

Chapter 2

Reasons for Acting

1. Introduction
 - a. Self-interest and self-sacrifice
 - b. Acting out of prudence
 - c. Doing what is legally permitted or required
 - d. Doing what is standardly done
 - e. Appealing to the Code of Ethics
2. Doing what is ethical
 - a. Getting the facts
 - b. Discounting our biases
 - c. Giving reasons
 - d. Finding true ethical premises
 - i. Tracking harms
 - ii. Ethical theories
 - Introduction
 - (a) Utilitarianism
 - (b) Deontology
 - (c) Virtue theories
 - (d) Comparing ethical theories
 - iii. True ethical premises

Questions

1. Introduction

In the first step of the method of tracking harms, we are to

- (1) Try to understand why the participants are doing what they are doing by constructing arguments that would justify their acts or omissions.

As we have seen, this step involves trying both to understand what a participant was thinking that could explain the action or omission and also to see how the participant could think that the action or omission was justified. We can understand why someone does something without thinking the action justified. One child may hit another who took a toy, and we can try to understand why the one child took the toy and why the other child hit the first one. But understanding why they acted as they did does not mean we think either child justified. If we are to determine whether a person is ethically justified in some action or omission, we need to assess the premises we can plausibly attribute to the person to determine if they are good reasons -- if, that is, they are ethically right.

So (1) -- the first step in our method of tracking harms -- requires that we be clear what kinds of reasons can be given for acts and omissions. Among all the possible reasons someone may have for doing or not doing something, we want to focus on the kind of rea-

son that provides good ethical grounds. We shall find that we can have all sorts of good reasons for acting which are not in themselves ethical, though they may be consistent with what is ethical. As we shall see, it is easy to confuse non-ethical reasons with ethical ones.

It is also easy to think ethical reasons are so distinct from non-ethical ones that one can never have a non-ethical reason for doing something ethical. That would be a mistake. We often have a number of different reasons for doing something. We may eat breakfast because it is good for us, because we like having breakfast, because we can see all the family, and so on. We will then have more than one reason for eating breakfast, and if any one of two or more of these reasons suffices to explain what we are doing, what we are doing is overdetermined. When our doing something is overdetermined, at least two reasons are sufficient to explain our doing it and we would do it even if one of those reasons disappeared. If our eating breakfast is overdetermined, we would still eat breakfast even if everyone else in the family is sleeping late, for example.

Just so, doing what is ethically right may be overdetermined. You may do something ethically right which also serves your self-interest, provides you with pleasure, is prudent, satisfies the Code of Ethics, is standardly what anyone in your situation ought to do, and so on. Though it would be wrong to do what is not ethically right just because you want to, for instance, it is not wrong to do what is ethically right if you do it for the right reason and it also happens to satisfy your self-interest, gives you pleasure, and so on. Indeed, we should all hope for a world in which doing what is right is also prudent, satisfies our self-interest, and gives us pleasure. We do not have such a world, and so, oftentimes, doing what is ethically right conflicts with other interests we have. But when it does not conflict, that does not make it wrong. It makes it overdetermined, reasonable to do, that is, on at least two different grounds.

So when we distinguish between various kinds of reasons for acting and argue that certain kinds of reasons are not ethical reasons for acting, we leave open the possibility that when we act ethically, we may also act for other reasons. It may be thought, for instance, that one must always give up one's self-interest in acting ethically, but that is not true. What is ethical may be just what is in one's self-interest.

The goal of sorting out reasons for acting is to put us in a position to distinguish the sorts of reasons individuals have in particular cases. But applying the distinctions to particular cases is complicated by at least two factors.

First, it is difficult to be sure we know people's reasons for acting. We must make judgments of what they are thinking and believing based on how they act, and these judgments are prone to error. The possibility of an act's being overdetermined means that people may do something ethically right only because it is in their self-interest, for instance. Since a single act may be explained by a variety of different reasons, that is, judging which reason or reasons count for the persons acting is a delicate matter in which it is easy to err.

Second, it is even harder to be sure we have got things right when we are ourselves parties to a case. It might seem easier. 'After all,' we may think, 'who is better positioned to know our own reasons for acting than us?' But what is needed to be sure what we are doing is ethical is objectivity. We need to back off from the case and, as an observer, come to understand the reasons each person involved in the case has for doing what he or she is doing. But if I am involved in a case, that means that I must (a) back off and try to see my behavior as others are seeing it and then (b) try to understand what my reasons are for doing what I am doing. Both of these enterprises is difficult.

(a) Consider, for instance, how difficult it can be to get someone to see that he or she is being abusive or passive-aggressive. Even well-intended men, who say they would not harm anyone for the world, have trouble seeing their anger or sharp words as abusive.

And we can all surely think of examples where someone thought we were angry or sad, say, when we had not realized we were until asked about it. Seeing our behavior as others see it is little easier, if at all, than seeing how we look, front or back or side, to others. When we are a participant in a case, that does not give us a special inside track to understanding the case. It seems as likely to cause us more difficulty in understanding at least our own behavior.

In addition, since we are involved in the case, and so presumably have an interest in the outcome, we may find it difficult as well to be objective about what others are doing. After all, the more disreputable the behavior of others, the better one's own behavior may look, and when the outcome of the case can cause real harm to one's reputation, or when one has not acted well, it seems only human nature to try to put a better face on what one has done. We can do that by downplaying the harm one has done or by magnifying the harm others have caused.

(b) Even if we come to see our behavior as others see it -- and that assumes it unambiguous enough that others see it in a single way -- we have the additional problem of trying to understand our reasons for behaving as we do. Sometimes we do not do well in understanding the reasons for our own behavior. We might think we are acting unselfishly when we are being selfish; we might think we are acting selfishly when we are being unselfish. These are not universal truths, but we find them true often enough in the case of others, at least, that each of us must admit that they may well be true in many cases for us too. The consequence is that we can never be sure we have come to understand correctly the reasons for our actions.

In addition, just as with understanding our behavior and that of others, we may find it difficult to be objective about the reasons others have for their behavior. After all, the more disreputable their reasons seem, the better one's own reasons may look, and when one does not have good reasons or when the case may cause one real harm, it seems again only human nature to try to put a better face on the reasons for one's doing what one has done. There may thus be a natural tendency to see the reasons others have as less than pure as well as a tendency to see one's own reasons in a better light.

But as difficult as it can be in a particular case to sort out our own and others' reasons for acting, we can do no sorting until we have a check list of the kinds of reasons we generally provide for our acts and omissions. We shall thus work our way through some of the kinds of reasons people have for acting, coming, at the end, to ethical theories and the reasons they provide. We are not concerned to provide an exhaustive list of all the kinds of reasons people can have for acting, but the subset of those reasons that we are likely to mistake for ethical reasons or that we will be tempted to use in difficult ethical cases. One aim is to make it clearer what count as ethical reasons and what do not.

a. Self-interest and self-sacrifice

Making an ethical judgment requires being like a judge. In deciding a case, a judge must come to see the situation from the point of view of the plaintiff, to see why the plaintiff has taken the case to court. Then the judge must do the same for the defendant, understanding why anyone should be willing to defend against such a charge. After seeing the case from the points of view of the defendant and the plaintiff, the judge must back off to make an objective judgment, determining not what is in the interests of either the plaintiff or the defendant, but what is right.

When we make an ethical judgment, we must extricate ourselves from the emotions

of a particular situation to look at the matter objectively. But when we are involved in a situation, it is extraordinarily difficult to back off and, just as important, to know that we have backed off and are making objective judgments. It is like trying to be a judge between a defendant and a plaintiff when one is the defendant or the plaintiff. Among other problems, our self-interest can get in the way.

It is particularly likely to get in the way when we are unsure what to do, as we ought to be when facing a hard case. If we are at a loss to know what to do ethically, but know we may cause great harm if we just act, without clarity about our aims, we may think it best to further the interests we know. When we are confused, our immediate self-interest may seem the only clear guide for action.

Some take the opposite view. Confused about what morality requires, and knowing that morality and self-interest are not identical, they adopt the principle that in such a situation we ought to act so as not to further our self-interest. They think we can fish some morality out of the confusion by self-sacrifice.

We may luck into what morality requires when we sacrifice our immediate interests or act on them. Acting out of our immediate self-interest may be just what is required in a particular case to achieve morality, for instance, but this would be luck, and we would do what is ethical for the wrong reason. Doing what is ethical because it is to our immediate self-interest is not to make an ethical decision, but one based on our immediate self-interest.

That is not to deny, as we have seen, that we can do what is ethical while acting in our immediate self-interest. In *Dancing a legal dance*, Mary had to consider that the law requires her to inform if Martha told her of child abuse, that we generally have an obligation to obey the law, and that a failure to inform might thus lead to her having legal problems. So besides having an obligation to obey the law, she had a concern about her own interests that might lead her to inform. She also had reason to believe that less harm would come to Martha if she informed, and appealing to causing less rather than more harm is an ethical reason for acting. So the judgment that she should inform could be backed by an ethical reason and yet also coincide with her self-interest.

Of course, it could also not coincide with her self-interest. Indeed, generally, acting in our immediate self-interest not only may be unethical, it is likely to be unreasonable. At the least we need to distinguish between acting in our immediate versus our long-term self-interest. Being inoculated would harm my immediate self-interest because the needle's going in is painful. Being inoculated may even harm some less immediate interests because sometimes an inoculation produces a painful persistent swelling. But my long-term interests are generally well served by being inoculated. Being inoculated generally means I am protected against having a particular disease and so will not have the pain and suffering and expense it causes. Those gains far outweigh the harms of the inoculation, and it would thus not be reasonable not to have an inoculation just because of the immediate and even less immediate harms it may cause.

b. Acting out of prudence

To be prudent is to act cautiously and in a way that does not harm one's long-term interests. It is imprudent to drink and drive, for instance. In *The death of a baby*, Deborah's not telling the authorities right away what she knew about Hal's suffocating the baby creates a problem for her that she has to consider if she is to act prudently. For her later to tell what she knows is for her to admit that she kept the information to herself, thus misleading

the authorities. A concern for her own well-being should not be a decisive factor in her determining what she ought to do, but the case illustrates how we can complicate our own position, and compromise our own long-term self-interest, by not acting when it may be thought most appropriate to act.

Of course, we may not know what we think we need to know to act in a timely way and may not even know what is in our interests. In Dancing a legal dance, Mary may not know whether it is more prudent to drop the case or not. If she drops it, she may get the reputation of not being able to complete difficult cases and so may not be hired again for such cases. Or she may get the reputation of someone who knows when to cut her losses and move on. The situation is complex enough that she does not know whether going on or staying is in her long-term interests and cannot know how others will see what she does.

She can gain some objectivity about her position by asking what professional interests of hers may be affected. Consider her interests in being Martha's social worker and about whether she ought to tell about the abuse. We tend ultimately to trust those who are honest and open with us. The core of the relationship of being a social worker to someone is trust. The social worker must be trusted by the client, and the client must be trusted to be telling the truth to the worker. Losing that trust damages the relationship. Mary's long-term interests in remaining Martha's worker, and so helping her eventually, argue for her being honest and clear. 'If you tell me about abuse, I will have to report it.' So focussing on her long-term professional interest of maintaining trust so that she can maintain the relationship allows her to back off and ask herself with more objectivity what she ought to do.

To act prudently, we must examine what interests we have in a situation and distinguish carefully between our immediate interests and our long-term interests. To act prudently is to give full weight and value to long-term interests, however powerful and compelling our immediate interests may be. We may be tempted to lie for immediate advantage, but, among other things, our concern about our reputation for honesty ought to weigh heavily against any such temptation.

Acting with prudence has value. For instance, we judge people in part by whether their conduct harms their self-interest. If it does, then, we think, we cannot be sure that they will not harm our interests. If they are not careful about their own interests, why would they be careful about ours? Acting in a prudent way also gives us a broader perspective on what ought to be done. Determining our long-term interests requires us to back off from the immediate situation and so creates some distance for us from the immediate situation. That distance provides some objectivity.

Yet what is prudent is not necessarily ethical, and vice versa. Acting in our long-term self-interest may be ethical, but not because it is in our long-term self-interest. Deciding not to help others because it is in our long-term self-interest is not to make an ethical decision, but a decision based on our self-interest. And acting ethically may harm our own long-term self-interest. There is no guaranteed match between what is ethical and what is prudent.

c. Doing what is legally permitted or required

In Dancing a legal dance, Mary was legally required to report the information about child abuse given her by Martha. So if, to keep Martha's trust, she had kept the information confidential, she would have broken the law. Mary had a conflict between keeping silent and doing what the law requires. Appealing to the law seems to lift the burden of decision-making from a social worker's shoulders and provide a way out of a complex problem.

In Doing what the judge orders, John thinks his interests are doubly protected by his

going through channels. He thinks he is protected because no one can fault him for having made a professionally questionable decision on his own. In addition, the 'right choice' becomes what the law requires, whatever John's self-interest may be and whatever others may think.

