later. I hope the reader will take my word for it that the sources of the need to qualify are not relevant here.

4. A survey and discussion of these and other difficulties may be found in Joel Feinberg, *Harm to Others*, vol. 1 of *The Moral Limits of the Criminal Law* (Oxford: Oxford University Press, 1984).

5. I stress that the idea here is that *where claims are concerned* the numbers not count. The idea here is not that the numbers never count in morality. Where matters of mere distribution are concerned, the numbers arguably do count. See John Taurek, "Should the Numbers Count?", *Philosophy and Public Affairs*, 6 (Summer 1977), for an argument that they do not, and Derek Parfit, "Innumerate Ethics", *Philosophy and Public Affairs*, 7 (Summer 1978), for criticism of Taurek's argument.

6. Alternatively, the doctor might instead go on to add "but you really ought not have starved the baby, you ought to have called a doctor". This says that while Bloggs did what he thought his best in the way of careful consideration, his consideration was not careful enough. But this too sheds doubt on the Requirement-of-Fault Thesis for Ought-Not, since it might have been through no fault of Bloggs' own that he thought his consideration was careful enough.

7. As is plain, I think the classical utilitarians (such as Mill and more recently Moore) were right to divorce the question whether X ought to do alpha from the question whether X would be at fault if X did not do alpha, and I follow them in that respect. On the other hand, the classical utilitarians did not allow probabilities to be relevant to what a person ought to do: on their view, what is relevant to what a person ought to do is only what in fact is and will be the case. (It is clear that Moore's way of explaining why probabilities can be ignored rests on his not having seen a difference between probabilities and an agent's estimates of probabilities; see his *Ethics*, pp. 118-120.) I am suggesting that we do well to depart from them in this respect.