Appealing to the law can resolve an ethical problem, but though we may end up doing what is ethically right and protecting the self-interest of those involved, in this case appealing to the law does neither. John's supervisor and the judge knew exactly whom he was talking about, and so, though he never mentioned Jane's name, he identified her as having AIDS. He thus can be faulted for breaching confidentiality. The judge can be faulted for having AI tested without his permission, and John and his supervisor can be faulted both for following the judge's unethical decision and for so manipulating things as to make it appear to observers that the judge is at fault, not they. The case nicely illustrates how appealing to the law can provide a way of resolving our doubts about what to do, but not guarantee an ethical resolution -- even though self-interest, and even prudence, may encourage that resolution.

What is ethical and what is legal may diverge. We may be ethically obligated to do something the law does not require. Someone who owes a debt, is not able to pay it for awhile and also able to convince the person to whom the debt is owed not to sue will not be legally obligated after a certain period of time. The legal obligation will lapse, but the ethical obligation to pay the debt remains. We may also be legally obligated to do something that ethically makes no difference.

But though what is ethical and what is legal may diverge, we are not suggesting that social workers, or any professionals, not do what the law requires. We are suggesting that the law is not the final word on what a social worker ought to do to be ethical. Legislators do not determine what is ethical, and we cannot justify doing what is ethical by saying, 'The law says I have to do that.' We need to add, 'And the law is ethically right in this case.' Some laws may not be ethically right. Martin Luther King certainly thought so when he protested state and federal laws supporting segregation. Thus, though we must take what the law says about a situation into account in deciding what to do, what the laws says is not determinative of what we ought to do.

d. Doing what is standardly done

A person trained as a professional is trained into a set of practices which constitute the standards of the profession and are normative for someone in the profession. A social worker's failure to honor confidentiality, for instance, subjects the social worker, at a minimum, to criticism -- both from the client and from the social worker's colleagues and others.

We get a sense of the normative power of these standard practices by noting what should be a lawyer's first question to a physician accused of malpractice. 'Did you do what any physician would do under the circumstances?' If the answer is affirmative, the lawyer has the first line of defense in hand. 'You are indicting the entire medical profession, not just my client, for any physician would have done what my client did. It's standard medical practice!'

To become a professional we must learn the standard practices, and because standard practices serve as norms for professionals, a professional must follow the standard practices to practice well, it may be thought. They are sanctioned, and so legitimized, by countless professionals having done them. Yet, for all that, they may be mistaken. It may be standard practice for someone to counsel all the members of a family, as Mary does in

Dancing a legal dance, but it puts Mary in a position where she is tempted, and even told by a colleague, to encourage the father to confess. Though confession may be good for the soul, it is not obviously her job as social worker to urge a client to confess. Standard practice does not guarantee an ethical practice.

A professional is always faced with a dilemma, obligated to follow standard practice because those are the norms of the profession and yet obligated to question those norms and to act otherwise if they are wrong. It takes courage to step outside the normal practice of the profession to do something contrary to that practice. We open ourselves to the charge that, as an opposing lawyer might put it, 'It is not the profession itself we indict today, but only this person, who, instead of following standard practice for such a case, insisted on following whims, untried and untested by the countless good professionals who would have saved my client harm by doing what the profession requires.'

An appeal to the standard practice of a profession has weight in deciding what we ought to do because that practice supposedly represents accumulated professional wisdom - the way in which many different persons responded to the same sort of problem, presumably with some success. If they all responded in the same way, despite their personal differences and the presumed differences in the situations they faced, that is some guidance to how anyone ought to respond. But appealing to the standard practice of social workers has its problems.

First, the situation we find ourselves in may be so unique that there is no accumulated wisdom and so no standard practice. How many times have social workers found themselves in a situation like that confronting John in *Doing what the judge orders*?

Second, even if there is a standard practice, we may not be able to find out what it is. Social work often takes place in private. The social worker's interactions with a client are generally unobservable by others, taking place in an office or during a home visit. The requirements of confidentiality may mean that cases never get discussed. We may not be able to discover any similar case.

Third, even if we do, we cannot be sure we have discovered the standard practice. Professionals are in a variety of communities. The professional community of a social worker stretches beyond the confines of the agency employing the social worker, the community within which the social worker actually lives, and the local or state professional organizations to which the social worker belongs. We cannot be sure these communities have the same standard practice for particular problems. What is acceptable in one community may be unacceptable in another. So even if we discover what others have done, we may not know whether it represents the standard way of handling such a situation, or whether it is the way those in a particular community handle such problems. Those who acted that way may have been as lost as we may find ourselves.

Fourth, even if there is a standard practice, and we can find out what it is, we may not know what it means in a particular case. The standard practice may be open to a variety of interpretations, or it may be clear, but we may be unclear how it applies in the case in question. In *Doing what the judge orders*, standard practice may require us to tell Al that he is to be tested for HIV, but we may be unclear whether that means we must ask him directly, 'Is it permissible to test you for HIV?', or whether it is acceptable to ask, 'Is it O.K. if we test you for everything you might have?' The standard practice may be consistent with both alternatives, but the latter would deny Al information he would need to give informed consent. In short, the standard practice may be itself conceptually problematic, consistent with different ways of understanding what is required or permitted, with different implications for what a practitioner ought to do.

So it is oftentimes not possible to appeal to a single accepted understanding of what

is the standard practice for a particular kind of problem. But even in the best of cases, where there is an accepted understanding, doing what is standardly done is no guarantee we are acting ethically. That a group of practitioners in a profession have come to act in a certain way in certain situations may provide the individuals in the group with legal protection, since they can cite the 'standards of the profession' as a justification for their actions, but it surely does not ensure that their actions are ethical. Otherwise, we should have to admit that because slave owners were a group, what they commonly did regarding their slaves was ethical. Or, more pointedly, even if social work practitioners commonly talked about their cases in public, that would not make it ethical.

This point about common practice not being ethical merely because it is the common practice can be difficult to accept for someone trained into a set of professional standards, particularly when those are embodied in a code of ethics that clearly contains ethical pronouncements. But codes of ethics have histories -- have changed, that is, as professions have changed -- and the changes are evidence that at least some of a code's pronouncements are not the truths of ethics some may imagine. It used to be common practice in medicine, for example, for medical practitioners to talk about their cases in public places. It is still common practice for those at admitting desks in hospitals and doctors' offices to ask, loudly enough for others to hear, whether one has insurance and if so, with what insurer. That such practices were or are common does not make them right, and it is a sign of the ethical maturity of a profession to critique its own code of ethics, as the social work profession has recently done, and make the changes it now sees are needed to ensure that the code better accords with what is ethically right. That a code of ethics is ethically better obviously does not mean that it is ethically the best it can ever be. History teaches us otherwise.

But where a code of ethics does embody the standard practice of a profession, one may think that at least it provides clear guidance as to what the profession thinks we ought to do in particular cases. Unfortunately, although appealing to a code of ethics for a profession provides some guidance about the sorts of ethical considerations we should keep in mind in trying to resolve an ethical problem, it is no better than appealing to the standard practice in guaranteeing true ethical premises and not much help in resolving the difficult ethical cases we are considering.

e. Appealing to the Code of Ethics

In the best of cases, codes of ethics contain the essential values of a profession, making explicit the ethical commitments of those in the profession to those whom the profession is to serve. Those learning about a profession may look to its code to understand the profession's mission and how its practitioners see themselves as working to fulfil it in an ethical way. They will also find prohibitions against certain sorts of behavior (like speaking about cases in a public place) and admonishments to further certain goals (like furthering the self-determination of clients).

But what neither those learning about a profession nor those in the profession will find is much guidance for the sorts of ethical problems we laid out in Chapter 1. These problems arise in specific cases where the details make a difference, and no code of ethics can be expected to guide us in everything we must do. And these are hard cases, about which seasoned professionals may disagree, and the Code recognizes that such cases exist and does not attempt to resolve them by fiat.

Consider what John in *Doing what the judge orders* would find were he to look at the

Code of Ethics of the National Association of Social Workers to help him determine what he ought to do when told by another social work practitioner that his client's mother, Jane, has AIDS. In the section of the social work Code concerned with confidentiality and privacy, we find the following rule:

The social worker should protect the confidentiality of all information obtained in the course of professionalen to pass it on in a professional setting unless she has a compelling professional reason. So the Code is helpful in giving us some guidance about what not to do. It does that in telling us what kind of reason we ought to have to deny someone's confidentiality.

It then provides an explanation of what count as compelling professional reasons when it says that

...the general expectation that social workers will keep information confidential does not apply when disclosure is necessary to prevent serious, foreseeable, and imminent harm to a client or other identifiable person...(1.07(c)).

This is the right sort of reason to justify denying confidentiality. Causing harm to others if we retain confidentiality is arguably unethical if the harm is foreseeable and imminent and of sufficient magnitude.

But requiring that we know that harm will occur sets a high standard to meet, and so the Code will be unhelpful in the most difficult cases -- the ones in which we know of a risk of harm, even a great risk, but do not know for sure that harm will occur. Deborah's problem in *1.1 The death of a baby* is of just that sort.

Second, if John is to act to prevent harm, what ought he to do?

- Tell Al that he may be HIV-positive?
- Tell Al that his mother has AIDS and that he may be HIV-positive?
- Tell Al that he may have something, but they are not sure what?
- Not tell Al anything, but tell others who may act for Al?

These are ethically different ways of responding to the problem. The first two respect Al's self-determination, giving him information he would need to make an informed decision. But they differ in that only the latter explains why John thinks Al may be HIV-positive -- although only by implication, leaving it to Al to figure out that John thinks he and his mother are sexually intimate.

The third provides Al with no information about the source of a possible infection or about what the infection is. So it protects the confidentiality of Al's mother, but leaves Al without the information John has that makes informed consent possible for Al. If I am told I have something, but have no suspicion it will lead to AIDS, I cannot willingly consent to a procedure that will test for my being HIV-positive. So when I make a decision, I will have the illusion of self-determination.

The fourth alternative denies Al even the illusion. It is the one John chose, and it denies his self-determination. Others are deciding for him what ought to be done about him.

These four options differ ethically, and John made an ethical choice when he chose

the fourth. Is there anything in the Code that might necessitate or justify that choice? Only one section is relevant:

Social workers may limit clients' right to self-determination when, in the social workers' professional judgment, clients' actions or potential actions pose a serious, foreseeable, and imminent risk to themselves or others (1.02).

This Code gives John permission to limit Al's self-determination, and for almost, but not quite, the same reasons that must hold for the general expectation of maintaining confidentiality to be lifted.

We are justified ethically in breaching a client's right to confidentiality when there is 'serious, foreseeable, and imminent harm to a client or other identifiable person.' We are justified in denying a client's right to self-determination when there is 'a serious, foreseeable, and imminent risk to themselves or others.' Denying confidentiality requires knowing that harm will occur; denying self-determination requires knowing that a risk of harm will occur.

The Code thus sets different standards for denying self-determination and confidentiality. We are permitted to limit a client's right to self-determination when we believe there is a risk that harm will occur and we are 'generally expected' to deny confidentiality when harm will occur. It is thus easier to deny someone's right to self-determination than it is to breach someone's confidentiality. For the former we need only know that a risk of harm will occur whereas for the latter we must know harm will occur. The Code thus makes breaching confidentiality ethically worse than denying self-determination because we need more to justify breaching confidentiality than we do to deny self-determination. So when John chose to deny Al's self-determination rather than breach the confidence, he was acting consistently with the Code.

But it is not at all obvious that confidentiality ought to have more ethical weight than self-determination. The Code implies that it does, but that does not make it true.

And that is one problem. Even if it were true, its being in the Code would not be sufficient to justify John's behavior. Citing a provision of a code does not give the ethical reasons for doing what it says to do. It is not enough to say, 'The Code told me so.'

Second, John cannot properly cite the Code. The Code tells him to break a confidence only for a compelling professional reason, but he only has a compelling professional reason if harm is to occur. But he cannot know that harm is to occur without knowing the results of Al's test. He thus cannot use the harm's occurring as the reason for having Al tested. Indeed, even if we were to lower the threshold and require only a increased risk of harm, he would not be able to act because he cannot know that Al is putting anyone at risk before Al is tested. So John must find some independent basis, something other than the Code to justify having Al tested.

In general, if a client has an ethical right not to have confidentiality breached except for compelling professional reasons, only an ethical reason can be the right kind of reason. In addition, the reason must be compelling enough, must outweigh -- to use a metaphor -- the ethical heft of a client's right to confidentiality. So we always need to ask the following two questions when we are considering denying a client's ethical right:

- (a) Do we have the right kind of reason for doing that?

Wanting your friends to know how difficult the case is that you are working on is not the right kind of reason for denying confidentiality, but even if we have an ethical reason, we

must ask,

(b) Is the reason compelling enough to deny the right?

We have the same reasons for ensuring that what a client says to a social work practitioner remains confidential as we do for ensuring that a patient's concerns are kept confidential by the physician and other medical practitioners who have access to them. Patients and clients need to be assured that they can speak freely, without having their remarks read or passed on to anyone, so that, among other things, physicians and practitioners can be sure that they have all the information they need to give informed assistance. That others may come to know of something the client finds embarrassing or wants to maintain private is an impediment to the client's engaging in the free and open exchange necessary for informed professional practice.

But besides this practical reason for maintaining confidentiality, clients have an ethical right, independently of the practical consequences, that what they tell to a social work practitioner be kept private. The information they provide is about aspects of their lives that are not, presumably, general knowledge to the world at large, and just as we all have a right that others not disseminate private information about us, clients have a right that those whom they see professionally keep private what was private.

But whatever we do, we must also -- always -- ask a third question:

(c) Is our act or omission firmly enough grounded with ethical reasoning that it can properly serve as a precedent for others?

Every decision -- to act or not to act -- creates a precedent others may appeal to in determining and justifying what they do, and so we need to be careful when we make decisions not to provide a precedent that can readily be misused. Otherwise, we may find our decision being cited as a precedent for someone's denying confidentiality, for example, for reasons we do not think relevant or for relevant reasons we do not think weighty enough. The latter difficulty is common to every distinction in which lines cannot be clearly drawn. If a particular reason is judged weighty enough to justify breaching someone's confidentiality, a similar reason almost as weighty may be thought weighty enough even if a closer call. And so on down the slippery slope.

The problem of our judgments about what we ought to do being precedents is compounded because judgments about what ethical weight various considerations have are subject to no clear calculus. Well-intended, well-informed and bright individuals may disagree. And so others may think your decision an obvious precedent for what they ought to do even if you do not, and what counts is not what you think or would say if asked, but how what you did is understood by others -- even if you understand it differently.

In general, to summarize, codes of ethics fail to provide solutions to hard cases. They fail because of what a code is designed to do and because of the nature of ethical dilemmas and problematic cases. A code of ethics is meant to define a profession. For instance, nurses ought to act for the well-being of their patients. That concern is of such importance that it outweighs the self-interest of a nurse. But those in business would not long survive if they took as a primary obligation caring for their customer's well-being rather than their own. A complete and well-articulated code of ethics for a profession would distinguish one profession from another ethically.

Yet professionals can reach agreement about what defines their profession only if the provisions are of a high enough generality that their articulation does not cause disagree-

ment. As soon as we begin to specify the reasons why a nurse is obligated to put the well-being of a patient above his or her own, or begin to specify exactly what it is that a nurse ought to do in such a situation, we discover disagreement. Is a nurse to care for a patient if the nurse is likely to get very ill, with AIDS, for instance? Such a question raises disputes about what count as relevant reasons and about how weighty various relevant reasons are. Social workers must encourage as much self-determination as possible and yet intervene to help others, sometimes making decisions for them about what is in their best interest. Determining how much to intervene to encourage self-determination or prevent harm requires giving reasons and assessing them, and no code of ethics can state how we are to resolve this sort of problem in the multitude of cases in which it occurs. We can see how this occurs in each of the three kinds of ethical problems we can face -- ethical dilemmas, factually problematic cases, and conceptually problematic cases.

In The death of a baby, Deborah's problem was that she did not know what happened and yet had to decide how much she should hold Hal responsible. But merely knowing of a possible harm to others is generally not enough to justify denying a right. When a case is factually problematic, its very nature -- the state of our ignorance -- makes appeal to the Code problematic.

The same is true for conceptually problematic cases. If John decided to break confidentiality, he would need to decide whether a 15-year-old is old enough to provide informed consent and then whether this 15-year-old is an exception to whatever generally holds. He did not do either, but, in any event, the Code of Ethics would be no help. What it is to be 'old enough to provide informed consent' is conceptually problematic, people disagreeing whether the concept is applicable even when they agree on the facts. In addition, codes of ethics fail to be specific enough to be helpful, and, obviously, no code could help decide whether AI is an exception to whatever the general rule is.

Codes also fail to be specific enough to help with ethical dilemmas. Dilemmas are especially sensitive to details. If the harm to one side were just a little more, or were this kind of harm rather than that, we might be justified in deciding one way rather than another. We cannot know until we look at the details, but no code of ethics can provide detailed prescriptions for every case that may arise, let alone do that and be general enough to obtain the general assent of a profession.

As we have stated, codes of ethics state, in one readily discoverable place, what a profession has presumably agreed upon as its essential values and norms. So codes can be used, among other things, to socialize members into a profession and to provide a basis for corrective measures and sanctioning members for violations in order to protect both individuals and the social good. And, as we saw regarding the situation where a co-worker tells you in a lunch line about a case, the Code of Ethics provides clear and helpful guidance about what social workers are and are not to do in some kinds of cases and provides a map to guide social workers about what is relevant and what is not in more difficult cases.

2. Doing what is ethical

When faced with an ethical problem, we ought to proceed in a systematic way to get a grip on what the problem is. No procedure will always lead to a single clear and obviously correct answer, but we ought to get into a position to see what the possible answers are.

a. Getting the facts

As our method of tracking harm requires, we begin by determining what the facts are. This is often not easy. In the death of a baby we may never know why Hal did not awaken while he lay on his son. Perhaps he had an epileptic fit. If he did, it would make a difference in any judgment we would make about him, but if he is unaware of what happened, access to it may be lost forever.

We may also think we know something we do not know. Although Mary Todd has no reason not to believe Martha's story about her father's sexual abuse, and has good reason to believe it because it was substantiated by Martha's sister, still, there is no clinical evidence of such abuse. So we can at best say that Mary has good reason for believing that Martha has been sexually abused, not that Mary knows that.

Getting the facts right thus requires sorting things in different piles -- what we know, what we think we know but may be wrong about, what we have good reason to believe, what we suspect but have no good evidence for, and so on. The borderlines between these categories are not clear, but even if they were, we can be unclear where to put what often does not fit neatly into a category. But we must make the attempt, for that tells us what we need to double-check and what more we need to know. We often find that we were assuming something we thought we knew or that the evidence we have for something is inadequate. For instance, Mary has good evidence that Martha is beating her mother. Her mother has bruises, and Martha says she is beating her. But her mother will not talk, and the evidence is consistent with the father's beating the mother and Martha lying to protect him -- perhaps because he is forcing her to.

Even what seems obvious may not be. A woman was charged with first-degree murder because she stabbed her live-in lover with a butcher knife as he was being held by two policemen. He had been beating her again, and neighbors called. When the police came, they grabbed him, and she went to the kitchen and came back with the knife behind her. As he stood with the policemen holding his arms, she killed him before they had a chance to intervene. It looks like an open-and-shut case of cold-blooded murder, stupidly executed. If she wanted to kill him, why kill him in front of witnesses and the very best sort of witnesses one can get?

Yet appearances can be deceiving. Suppose that he had been beating her for a long time, that this time he said he was going to kill her, and that she felt powerless to prevent him. She knew the police would not keep him long and believed he would return and beat her unmercifully, perhaps to death. She did not have the capacity to keep him away from herself and her children, whom he was abusing. So she killed him while he was helpless -- the only way she thought she could. What at first appears stupid cold-blooded murder may be a desperate act of self-defense. This scenario does not justify her act, but does put it in a different light.

The context makes a difference. Often without really thinking about it, we assume a context for a particular set of facts, but with a different context, the facts look different. One can tell a variety of different stories that are consistent with a single set of facts, and yet those different stories make the facts different. What looks irrational may turn out to be the smartest thing to do in such a situation.

This sort of problem can be particularly acute in social work where a social worker from one culture may have a client from another. What is commonplace and obvious for one person, in one culture, may be odd and obscure to another, in another culture. The Code of Ethics thus mandates that social workers 'should have a knowledge base of their clients' cultures' and 'seek to understand the nature of cultural diversity' (§1.05(b) & (c)).

The same sort of problem can arise when a practitioner is from one social or economic class and the client is from another or when the client is a person with developmental disabilities, for instance, and what seems obvious to the practitioner is not at all obvious to the client. Getting the facts does not mean listing bits and pieces of information, but understanding the full picture, understanding as fully as we can the context within which those bits and pieces fit.

b. Discounting our biases

In *Adoptive Children* Dena had to make judgments about who should get her help and who should not. She decided to give help 'only to those who would be sensitive to the needs of their natural parents' and then had to determine who was likely to be sensitive and who not. She decided that those who work to change the law show they are not selfish and that those who work in support groups for adoptive children show they care about others. So she used those two factors to determine who should get information, but the way she used these factors in regard to the brother and sister creates a concern about bias.

She decided they should not get information, and one of her stated reasons was that they came to her without an appointment. But that seems a trivial and mistaken reason for treating them differently than others. It is a presumption of justice that like cases be treated alike, and so it would be unjust to treat anyone needing help differently from others for reasons that make no ethical difference. Dena did treat the two differently, and so she needs a good ethical reason for doing that.

But coming to see a public official without making an appointment is not evidence of a lack of concern for the official and thus of a lack of concern for others. They may not have thought about needing an appointment to see a public official. Or, having thought about it, they may have been hesitant to call because they were afraid she might not see them given what they wanted.

If we are to come to a proper ethical judgment, we must discount any feature of the situation or of ourselves that may cause us to bias our judgment. But Dena seems unduly sensitive to slights. So some feature of her may be biasing her judgment. And without checking, Dena does not know why the two did not make an appointment. Their failing to do so may be due to an insensitivity to the interests of others, but it also may be due to their youth, or to their concern about being rejected should they call first, or to their failure to realize one calls a public official to make an appointment -- a failure common, we should think, to many of us. So a feature of the situation may be misleading Dena. The two may seem impetuous and selfish when they may be cautious and concerned.

Dena's other criteria for helping is that those who ask her are working to change the law and in support groups. But that may be to hold the brother and sister to too high a standard. They are young and perhaps politically inexperienced. Working to change the law may not have occurred to them, or they may already be jaded enough by the political system to be skeptical about the possibility of changing the law. Why should their immaturity, or their competing political judgment, be held against them? Similarly for not working with a support group. They might think it important, but going to school and working to support themselves, they may judge they should do it later. Why should a judgment about timing or priorities be held against them?

What Dena needs to know are the reasons the two have failed to do what she thinks they ought to do. Otherwise she is not entitled to make a judgment that excludes them from access to information she is giving others.

We may explain Dena's hesitations by her prudential concerns. Setting high standards for those to whom she gives illegal information is a way of protecting herself. She may think that those who care about the well-being of their natural parents, and are concerned to change the law, are less likely than others to tell on her.

Yet the remark about the brother and sister not calling to make an appointment arouses a suspicion that she is not properly applying her own stated criteria for determining who is to get help and who is not. The judgment she is making not to give them help may be determined by her irritation at being treated in a way she considers impolite. But if she has decided to give information to those who deserve it, she should take into account only what is relevant. What is required in making moral judgments is that the judgment be objective, not in any way biased by subjective factors.

It is not just ethical judgments that require objectivity. We ought to examine how much money we have in our checking accounts by checking on what we have put in and taken out, not by fantasizing about what we want there. Objectivity is not usually problematic when checking how much money we have in our account, but checking our ethical judgments is more difficult. They tell us how things ought to be, not how they are, and so we cannot check them against the world to see if they are true. They are also peculiarly prone to be biased by individual, social, and cultural biases.

A major difficulty in making ethical judgments is that some feature of ourselves may get in the way of being objective, and we are the worst positioned to realize it. In making ethical judgments about others, it may matter that one is white, or that one is female, or that one was born of relatively well-to-do parents. We may so look at the world, and thus so look at the problems we face, that even what we see is biased by those features of ourselves. What we must do to make a proper ethical judgment is to discount anything about ourselves that might bias our judgment. We must ignore our own economic position in making a judgment about someone with more money or someone with less. Neither envy nor sympathy ought to bias us. If we are having a bad day, that ought not to make a difference in how we treat our clients. Someone who comes on a bad day ought not be penalized for such bad luck.

The question to ask oneself in coming to an ethical judgment is always whether others, differently situated regarding the relevant feature, would make the same judgment. Dena should thus ask herself whether someone else, not offended by someone's not calling to make an appointment, would judge as she did. Asking whether someone else, without that feature, would agree is a way of assuring a judgment's objectivity. Just as we want, as it were, to look past our own personal characteristics to see a situation as others would see it, so we want to examine and discard whatever external features may prevent us from seeing the client in an unbiased way. We are trying to get rid of the contingent features of the situation that ought not to make a difference to the ethical judgment we make -- that we got up on the wrong side of the bed in the morning, that the person is from another culture, and so on.

The problem of discounting our biases to try to obtain objectivity permeates every aspect of understanding an ethical problem. Even getting the facts right depends upon understanding the context within which they occur, and that context can be determined by cultural features, for instance, not easily accessible to someone from another culture. We must try to understand the situation from within that context and then back off from it, and from our own, to determine what we ought to do.

c. Giving reasons

When we give reasons for doing or not doing something, we provide arguments. That word may make us think of people with loud voices disagreeing to no good effect. But an argument is a set of premises that provide support for a conclusion, and though people do use arguments in disagreeing with each other, they also must use arguments to provide reasons for doing what they think they ought to do. If Mary, in *Dancing a legal dance*, is not to engage in therapy with both the father and Martha, she needs reasons for choosing one or the other.

But the reasons given in the premises of an argument can fail to support a conclusion in two different ways. They may fail to support the conclusion because even if true, they are not relevant to the conclusion. Or a premise may be relevant, and so could support the conclusion, but does not because it is false.

The first kind of failure is one that depends upon the relation between the premises and the conclusion. Could those sorts of premises support that conclusion? This is a question of the form of the argument and has nothing to do with whether the premises are true or false. For example, in the following example of what is called a deductive argument the premises could support the conclusion:

All social workers are overpaid.
All overpaid workers are lazy.
Therefore all social workers are lazy.

This argument is valid. The premises are such that if they were true, it is necessary that the conclusion be true as well. A mark of a valid argument is that we can put different premises of the same form in place of those in the argument and, no matter what premises we put in, we can never have true premises and a false conclusion. A valid argument preserves the truth of the premises so that we never go from truth to falsehood.

But we must start with true premises if truth is to be preserved, and the premises in this argument are false. Social workers are not all overpaid, and not all overpaid workers are lazy. So these premises do not succeed in proving the conclusion true. A good argument, one that gives good reasons for a conclusion, has true premises that provide support for a conclusion.

Consider John's conclusion in *Doing what the judge orders*. He concluded that he would not be responsible for Al's being tested for HIV because the judge told him to do it. So his argument goes roughly like this:

The judge told me to have Al tested.
We are not responsible for what we are told to do.
Therefore, I am not responsible for having Al tested.

These premises, if true, would support the conclusion. If we are not responsible for what we are told to do, we are not responsible for doing something we are told to do. So the argument is valid. But is it true that we are not responsible for what we are told to do? Surely John does not think we should never question orders. Suppose someone ordered him to murder Al. If we are to understand why John thinks he is not responsible, we must understand him as implying something like this: when we are told to do something by someone who has the authority to tell us what to do, and what we are told to do is ethically right, then we are not responsible. But though this new premise is more likely to be true, the con-

clusion is not supported by it unless it is right to have Al tested without his knowledge and permission. John must be assuming that it is. But a judge's ordering it does not make it ethically right.

Trying to understand how John could argue for his conclusion allows us to see more clearly what is at issue in the case. As we try to see how John could have true premises for the conclusion he draws, we uncover the assumptions John is making that we must assess - including the premises John must take to be true. Such premises purport to be true of the world, and so we can look to the world to see if they are true. Premises that tell us what we ought to do does not tell us how the world is, but how it ought to be. The question arises, then, how we determine the truth and falsity of such premises.

d. Finding true ethical premises

As we have seen, we cannot be sure that what prudence, or the law, or standard practice tells us is ethical. Even if what is ethical is prudent, legal, and standard practice, it is not ethical for those reasons. And appealing to the Code of Ethics does not give us sufficient guidance in difficult cases. So we need to determine on what basis we are to judge what is ethical.

i. Tracking harms

In each case we have had so far, the practitioner was trying to prevent something harmful from occurring. In *Dancing a legal dance*, Mary was concerned, among other things, with preventing the father from further harming Martha. In *The death of a baby*, Deborah was concerned that no matter what she did, she would cause harm to the family, either by saying nothing when something should be said or by saying something when nothing should be said. In *Doing what the judge orders*, John tried to work through a minefield of different harms to his client, to his fellow social worker, and to the innocent persons who might become infected with AIDS. And in *Adoptive children*, Dena thought the law harmed those children and their natural parents who would want to be reunited, and she tried to give information only to those she thought would not cause harm.

It is no accident that each case involves minimizing harms. Social workers and their clients are not often graced with the problem of choosing between different goods, and when they are, though the choice may be difficult, it is hard to complain about a mistake. The problem that most often confronts them is that of preventing harm and, unfortunately, of trying to determine which is the lesser of two or more harms.

One helpful way of entry into difficult ethical problems is thus to track the harms that have occurred and are likely to occur. This is why our method is designed to ferret out harms. In *Dancing a legal dance*, the number and varieties of harms are enormous -- the physical and psychological harm by the father to his daughters, the deterioration of what is left of the family, the physical and psychological harm to the mother as Martha batters her, the harm to Martha and others caused by Martha's deterioration, what Martha took to be the deception by Mary in reporting what Martha thought she was telling her in confidence, the consequent loss of the relationship of trust between Mary and Martha, Martha's possible lack of trust in anyone who may now try to help her, the harm to others who are not being served because of the time spent on this family, and so on. Many different persons -- Martha, Kathy, the mother, perhaps even Mary, who may get a reputation for not being able to

do such therapy well -- were harmed or are at risk of being harmed, in a variety of different ways -- physically, psychologically, socially, and so on.

Suppose Mary were considering whether to deceive Martha about having to tell. What ought she to take into consideration? First, there is the deception itself. We are harmed by deceit whatever the consequences. Second, the relationship a social worker has with a client requires trust. Mary must weigh whatever good reason she thinks she has for deception against the likelihood that Martha will no longer know when she should believe Mary and when not. Third, Martha has already been deceived by her father, someone she should be able to trust. So in deceiving her, Mary is making it that much harder for Martha to trust anyone. Martha will learn that even someone professionally obligated to be trustworthy is not. Fourth, Martha will have been denied the opportunity to decide for herself whether to tell Mary what her father is doing.

As the list of harms grows, some immediate and some long-term, it becomes harder to justify Mary's not telling Martha what she is legally obligated to do. As we saw, the harm she is trying to prevent -- Martha's being abused by her father -- may not be prevented even if Mary gets the information she needs to go to court, and the harm that may result from the deception may be much worse -- the continuation of the abuse and the loss of anyone Martha feels she can trust to help her. The dilemma Mary may have had, to tell or not to tell, is easier to resolve if we list the harms and consider whom they will affect, for we can more readily determine what course of action is likely to minimize the harms that may occur. Tracking the harms allows us to see that she ought to be honest with Martha because the gain in not being honest does not outweigh the harms caused by the deception.

In making this judgment we are implicitly appealing to an ethical theory, and it is time to turn to an examination of the most plausible candidates. It is these theories which provide backing for the ethical premises we must use in making ethical judgments.

ii. Ethical theories

Suppose again a situation that has as many footholds for immorality as we can find - a person intentionally doing something to someone else that is itself harmful to that person's interests and will predictably cause other harms, the harms are of great magnitude, the person causing the harms knows that what he or she is doing is harmful and does it because it is harmful, and what is done sets a precedent for others to cause similar harms in similar situations. We need to consider the harms of both the act itself and its effects, including the systemic harm that may result from an act's being a precedent. Within the act itself, we need to separate out the intention of the person who acted (or omitted acting when acting was called for), the knowledge or ignorance of that person regarding what was done (because we might do something intentionally, but without realizing that what we did was harmful), including the capacity to predict the normally expected consequences of the act, the kind of act that was done, and the magnitude of the harms committed.

With so many ways in which harms can occur, and so many different kinds of harm to compare with and weigh against one another, we can understand how saying we ought to minimize harms would be unhelpful in some cases. Not every issue will work itself out in the way we worked out what John ought to do regarding Al in *Doing what the judge orders*.

A major source of difficulty is that different ethical theories hold up different visions of what we ought to do. An act may further a particular vision, may fail to further it, or may positively harm its realization. But what sets back the realization of one vision, and so is a harm to it, will not necessarily set back another vision, and so will not be a harm according

to it. Indeed, what is a harm by one theory may not even look like a harm by another or may look like a different sort of harm, of a different kind and different weight. So a participant who sees a situation according to one moral theory may not even see a problem -- see something as a harm -- that someone else, with a different vision, takes to be central. The situation is like that in the example we considered where a woman stabbed her lover while he was being held by two police officers. What appeared irrational on one interpretation appears more reasonable, though desperate, under another. So one source of disagreement people may have in grappling with cases is that they presume different ethical visions, perhaps without realizing it.

We will find that many of the cases we consider lend themselves to resolution even when they are viewed in different ways by different theories, for often the theories produce the same result, though for different reasons. Each of the theories we consider provides reasons for minimizing harms, for instance. But sometimes the best we can do is to lay out how the competing ethical visions produce different ethical results. The difficulty we will have in resolving some hard ethical problems is that we have no way to choose between these competing visions.

One main divide between ethical theories regards the ethical import given to what happens in the world versus what is internal to the person acting in the world.

On the one hand, ethical acts are acts. When I turn the knob of a door, I act for some end -- to open the door to go out, to see if the door is still stuck. By turning the knob, I may cause it to break, or to become stuck, or to open the door so a person lurking outside can enter. Ethical acts and omissions are like any other acts and omissions. They have consequences in the world, intended or not.

An ethical theory may take as ethically crucial what happens in the world and thus what is done to other people. After all, the argument would go, if by our acts we harm others, we have acted unethically. Any truly ethical theory, it is claimed, must take as of first importance how it is that we treat others when we act.

Yet, on the other hand, it is arguable that we can only be held ethically accountable for what we are responsible for. If I become so ill that I faint, and in fainting fall against you, I cannot be held responsible for falling against you or, if I am, only to the extent that I am responsible for making myself so ill that I faint. But if we are only ethically accountable for what we are responsible for, then, the argument goes, we can only be held accountable for what we decide to do. We surely cannot be held accountable for the consequences of our actions that no one could reasonably have predicted, and it is not even obvious that we should be held accountable for our actions. We have all been in situations where we thought we were doing one thing, but were interpreted as doing something else. We cannot be held responsible for how people interpret what we do, the argument continues, and so, at best, we can only be held accountable for what we decide to do -- independently of our act and of its effects in the world. What is ethically crucial on this view is not what happens in the world -- not even the act of the person, but something about the person who acts.

But if we examine what happens when someone decides to act, we discover much complexity. The person may think through what they are doing or not, be motivated by hate or by love in doing the very same thing, act in character or out of character, and so on. A main divide within this complex terrain of a person is between what a person intends to do and the person's character. Even a person with a bad character may intend to do something good, for the right reasons, and the person should get ethical credit for the intention, despite the bad character. So, it may be argued, it is the intention that matters ethically. But, someone may argue, anyone can do something good. It takes real commitment to forge a good character, to make of oneself a person whose very constitution is good and who acts

ethically out of a deep abiding commitment and not some transient intention. On this view, what matters is the character on the basis of which we act.

These disagreements turn on disputes about what kind of world we ought to live in, on competing visions of an ethically ideal world, and we can best sketch out the competing theories by imagining ourselves being asked to create an ethical world, a world that we would recognize as an ethical ideal, but a world that we could achieve, a realistic possibility given the nature of human beings. There are three main contenders -- utilitarianism, deontology, and virtue theories. Each has its strengths, and each has its problems, some of which we shall note.

We shall distinguish these theories in terms of the following different features they bear:

- competing visions of the kind of world we ought to live in
- based on different natural features we have,
- picking out different facts in a situation as being ethically relevant and important,
- giving different sets of individuals ethical standing, excluding some kinds of beings from having a say in how they are treated ethically and including others, and
- providing different ways of justifying what we do or omit doing.

It will help to keep these differences in mind as we go through these theories, marking off one from the other in terms of their competing visions, with all the differences those entail.

We should remark that it is impossible to state these theories in a way that is not contentious. Almost every word and phrase has been subjected to extensive criticism, and one is hard-pressed to find a neutral, natural and simple way to state the theories without appearing naive. In addition, we are slicing through these theories in a way that is different than the usual because we think that angle of cut best exposes their complexities and differences, opening up the theories for examination and comparison. We think our analysis particularly helpful for those engaged in using these theories for cases they face in their practice.

(a) Utilitarianism

Suppose that you are awakened one night, or think you are, by a voice representing itself as from some superior being, and the voice says, 'I have chosen you to create a new ethical world. I tried it once before with a flood, but I failed. I want you to tell me what the ideal ethical world should be.'

The power to create a whole new world does not come to us, obviously, but we do have the power to help create a different world by what we do, and fail to do, in this world. But without a vision of the kind of world we ought to have, we cannot know the point of our acts and omissions or what we ought to do, and not do, to achieve this new different ethical world. So this voice we hear one night is asking us a question we all need to ask ourselves in any event, 'What kind of world, ideally, ought we to have?'

This is a question about an ethical vision. What kind of world ought we to live in? It is also a question the answer to which presupposes something about us, about the kind of world we human beings are capable of having. Whatever ideal we are to strive for, it must be an ideal we can achieve, a realistic possibility, as we have said, given the nature of human beings.

We all share an aversion to pain and a preference for pleasure, for happiness, we

may say. And we have the capacity to contribute to the happiness of others -- and, unfortunately, their unhappiness as well. Surely it would be a better world than it is if everyone were happy and no one were unhappy. But that seems impossible given the nature of the world and human nature. As long as resources are scarce and human interests conflict, not everyone can be happy. The triangles of love are not going to disappear, and when things are settled, one (or more) of the lovers will be unhappy. So a more realistic end would be a world in which the greatest number have the greatest amount of happiness they can have. If we cannot make everyone as happy as can be, we can arrange the world so that many more are happy than are now happy. That is an ideal, and it looks achievable.

It is the vision that utilitarians hold up before us as the ethical ideal, the end we all ought to strive for when we act or think about not acting. It is a vision we all presuppose in regard to at least some of what we do. Suppose that on a Friday evening you and six hungry friends each contribute \$3 for a pizza, and you are to call in for it. Suppose as well that you love sausage and that your two best friends there love pepperoni and that you have before ordered pepperoni and sausage for the three of you. So you do that again. But four of the group are vegetarians, and you know it. They would rightly have a complaint against you. You took their money for the dinner and spent it on something they will not eat. You might have been so caught up in the idea of having the pizza you and your friend wanted that you just forgot about the others. That would be bad enough. But suppose you remembered what they would not eat and said to yourself, 'Well, it's their fault. They should have known better than to let me order.'

Your four friends -- or former friends -- would think you wrong in what you did. Presumably you had a choice of other toppings and a choice to order two or more smaller pizzas, one with pepperoni and sausage and the other vegetarian. But if you knew that four in your group were vegetarians, your last choice ought to be pepperoni and sausage for everyone, even if you love it. That is a choice guaranteed to produce unhappiness for the greater number -- four of the seven of you.

You could do worse. It could be that everyone in the group dislikes anchovy, even you, and that an anchovy pizza is delivered without your realizing it. Everyone's money is gone for something no one will eat. That would be a worse situation -- more would be unhappy -- than choosing a pizza that would at least make three of you happy. That if you had the choice you would reject without even thinking about it is evidence of how powerful the utilitarian vision is. It permeates much of our ethical life.

The judgment your friends would make about you were you to choose pepperoni and sausage is the same judgment everyone would make about someone given the choice to create a new world who chose a world in which as many as possible were as unhappy as possible. We would be almost as upset at someone's choosing a world in which there was less happiness than there could be. We might not conclude the person was evil, only stupid or careless, but our judgment would have moral weight. Given the choice to create a new ethical world, who would choose anything other than a world in which everyone was as happy as could be. If you chose a world in which everyone was as unhappy as could be, we would think you either evil or crazy. What else could explain your choosing such worlds over a world in which there was the greatest happiness for the greatest number?

The power the vision of a world where there is the greatest happiness for the greatest number comes in part from our being the kinds of beings who, like all sentient beings -- whatever living things can feel -- avoid pain and strive for pleasure. Pleasure is a good and pain an evil for all sentient beings, and it is thus ethical, where one has the power, to produce pleasure for sentient beings and mitigate pain. It is thus unethical to kick dogs, and to teach a child to gain pleasure in torturing animals and other children is to go a long way to-

wards creating an ethical monster. Kicking dogs or teaching children to torture helps create a world with less pleasure and more pain than need be -- a world that is unethical to the extent that it could have more pleasure and less pain for more beings than it does.

The utilitarian vision is powerful, giving us a view of what could be that is itself attractive and would require immense changes in our world to realize. It provides, that is, an appealing target to aim for in going from where we are to where we ought to be. We live in a contingent world, a world that will change through our actions and inaction, and we can readily imagine living in worlds we would want to change and ought to change -- worlds that are less than ideal according to the utilitarian vision.

We could have a world in which a large percentage of the population does not have basic health care. Or a world in which the distribution of wealth is so skewed that some have more money than the gross national products of most countries while others are homeless. Or a world in which children go to bed at night without enough to eat. Or a world in which many die before they need to because of preventable diseases. These are all possible worlds, a kind of world we could live in. But their features are contingent. That is, they can be changed, and to the extent that we are committed to the utilitarian ideal, they ought to be changed so we can live in a world closer to the ethical ideal.

As the examples just given illustrate, it does not take conjuring for us to imagine a world that is less than the best. Our world will do quite well as an example of a world less ethical than it need be. Far too many in our world have far too little chance at happiness. We could distribute what resources we have in an ethically better way, and we ought to do so, utilitarianism tells us, ensuring that we help create a world in which as many people as possible are as happy as possible. Whenever we act or fail to act, our acts and omissions could be, and ought to be, means to realize that end.

The utilitarian vision thus gives us a test for determining what is ethical and what is not. What is ethically relevant -- what we should take into consideration in deciding what to do -- is what helps or hinders the realization of that end. If an act or omission decreases the likelihood of achieving the greatest good for the greatest number, it is harmful and unethical; if it increases the likelihood, it is beneficial and ethical. Everything else is ethically irrelevant.

So it does not matter what you intended to happen. If your child yells at a sibling, and your ear happens to be next to the yell, it is no good your child saying when you yelp, 'Well, I didn't intend to yell in your ear.' Your child did yell in your ear, and it is the effect that counts, not the intention. This is a vision that emphasizes what happens in the world. Good intentions are not enough. Utilitarianism tells us to ignore people's intentions and concentrate instead on how actions affect the participants, determining whether what is done in fact increases or decreases the happiness of the greatest number, whatever the intentions. We need to see results, always acting to maximize the happiness of as many persons as we can.

Because we are concerned with what actually happens in the world, utilitarianism argues, we can have objectivity in our ethical judgments. If intentions were to matter ethically so that the same act could be good or bad depending upon the intention with which the agent of the act acted, we would have no way of knowing for sure what would be a good act and what a bad act. The person who acted might know -- although we all sometimes act without a clear intent sometimes -- but we could not know. But if we judge whether an act is ethical by what we can all see -- it and its effects in the world -- then objectivity in ethics is possible.

We could also learn to be more ethical. One of the most difficult things to do when we are involved in an ethical problem is to back off and not to let our self-interest dictate or

skew the right solution. In deciding what to do, that is, we are not to let our own happiness count for more than anyone else's, and that can be difficult. Nothing in utilitarianism, or in any ethical theory, allows us to be sure that when we judge, we are not biasing our judgments in some way. But because we can see the results of our actions and determine what effects they had on the amount of happiness produced, and for whom and how many, we can at least judge after the fact whether we made the right choice. That will help us make better judgments the next time.

Looking at cases as matters of utilitarian calculation can provide us a handle on how to respond. For instance, Dena's problem in *The death of a baby* can be readily construed as utilitarian. She is faced with telling or not telling what she knows, and the difficulty she has in making that choice turns on the consequences -- the additional harm to a family already traumatized by the death, for example, or the risk of harm to the other children. Her intention is to do good and will be the same regardless of what choice she makes. What matters is not her intention, but what to choose. What utilitarianism tells her is to choose that alternative with the least harmful consequences for the fewest number. Only that choice will help produce more rather than less happiness for the greatest number.

So utilitarian theory gives her guidance, it would seem, about what she ought to do. It sets an end point that is admirable and tells us to look at our actions and omissions as means to that end. So Dena should ask, 'Will doing this or doing that best serve to that end?'

Unfortunately, things are not so simple, and utilitarianism is open to a number of different problems which make it far more difficult to use than we have so far made it appear.

First, the vision is murky. Suppose, when we are going to order pizza, that we discover that the only choices available will make some members of the group unhappy, the different choices producing differing degrees of unhappiness for differing numbers of individuals. One choice is non-vegetarian, and four will be unhappy because they will not be able to eat, but the other three will be extremely happy because that is their favorite meal. Another choice is broccoli, which no one particularly likes, but which all will eat. So the greater number -- all seven -- will be somewhat happy, but not particularly so. Another choice is a topping which five like fairly well and which the two others will eat, but do not like at all. Which choice is mandated by the end of achieving the greatest happiness for the greatest number?

Are we to work to secure extreme happiness for a few, letting their great happiness outweigh the failure of others to have any happiness? Are we to work to secure some small measure of happiness for as many as we can, letting the numbers provided happiness be the determining factor and ignoring the degree of happiness achieved? Are we to work to secure some higher degree of happiness for the majority of the individuals involved as long as the rest are not unhappy?

All these options are compatible with the goal of acting to produce the greatest happiness for the greatest number, and so that goal, far from being a clear target we can aim for, is consistent with so many different acts that it gives little guidance at all.

Second, even if the end were clear, utilitarianism tells us that we are not to concern ourselves with the good of each and every individual, but with the greatest good for the greatest number. It tells us that it is ethically permissible, even ethically required at times, to choose an arrangement that will produce unhappiness for some if doing that will provide the greater happiness for more. Some have tried to justify slavery, for instance, by arguing that it made the slave owners particularly happy and that the slaves were better off too -- a way of distributing the benefits and burdens of freedom and wealth that uses the slaves for the good of others as though they were simply instruments to provide greater happiness for

their owners. But that seems wrong on its face -- like lying to someone to get them to do what you want them to do.

Third, utilitarian theory draws part of its strength from its vision of the kind of society we ought to live in and part of its strength from our desire for happiness. It seems happily to combine idealism with a realism about what we truly desire. But it fails to be realistic about how we can make decisions about what we ought to do. A pure utilitarian world would have in it only acts and omissions that furthered the utilitarian vision, but we simply cannot calculate the effects of each and every act before we act. Sometimes we must make split-second decisions about what to do -- whether to swerve to avoid an animal that darts out between parked cars -- without taking the time to calculate out the effects of that act versus any other. Swerving may put us in the path of an on-coming vehicle, for instance, and so cause worse problems.

In response to that problem, among others, a distinction is made between act- and rule-utilitarianism. We have been examining what is known as act-utilitarianism. It tells us that every act or omission ought to be directed towards the vision utilitarianism holds up before us. But we often act on principle, act, that is, on such admonitions as 'Honesty is the best policy' without stopping to calculate whether acting on honesty in the particular case in question will in fact lead to the greatest happiness for the greatest number. The admonition acts as a rule for our behavior, telling us that we ought to be honest. Following such rules, it is argued by rule-utilitarians, will generally lead to the greatest happiness for the greatest number. So we need not calculate out the effects of every act and omission. We instead act on the basis of rules. One effect of being a rule-utilitarianism is that we have a usable guide to action. Another effect is that we will sometimes act in a way that is ethically wrong. Honesty may generally produce the greatest happiness for the greatest number, but sometimes it will not. Lying will and is thus the right thing to do. But we must follow the rule, and it tells us to tell the truth. We thus do what is ethically wrong. So we gain a theory we can use at the cost of countenancing some unethical acts.

Fourth, assessing the full effects of acts or rules seems beyond our capacities. We seem always to be discovering unintended consequences of even what seem to be the best of acts; even the smallest of acts can have enormous and unpredictable consequences far into the future; and we are often surprised by how others respond to what we do. If we are to determine what we ought to do by assessing the consequences of alternative courses of action, that is, it appears we are being asked to do the impossible.

Fifth, even if we could predict the future as well as what would be the future were we to act differently, utilitarianism asks us to make comparative judgments we cannot readily make. When utilitarianism was first introduced as a full-blown theory, the assumption was that pleasure was quantifiable. Having our favorite dessert is more pleasurable than many other things, and on a scale of 1 to 10, with 10 being the most pleasurable, it is closer to 10 than 1. But when John Stuart Mill articulated the theory, he made a distinction we all make between the quantity and quality of pleasure. Reading a good book can be pleasurable, but it is a far different kind of pleasure than a runner's high. Trying to compare the two by quantifying the pleasures is like trying to compare apples and oranges. They are different kinds of pleasure, qualitatively so different that we do not capture the pleasure of either by treating them both as being of the same kind and quantifiable on a single scale. So the goal of acting so as to produce the greatest amount of happiness for the greatest number is unclear in another way: are we to choose higher quality pleasures, perhaps less intense, over lower quality pleasures of high intensity, or vice versa?

Sixth, although utilitarianism purports to be objective, whether it includes the quality as well as the quantity of pleasure or not, what pleasure or happiness we get from some-

thing depends upon the beliefs we have. If we believe it wrong to drink coffee, no coffee no matter how tasteful will produce pleasure or, if it does, will produce it along with guilt or shame for having drunk it against our beliefs. But then we cannot determine what is ethically right by appealing to the end utilitarianism envisages. For that end will itself be determined by our beliefs, some of which may be ethically suspect. If I believed that men and women were fundamentally different and that society ought to be arranged so that each is 'in their proper place,' women in the home, for example, men at work, then I would take no pleasure from a proposal that required equality of opportunity regardless of sex. And if enough people believed that, then a comparison of a social arrangement that guaranteed equality of opportunity with one that put men and women 'in their proper places' would tell us that the latter, not the former, produces the most happiness for the greatest number. In short, the utilitarian criterion does not look to be objective and so independent of false ethical beliefs, but itself determinable by beliefs -- which may well be unethical.

Seventh, early on in the description of utilitarianism, we wrote about 'pleasure, or happiness.' The two are not the same, and utilitarian theorists may differ as to which is the proper ethical end: are we to try to produce pleasure or happiness? Every sentient being -- from chipmunks to humans -- can experience pleasure and pain. But only some sentient beings -- people certainly, perhaps some other animals -- can experience happiness. So if we are to provide a world in which the greatest number of beings have the greatest happiness they can have, we leave out of consideration those beings who experience pleasure and pain but not happiness. It is difficult to pick a non-contentious example of such a being. Perhaps chipmunks and cows cannot be happy, but can experience pain and pleasure. But they would not have ethical standing in a world designed to secure happiness. They would not be part of our ethical world and would have no claims against us if we arranged matters in such a way that they were caused pain. What matters ethically in such a world is happiness, and if it furthers the happiness of those beings capable of happiness to cause pain to beings incapable of happiness, so be it. Carnivores may hold this view. Vegetarians may not. In short, the scope of utilitarianism is unclear. Which beings does it cover and which does it not? It matters enormously to such issues as how we human beings are to relate to the rest of sentient creation.

So utilitarianism is not without serious problems, but it presents a powerful vision of what kind of world we ought to live in, a vision that, were we to accept it, would obligate us to change much in this world. We can summarize it as follows:

- It tells us that we should hold before us as a vision to create a world in which the greatest number of beings have the greatest pleasure (or happiness).
- It bases its vision on sentient beings having the natural capacity for pleasure and pain (or happiness and unhappiness), wanting to avoid pain (or unhappiness), and wishing pleasure (or happiness).
- What it picks out of a case as being important and ethically relevant are those features of the case that make a difference to creating a world in which the greatest number of beings have the greatest pleasure (or happiness).
- It gives an ethical hearing to all those beings who are capable of pleasure and pain or, if it is restricted to those beings capable of happiness, to all those beings with that capacity.

- It justifies an act or omission in terms of whether it is the choice most likely to produce the greatest pleasure (or happiness) for the greatest number.

(b) Deontology

The death of a baby is problematic because Dena must act in ignorance. That is why she is hesitant. She does not know which alternative is least harmful. She does not have enough evidence to determine the consequences of the alternatives facing her, and even if she knew which was most likely to cause the least harm to the fewest number, her choice may not have the effects she hopes for. Too many factors not in our control can intervene to change the course of events in ways we could never have foreseen.

We cannot even be sure how our best-intended acts will be received. What we mean to be kindness may be construed as meddling. Empathy may be taken as pity. Indeed, we cannot even be sure that what we end up doing will be what we intended to do. We all have made mistakes where we intended to do one thing (leave the car door unlocked while we check out something) and do something else (e.g. lock ourselves out).

So, the objection goes, why ought we be held responsible for what actually happens because of what we intend to do? What matters, according to deontology, is what we intend to do -- what we will to happen, as Immanuel Kant would put it. If what we will is ethical, we ought to get ethical credit, whatever actually happens. Indeed, the argument continues, it is an ethical mistake to consider consequences in deciding what to do.

Suppose a small child comes to a candy store, purchases two lollipops for ten cents apiece, and gives the cashier a quarter. What ought the cashier to do? The answer may seem obvious. Give the child the nickel change. But why ought the cashier do that? A utilitarian would consider the consequences of the alternative courses of action -- giving the child a nickel, giving the child nothing and saying there is no change, giving the child a penny and saying that is correct, and so on -- and choose whatever alternative produces the greatest good for the greatest number. Considering the matter that way, we might choose to give the child a nickel, deciding that the store's reputation for treating customers honestly would be well served by evidence that it treats honestly even those it could easily cheat. Yet deciding what we ought to do by considering the consequences of our act means using the child to gain a reputation for honesty. But, deontologists say, the cashier ought to give the child a nickel because the nickel belongs to the child. Doing the right thing for the wrong reason is not right. The store may gain a reputation for honest dealing by doing right by the child, but, on this view, the cashier ought to give the child the nickel even if giving the correct change would not produce good results.

In *Doing what the judge orders*, John seems to have thought that if Al's test were negative, he would not have to do anything and that if the test were positive, then even if Jane found out that her confidence had been broken, he would not be held responsible. John thus attempted to use Al's being ordered to be tested by the judge as a way to keep himself from being held responsible. Acting that way is like lying to get someone to do what you want them to do. If you are believed and the person does what otherwise he or she would not have done, you will have manipulated the person. Because the information the person was acting upon is false, the person will not have made an informed decision, but will have the illusion of making such a decision. Yet you will have used the person for your own end. The cashier who gives the child the nickel change so that the store will gain a reputation for honesty is using the child for the store's own end. The child and everyone else may be happy, but the cashier is not giving the child proper change for the right reason. John has

so manipulated things that Al will be tested, without any say, and in doing so, John is using Al's being tested to further his own self-interest. John may have made a decision in which things work out for the best, but how decisions work out is not relevant to whether it is ethical, according to deontology.

The aim of morality, on this view, is not to make people happy, but to accord them the respect and dignity they deserve as persons. Lying to people to get them to do what you want done is not to treat them with respect, but to attempt to manipulate them -- to treat them as instruments for your end and not as persons capable of having ends and capable of deciding for themselves what they ought or ought not to do.

So deontology holds up a very different vision for us from that of utilitarianism. If we try to create the greatest good for the greatest number, and so take consequences into consideration, we treat our acts as means to the end of having a certain kind of world and so, deontologists claim, end up treating people as means to that end -- using them to produce a world in which they may not be among the greatest number who have the greatest good. We ought instead to work for a world in which everyone is treated with respect, and the way to achieve that end is to treat everyone with respect. We ought to act on principle, that is, regardless of the consequences. We may in fact produce more happiness for more people if we have as our end treating persons with respect, but happiness ought not to be our end.

This is a powerful ethical vision. What matters ethically is that we act out of good will, and 'the good will is not good because of what it effects or accomplishes or because of its competence to achieve some intended end; it is good only because...it is good in itself.' So, as opposed to utilitarianism, the consequences do not matter ethically.

But not just every act of will is good. We might will to kick puppies or to teach children to torture animals and other children, and then we would be acting out of an evil will. So what is it to act out of a good will? Kant's answer is somewhat convoluted, but we do not need to get into it in all its complications to explain what it is to act from good will. We do, however, need to make a distinction between acting from inclination and acting from duty.

We are all familiar with acting from inclination. We do it all the time. When we are so inclined, we stay in bed longer than we should, indulge ourselves in a large dessert, and so on. We have inclinations relevant to everything we do, and, more importantly, some of our inclinations, at least, are part of our character. Some of us are inclined to help others, inclined to do good whenever we can, inclined to tell the truth, and so on.

We are also familiar with what it is to act from duty. When we were children, our parents often gave us things to do -- cleaning our rooms, walking the dog -- that we were sometimes not inclined to do, but had to do anyway. Our parents made it our duty. They did not say, 'Keep your room clean if you want,' but 'Keep your room clean.' As adults, we all have duties -- as parents, as citizens, as professionals. A social worker has a duty to inform needy clients what the welfare system provides, for instance. A social worker has a duty to help clients.

A social worker may be the sort of person who wants to help others and so be inclined to help clients. But, Kant says, we deserve no moral credit for doing what we are inclined to do because we are inclined to do it. We deserve moral credit only for doing our duty, only for doing what ethics tells us we must do, whatever our inclination.

It may be that we are inclined to do what duty tells us we ought to do. It may be that we are inclined to do something other than what duty tells us we ought to do. The latter are the dramatic cases we are likely to notice -- the person who wants to kick dogs but, with great force of will, is able to pull back each time. But the person who loves to help others out of an inclination to help deserves no more moral credit, on the view of deontologists, than the person who kicks dogs out of an inclination to do so. Doing good or ill out of incli-

nation has no ethical value. Only acting from duty has ethical value.

So how are we to act from duty? Kant assumes that rational beings like us always ought to have reasons for doing what we choose to do. These reasons can be stated in the form of maxims which provide the reasons for doing whatever it is we do. Someone who eats fruit everyday is committed to some maxim like 'An apple a day keeps the doctor away,' and that maxim provides the reason for the person's eating fruit every day. This maxim is not (in any obvious way) an ethical principle, and in *Dancing the legal dance*, the maxim Mary would have to adopt to justify deceiving Martha is not ethical either. So how are we to tell when we have an ethical principle and when we have something else?

In *Dancing the legal dance*, suppose that when Martha asked Mary not to tell anyone about her father's having sexually abused her, Mary said, 'No, of course I won't' -- while all the time knowing that she would have to inform the authorities. Mary would be committed to an argument roughly like this:

Martha will not tell me what her father is doing unless I promise that I will keep the information confidential.

I cannot help Martha unless she tells me what her father is doing.

But I cannot keep the information confidential.

Yet it is permissible to promise that you will do something you know you will not do when the deception will benefit the person you deceive.

She then concludes, from these premises,

Therefore, I will deceive Martha by promising I will not tell when I know I will.

The crucial ethical premise is the last one: is it permissible to promise to do something you know you will not do on the ground that the false promise will produce a benefit to the person deceived? A rule-utilitarian would consider the consequences and judge according to whether acting on that maxim would produce the greatest good for the greatest number. But the proper test, deontology claims, is whether you can universalize the maxim. Can you consistently maintain that everyone act on that maxim?

The question asks whether, among other things, you would be willing that others make a promise to you they know they will not keep because they think it will be in your benefit not to keep it. If everyone can act on that maxim, others can act on it regarding you. Are we willing to be treated as we are contemplating treating someone else? Mary must ask whether she would also approve if Martha deceived her for what Martha thought was Mary's benefit. Universalizing the maxim forces Mary to be objective, to consider how anyone similarly situated ought to act. It especially forces her to look at how she is treating others because it forces her to see her own action from their point of view.

In this case, we cannot consistently maintain that everyone act on the maxim. It would thus not be rational to adopt that maxim. For to make a promise is to ensure the person you are promising that you intend to do what you have promised. That is the point of promising, and that is why people ask, when in doubt about your intentions, 'Do you promise?' But we cannot be promising while intending not to keep our promise. That is not what it is to promise. We can say, 'I promise...', and people may believe us in part because there is a practice of promising in which those words mean that we intend to do what we say we are going to do. But to say the words, without the intention, is to take advantage of the expectation created by the practice and thus to use the practice, and thus all who engage in it, for your own end. You put the practice at jeopardy and so harm others because if

others acted on your precedent, no one could ever depend on anyone's saying, 'Yes, I really do promise.'

So the maxim we are suggesting Mary may have thought she should act upon is not an ethical principle. She cannot consistently maintain that everyone ought to act upon that maxim. What is ethical is that when we make promises, we intend to keep them. Universalizing a maxim is thus, it is claimed, a test for determining whether a maxim is an ethical principle or not. We can justify a variety of ethical principles by universalization. 'Do not lie' is another ethical principle. So is 'Do not cause unnecessary harm.'

We thus have a vision and a test for ethical principles, and we know what is ethically relevant in a case. It is not what happens, but what is decided and how the decision is made. It is not enough just to intend to do what is right. We might intend to do right for the wrong reason. What is required is that we have the right intention for the right reason. Only then do we decide to do what is right because it is right.

Deontology requires that we think through every ethical decision we make to ensure that we are acting on principle. Some people are naturally empathic, for instance, but if they act on empathy, they ought to get no ethical credit for that. We only deserve credit for choosing to be empathic. What matters ethically is that we choose to do the right thing for the right reason. Central to Kant's theory is thus what he calls a good will. We must will what is good -- will, for instance, that the young child get the nickel change deserved. It is ethically wrong to do anything else.

Making a promise without intending to keep it is ethically wrong even if it should work out that it has good consequences. A wrong has been done regardless and so a person has been harmed, in that the person is wronged, even if the person is benefitted by the consequences of the wrong. So Mary ought to intend to keep any promise she makes because it is ethically wrong to make promises without intending to keep them.

Just as the cashier would be using the little boy if the boy were given the correct change in order to further the store's reputation for honesty, so Mary would be using Martha by lying to her. Martha would be making a decision about what she ought to do based on the false information Mary had given her. Mary would thus be playing with her as though she were an intelligent instrument, knowing that if she made Martha believe certain things, Martha would act in the way Mary wanted her to act. That would not be to treat Martha with the respect due another person.

In this vision, each of us is to be treated by every other one of us with the respect due an autonomous person, capable of self-determination. We would thus have a world in which none use any other for their own ends, but each of us is treated as what Kant calls an end-in-itself. Were this vision to be realized, we would live in a world of mutual respect, in which each of us would respect all others.

This is a powerful vision, resting on the capacity for rational thought, and it includes among those within its scope all those beings capable of rational thought. We are thus to consider ourselves to be aiming towards a moral universe which we share with any other rational being, but not with all sentient beings. That a being is capable of pleasure and pain is not itself sufficient to make it rational. The scope of deontology -- the beings who have primary ethical standing in its universe -- is thus significantly different than the scope of utilitarianism.

Deontology is not without its problems, however. Among other things, it is almost impossible for us to know whether we or anyone else has acted ethically. If you have an inclination to be helpful and you are helpful, how can you be sure that you acted from duty and not from inclination? How can the rest of us be sure? Even when you are inclined to be mean, but instead are nice, how can you be sure that it is not your inclination that is oper-

ating -- an inclination to manipulate, for instance, and trick others into thinking you are nice when you are mean -- and how can the rest of us be sure?

Second, if, in 1.2 Dancing the legal dance, we conclude that Mary ought not to promise Martha anything without intending to keep her promise, we have really not done justice to her ethical problems. She cannot just act on one principle, about keeping promises, but must weigh that principle against the principle that she has an obligation to help keep Martha from being sexually and physically abused. The principle that when we make promises, we intend to keep them is thus at best a prima facie principle -- a principle we ought to keep unless compelling ethical reasons exist not to keep it. It would be wrong for her to make a promise without intending to keep it, but it would also be wrong to allow harm to continue when she could prevent it. It is just her problem that she does not know what her maxim of action ought to be because she has no clear way of choosing between the sets of harms. Deontology would help only if we could universalize one, but not both the contending principles or if we could universalize some compromise between the two competitors. But it seems we can universalize both principles, and it is certainly not obvious what compromise between the two can be universalized. So the theory would seem to give Mary no help.

Third, deontology's vision seems cold. We are to treat everyone as ends-in-themselves, but are we permitted to show more love to one person than to another? To favor one over another by giving that person gifts? How do we avoid making invidious distinctions between individuals while distinguishing between them in ways that reflect our very human capacity for caring and loving? Deontology seems to elevate the rational part of us to the highest ethical level while leaving to one side the sentient and the emotional part that allows us to connect with others as more than simply rational beings.

Fourth, it thus excludes from having an ethical standing those sentient beings who lack the capacity for rational thought. Utilitarianism and deontology stand in stark contrast on this issue.

So, as with utilitarianism, deontology has its problems, that is, but it does present a very different vision of the world. We can summarize it as follows:

- It tells us that we should hold before us as a vision to create a world in which each and every one of us is treated with respect and dignity.
- It bases its vision on rational beings having the capacity to reason about why they do what they do.
- What it picks out of a case as being important and ethically relevant are what those involved have decided to do and what reasons they have for their decisions.
- It gives an ethical hearing to all rational beings and only rational beings. Those other beings who feel pleasure and pain but are not rational cannot be ethical agents and are not to be accorded the respect and dignity we must accord rational beings.
- It justifies a decision about an act or omission in terms of whether it is a principled decision, one based on a maxim which can be universalized.

(c) Virtue theories

It is not enough just to decide to do what is right, for the right reason. What happens matters, and being ethical means doing the right thing in the right way. Yet the ways in which one may fail to do something right in the right way are many. Children asked to apologize to a sibling sometimes yell in anger, 'I'm sorry!' They have said what they ought to say, but not apologized. Just saying the words is not enough. We must feel apologetic and display that feeling with the appropriate behavior. We can fail to apologize by doing it too late, by using the wrong tone of voice or the wrong facial expression, and so on. What virtue-theory tells us is that we ought to act virtuously -- do the right thing in the right way -- and it emphasizes that persons have capacities for kindness, generosity, empathy, and so on and that they can develop and hone these capacities to become ethically better persons. Just as we are born infants and, in growing to adulthood, can affect how we develop physically through exercise, eating well, and so on, so we are born with various capacities that will develop or not come what may. Virtue theory tells us that we ought to develop those ethical capacities we have to become the very best persons we can be, ethically.

Whereas deontology and utilitarianism concentrate upon particular acts and principles we follow, virtue theory concentrates upon our character. A person can do what is right without being a good person, and neither deontology nor utilitarianism would allow us to distinguish a bad person who did what was right from a good person who did what was right. What they measure is whether what was done was right. Virtue theory tells us to look not only at our acts and at the principles which animate our behavior, but at our underlying characters -- the kinds of persons we are. And it tells us that we should strive to become the very best we can be.

Virtue-theory thus provides us with an ethical vision. Because we are born essentially undeveloped and must develop somehow, and because we have the capacity to develop ourselves to become far better persons than many of us do become, we ought to hold up before ourselves the vision of a world in which persons are as fully developed ethically as they can be. We can choose what sort of persons we ought to be in this world, and we ought to become as kind, as generous, as empathic, and, in general, as virtuous as persons are capable of becoming. Morality is best conceived as exercising a skill, something we can learn to do better and better.

This is a vision that requires us to take into consideration both the reasons for and the consequences of our actions. Just as we can always make mistakes, no matter how good we are at something, we can always make an ethical mistake, no matter how virtuous we may be. The goal, however, is to strive to become as virtuous as we can be -- minimizing the ethical mistakes we make and learning how to handle ethically the consequences of the ethical mistakes we do make. The test for our successes and failures turns on both how we act and why we act as we do. To be truly concerned for Martha's well-being in Dancing the legal dance, Mary has to feel concerned and act accordingly. And she has to take into consideration how what she does will affect Martha -- how deceiving her will mean that Martha will no longer be able to tell, by looking at Mary's face and listening to the tone of voice, whether to believe her or not.

Virtue theory thus tells us what is relevant in any ethical case. Ethical perfection is achieved when someone is so attuned to the specifics of a situation that he or she does exactly what is right, both in the immediacy of the situation and for the long-term. Virtue-theory thus gives us as well an ethical vision and a way of testing for what fits that vision.

Some recent work in ethical theory has reinvigorated virtue theory, which long lay dormant, relatively ignored by philosophers. Utilitarianism and deontology are both rule-

driven decision procedures. They are decision-procedures because they tell us how to go about making decisions about what to do ethically. Utilitarianism, for instance, tells us always to look to see whether our act, or the rule we follow, will or will not produce the greatest happiness for the greatest number. 'Do that which provides the greatest happiness for the greatest number!' is a rule, and it tells us to perform those acts or omissions which satisfy it. It thus serves as a guide for telling us what we ought to do. But ethics consists of more than just following rules. In our five-step method of tracking harms, the last step tells us:

- (5) Determine how to do what you have determined ought to be done and do it in a way that will itself produce more good than harm.

The point of this last step is to remind us that making a decision about what to do is not the end of what we have to do ethically. We need to be the kind of persons who, deciding to do the right thing, do it in the right way.

That emphasis on how we are to relate to one another is the emphasis that has been picked up by recent work about what has been called the ethics of caring. This work tells us that the end of ethics is not to follow rules, but to be a caring person. It thus tells us that what matters ethically is what kind of person we are, and it gives us guidance as to which of our characteristics we ought to develop.

But virtue theory, even in its most recent form, has its problems. First, we need an ethical basis for determining what characteristics we ought to cultivate as virtues. A warlike society may think killing without regret worth cultivating, and, indeed, different professions in our society value different characteristics. A lawyer must be as tough with a client as any opposing lawyer would be so as to ferret out any weaknesses in the client's story, the better to provide a defense. A social work practitioner who treated a client that way would be acting inappropriately. Compassion is a virtue for social workers, not lawyers. But it would appear that no version of any virtue theory can tell us why we should develop the characteristics it takes to be of ethical import. If we are to be caring, for instance, what ethical basis can there be for that judgment that does not appeal to the consequences of caring, for instance, and thus to utilitarianism?

Second, even if we could agree on what virtues social workers ought to encourage, for example, virtue theory fails to give us guidance. For instance, it is a virtue of social work that practitioners should listen to their clients, try to understand how they are perceiving their problems. But one practitioner's 'listening' -- by asking questions, querying responses, testing for consistency -- may be another practitioner's understanding of what it is to interrupt. It is not obvious that any one way of listening is ethically preferable to another, and virtue theory gives us no guidance about how to choose between alternative practices of such a virtue.

Third, virtue theory has a more important failing. It gives us no guidance about how to proceed in ethical dilemmas and problematic cases. To be virtuous we must act in virtuous ways, but the problem these kinds of difficult cases presents is that we do not know what it is we ought to do. Whether Dena tells the physician what she knows or decides to withhold that information, she ought to be so attuned to what she is doing that she does it well -- in the right way, with the right tone of voice, at the right time, and so on. But virtue theory gives her no guidance as to which choice is the ethical choice.

We can summarize virtue theory as follows:

- It tells us that we should hold before us as a vision to create a world in which each

and every one of us has honed our capacities to become the most virtuous individuals we can be.

- It tells us that what matters ethically is what kind of persons we are, not just whether we act rationally or whether our acts produce the greatest happiness for the greatest number.
- It bases its vision on human beings having various capacities -- to be kind, to be generous -- that will develop in some way regardless of whether we work to develop them to make ourselves ethically better.
- What it picks out of a case as being important and ethically relevant are what those involved do and feel.
- It gives an ethical hearing to all human beings.
- It justifies a decision about an act or omission in terms of whether it is done in accordance with virtue -- whether someone telling the truth honestly feels honest as well as tells the truth, whether a person being apologetic is really apologetic as well as apologizes, and so on.

(d) Comparing ethical visions

Each of these three ethical theories appeals to some essential feature we have and erects a vision for us to achieve, and yet these are visions of distinct ethical worlds with radically different conceptions of what is ethically of value. Utilitarianism emphasizes that we are creatures who desire happiness and could have more for more of us. We should thus strive, utilitarianism tells us, to create a world in which the greatest number have the greatest amount of happiness. Virtue theory tells us we have capacities of empathy and kindness, courage and caring, and so on which we can each hone to their fullest expression. Deontology tells us we are rational beings capable of conceiving of and having a world in which everyone is treated with respect and dignity.

If you had to choose between these visions, how could you choose between a world in which there was the greatest happiness for the greatest number, a world in which each of us developed and honed our capacities for such traits as kindness and caring to their fullest, or a world in which we were all accorded dignity and respect without regard to our color, our race, our sex or any other natural or social characteristic we might have? Each vision is enticing. Yet they compete. Most obviously, what provides the greatest happiness for the greatest number may not only preclude some -- the smaller number -- from having even a modicum of happiness, but also the sort of respect that deontology emphasizes. It is one worry about utilitarianism that because concerns about the greatest number outweigh concerns about minorities, minorities, already heavily burdened by a disproportionate share of problems, will find their problems further aggravated by appeals to utilitarianism, losing any claim to dignity along with a chance for happiness.

We can get a sense of how different are the visions these theories hold up for us by asking what we ought to do to achieve them. Virtue theory tells us to work on ourselves, make ourselves better people, letting our relations with others and the structure of our society come to whatever they will given the changes in us. Utilitarianism requires us to work

on the social structures that preclude us from achieving a world in which the greatest number are the happiest. And deontology requires us to ensure that each of us as individuals respect all others as equals and that the social structures do that as well.

That these ethical visions are so different might not be thought to matter if achieving any one of them led to the same practical results. A utilitarian, for instance, might argue that the only way to achieve the greatest happiness for the greatest number is to accord each and every individual dignity and respect. But consider the situation from your point of view if you were denied dignity -- by being enslaved, for instance. To claim your freedom, you would have to show that your being a slave in fact did not provide the greatest happiness for the greatest number and that some other distribution of freedom and slavery did. This is an empirical claim and hard to prove, especially for someone not well-positioned to prove it because enslaved. So you would have a difficult time obtaining your freedom. But, on deontological grounds, you deserve dignity and thus freedom as a person -- quite independently of whatever effects the denial has on others.

It matters, practically and ethically, what sort of justification we must give for what we do. Both deontology and utilitarianism may prohibit slavery, but if the reasons for the prohibition differ in the way suggested, the identity of result is not as important ethically as that it is more difficult, under utilitarianism, to justify eradicating slavery. What distinguishes utilitarianism, deontology, and virtue theory is that they have differing visions of the ethical life and so differing conceptions of how we ought to justify what we ought to do. A deontologist appeals essentially to the rights of individuals, and if the right of an individual to liberty should happen to cause less happiness for the greatest number, so much the worse for the happiness of the greatest number. For a deontologist, the appeal to rights trumps any appeal to the greatest happiness of the greatest number. For a utilitarianism, the opposite is true. We have no rights but what will further the greatest happiness for the greatest number, and should someone claim a right that does not further that end, the claim is denied. The end of the greatest happiness for the greatest number trumps any other vision of the world. So, for utilitarianism, we ought to encourage only those virtues that further that end. Any other virtue can have only a negative ethical import. Virtue theory, on the other hand, seems not to accept that limitation. What matters is that we strive to become as virtuous as we can be, and we are not to limit ourselves to whatever might happen to provide the greatest happiness for the greatest number.

It thus matters which theory one chooses. But how are we to choose between these ethical visions? The three we have considered are the main contenders, and yet the history of philosophical thought is filled with critiques of them. No criticisms of any one theory have been so devastating as to make it obvious that it is ethically bankrupt, and yet no overarching ethical theory allows us to choose between these competitors. We thus have no ethical basis for choosing between them.

Yet the theories are ethically different, giving us different answers to the ethical problems we face and opening up ethical situations in different ways, allowing us to see different ethical features. If we do not familiarize ourselves with these different theories, we may find ourselves looking at ethical issues in social work with one vision in mind, not realizing that there are competing ways of perceiving the same situation. That is a recipe for ethical confusion and mistakes. Two or more people may consider the same ethical issue with different visions in mind, without realizing it, and so come to different conclusions about what they ought to do. Or they may consider the issue with the same vision in mind, not realizing that the situation cannot be fully understood unless seen in terms of another ethical vision.

One way to understand the problems with John's decision to have AI tested in Doing

what the judge orders is to ask whether that decision respects Al's right to determine for himself whether he wants to be tested. If you look at the issue as a virtue-theorist, you will concentrate upon how skillful or unskillful John was in exercising the capacities to empathize with Al's potential problems or with the problems of Al's mother. And so you may miss the issue of Al's self-determination. If you look at the issue as a utilitarian, concerned only to work things out so that most people involved in the case are as happy as they can be, you may miss seeing the case from Al's point of view. Looking at the case as a deontologist would, and so asking, 'Are everyone's rights being respected here?', forces into the open, for examination, the question whether Al should have the right to determine for himself whether or not to have a test.

Even if all theories lead to the same answer, they provide different reasons, and these differing reasons can make an enormous difference in how we think about what we ought to do to prevent recurrences, for instance. Consider lying. It is generally wrong for a utilitarian because it leads to more unhappiness for more people than telling the truth. It is wrong for a deontologist because it is a form of manipulation, treating the person as an object rather than a person. It is wrong for a virtue theorist because it runs counter to the virtue of honesty. So all three theories may agree, in a particular case, that we ought not lie. But because the reasons they provide for not lying differ, they direct our concerns in different directions. Virtue theory tells us to work on our capacity for telling the truth. Deontology tells us to treat others as we should treat all persons -- and as we would like to be treated ourselves. And utilitarianism directs us to change the system in ways that ensure that lying does not further the happiness for the greatest number.

As the fourth step in our method of tracking harms, we need to:

- (4) Back off from the case and judge what is best to do: what will minimize harms?

We can now see that this question may not produce certainty about what we ought to do. All the theories agree that we ought to act to minimize harms, and they all agree that in making such judgments we need to be sure that each person counts the same as every other person, with no one person's interests being elevated above any other's. But they provide competing interpretations of harm -- a harm for a utilitarian being what impedes or at least does not further the creation of the greatest happiness for the greatest number, a harm for a deontologist being a failure to act from duty (and so, for instance, the manipulation of someone by lying), a harm for a virtue theorist being what slows or arrests or reverses the development of an ethical virtue. With competing conceptions of harm, the fourth step in our method becomes itself conceptually problematic. We may thus end up with a case where the competing theories tell us to do different things and have no way of choosing, ethically, between the different answers.

iii. True ethical premises

Even though the different theories provide different reasons for not lying, they all agree that it is still generally wrong to lie, whatever the reasons, and as we have seen, we can do a great deal without waiting for philosophers to settle their disputes about which of these ethical theories is ethically best. We have illustrated how much can be done as we have worked through the cases in this chapter.

In regard to *Doing what the judge orders*, we explored whether John had gained anything by having Al tested. He began the case with the ethical problem of whether he

should tell whether Jane had AIDS. He decided to have Al tested, and the consequence was that he caused harm without gaining any leverage on the original problem he had. In giving that analysis, we assumed it ethically wrong to cause harm for no good reason.

In analyzing *Dancing a legal dance*, we came to see that in passing on to the authorities the information Martha gave her, Mary caused harm, but that she had an obligation to pass on the information. We were unsure what Mary had told Martha, if anything, about what she would have to say to the authorities if she were told about abuse, but in considering what Mary ought to have told Martha, we were presuming that Mary ought to minimize harms and make sure that Martha understood that any information she gave Mary about abuse would have to be reported to the authorities.

In several cases we were able to come to conclusions about what ought to be done, and we came to these conclusions in ways we think are non-contentious. Who would think that Dena should treat people who are adopted and want information about their parents differently when they make appointments ahead of time with her and when they do not? We also appealed to non-contentious ethical premises in analyzing these cases -- that it is ethically wrong to cause harm for no good reason, that we ought to act to minimize harms (when we can), and so on. We will find other such premises as we work through the cases we shall consider in later chapters.

Are such ethical premises true? They are certainly not contentious. No social work practitioner would claim, we can safely assume, that we ought to cause harm for no good reason or that we ought to maximize harms when we can. And each premise is supported by the ethical theories we have examined. A utilitarian would argue that if we are to achieve the greatest happiness for the greatest number, we ought not to cause harm -- especially for no good reason. A deontologist would ask whether we could universalize the maxim of causing harm for no good reason and find that we could not. We cannot consistently maintain that everyone ought to act on such a maxim, for, among other things, we would be agreeing that others can harm us for no good reason. And a virtue-theorist could not agree that it is a virtue, something we should become skilled at, to cause harm for no good reason.

Some of the premises to which we have appealed are ethical principles, and it may be tempting to think one only need to appeal to them to resolve ethical issues. We can readily draw out of the cases we have discussed so far at least the following principles:

- the principle of self-determination (or autonomy), that a person has the right to decide what to do or omit from doing regarding those matters that affect him or her and so has the right to be fully informed about such matters
- the principle of nonmaleficence, that one ought not to harm others
- the principle of beneficence, that one ought to help others, either by mitigating harmful conditions or by providing benefits
- the formal principle of justice, that we treat like cases alike and unlike cases unlike

These principles can serve as guideposts for us. First, they give us guidance about the kinds of considerations we need to keep in mind when assessing a case. They remind us of the interests of the various parties involved that we need to consider -- an interest we all have in self-determination, an interest in not being harmed, and so on. Second, that we have these interests enshrined in principles reminds us that we shall need very compelling ethical rea-

sons to do something that causes harm to others or denies someone's self-determination, for example. That is the point of these being principles. There is a presumption against breaking them, and the presumption can only be overridden if there are compelling ethical reasons to do so. We might deny someone's self-determination, for instance, if permitting it would cause great harm to others.

But it would be a mistake to think that all we need to do is to appeal to these and other principles to determine what we ought to do. It would be a mistake for three different reasons.

First, consider *Dancing the legal dance* again. In deceiving Martha, if that is what Mary did, Mary lied to her and denied her the information she needed to make a decision for herself about what she ought to do. Mary thus broke the ethical principle that we ought not to lie and the principle that we ought to make every effort to foster self-determination on the part of clients. She may have thought she was obtaining information she could use to help Martha, and so she may have been acting on the principle of beneficence. If we conceive of her problem as a conflict of principles, we must weigh the principle of beneficence against the principles of not lying and fostering self-determination, and it is hard to know where to begin. Is fostering self-determination more important than helping someone? If so, when is it more important, and is *Dancing the legal dance* one of those cases in which it is more important or one in which it is not? Does it matter that Martha is so young? In short, we seem not to find a ready resolution of problems when we take them to be conflicts between principles.

Second, even when the principles themselves seem clear, we may disagree on their application. Even an appeal to the formal principle of justice can be contentious. One concern we had about Dena in *Adoptive Children* is that she seemed to be treating the young brother and sister differently from the way she treated others because they were not thoughtful enough to call her to make an appointment. She thinks that their not being thoughtful is a relevant difference between them and others in need. We might disagree. We could agree that we should treat them differently if they were relevantly different in their concern for the interests of others, but not think they are really relevantly different. Just appealing to the principle of formal justice will not help.

Third, the ethical principles are contentious because their meaning can vary from theory to theory. Two people might appear to be using the same ethical principle when they have very different sorts of harms in mind. In arranging to have Al tested for HIV, John denied Al his autonomy. Where a deontologist sees a denial of respect that can only be set right by according respect, an utilitarian sees a harm that could be set right if things worked out for the best. They would both appear to be using the same principle, of not causing unnecessary harm, but by concentrating upon it, we obscure the differing kinds of harms that really are at issue in a case. Different ethical theories create different contexts for cases, that is, and approaching a case with principles rather than theories in mind may blind us to different ways of seeing a case.

So we recommend that we range over the ethical theories to understand a case fully, holding a case up first to one theory and then to another to see it in all its aspects. This has the advantage of forcing us, if we are to do a full survey of a case, to lay it out in all the ways conceivable by the competing theories. Just as we can bias our judgment by some feature of ourselves we bring to a case unawares, so we may bias our understanding of the problems in a case by a particular conception of what counts as an ethically relevant feature.

We do not suppose that holding a case up to survey it in this way will allow us completely to overcome either the personal or social values we bring to the case, but it will help

as will consulting with colleagues, a supervisor, and/or an ethics committee -- as long as confidentiality is not broken in the consultations. Bringing another professional's point of view to a difficult case is to bring fresh eyes to the case. Getting their views on the case -- both on what makes it difficult and on what you might do to resolve it -- may help open up the case up so that you can see what it was that made it difficult and help ensure greater objectivity, less reliance, that is, on any personal or social values you may be bringing to the case.

For example, in *Adoptive Children*, someone may be taken by Dena's apparent inability to relate well to the brother and sister who came to her. They irritated her, but focusing on that aspect of the case is to concentrate on a feature of the case virtue theory makes prominent, namely, how properly to relate to others. We may then ignore the issue of whether her criteria for breaking the law are adequate or unfairly discriminate -- issues we are more likely to notice if we approach the case as a deontologist or a utilitarian. If we approach the case as a utilitarian, we will concentrate upon the consequences of what she does and so get a better view of the implications of her behavior for her and for others. If we approach the case as a deontologist, we will concentrate upon whether she lies to those who seek help from her when she does not want to give it. To be sure that we have thoroughly understood the ethical issues in a case, we ought to look at the case from all those perspectives -- examining the intentions of those acting, the consequences of the actions, and the virtues displayed in the action.

The procedure we are recommending, as we have seen, can lead to clarity about what is at issue and about what we ought to do. We think that most cases can be resolved by getting clarity on what went wrong and what goals the social worker in the case ought to have. But there is no guarantee that we shall always be able to resolve difficult ethical problems. Among other problems, we cannot always be sure that we have fully analyzed a case and seen all the ethical issues it raises. In addition, we may find that even when we think we have, we cannot resolve the case without making some choice among the competing ethical theories -- a choice, as we have claimed, for which you can have no ethical basis. But as with any complex problem, our failure to resolve it does not mean that it is not resolvable by others or by ourselves with more information or a fresh look at it or a different way of thinking about it. We should always presume that cases are resolvable, and when we find ourselves stuck, we should back off and find out new ways of proceeding to untangle the case. We might consult with colleagues, for instance, if that will not breach confidentiality. Only after all available resources have been exhausted are we permitted to say, not that the case cannot be resolved, but that we cannot now resolve it.

Questions

- (1) Provide an example from your own experience of where you would have engaged in unethical behavior if you had acted in your self-interest.
- (2) Provide an example where someone harmed their self-interest in doing what was ethical.
- (3) Give two examples of when you have acted prudently in your job or in an internship. If you can, give two examples of when you acted imprudently in the same sort of situation or give two examples of when you saw others act in a way you thought imprudent in their job or internship.
- (4) Examine the acts of prudence you provided for the first question and determine why

- they are acts of prudence. Are they also ethical? If not, why not? If so, why?
Distinguish the reasons for which they are ethical (or unethical) from the reasons for their being prudent.
- (5) Look back to 1.3 *Adoptive Children* and sort out among all the things she is doing what is prudent and what is not. You should do this both from her point of view -- what must she think is prudent and what not? -- and from your point of view -- what of the various things she is doing do you think prudent and what not?
 - (6) Look back at 1.1 *Death of a baby*. After sorting out what the law required Deborah to do and not to do and what she actually did, determine in what ways she acted illegally. Then determine whether she was prudent in breaking the law and whether, given her decision to break the law, she acted prudently to protect herself and/or to protect others.
 - (7) Provide an example, preferably from your professional practice, of where, unsure what to do ethically, you did what was legally required as a way to resolve the issue you faced.
 - (8) Lawyers often put their clients through grueling questions in order to ensure that they know exactly where the weak spots in their clients' cases are so that the lawyers can better defend them. Nothing in their code of ethics specifies the manner in which lawyers are to do this and the standard practice varies: some are nice about it, some are not. Provide a similar sort of example about variations in standard practice from social work.
 - (9) Provide an example, if you can, of someone stepping outside the standard practice. What was the reason for doing this? Did the person not know what the standard practice was? Or did the person have an ethical reason for this?
 - (10) Does social work need a code of ethics? Why, or why not?
 - (11) Explain why appealing to a code of ethics does not resolve our problem when we face an ethical dilemma. You may use examples cited in the text.
 - (12) Explain why appealing to a code of ethics does not resolve our problem when we face a problematic case. Again, you may use examples cited in the text.
 - (13) Provide an example, out of your own experience if you can, of how something looked to be one thing but was another -- like the example of the woman who killed her boyfriend. What differences did looking at it differently make?
 - (14) In 1.2 *Dancing a legal dance*, Mary thought that Martha was being defiant in running away from home and staying with friends. She made an inference about why Martha was doing what she was doing based on the fact of Martha's staying with friends. We often make such inferences -- that the mail carrier has come because we hear our mailbox open and close, that there is orange juice in the refrigerator because we left some the last time we had it, and so on. But these are inferences from facts, and though they may be plausible, they may not be true. Give five examples of plausible inferences from facts where the inferences are in fact false.
 - (15) Look back at one of the cases we have examined or to one of your own and provide an example of something in that case that turned out to be relevant to its resolution that you did not pick up in your first reading of it.
 - (16) In 1.3 *Adoptive Children*, what would Dena have had to do to make an objective judgment?
 - (17) Give an example of how you had to discount your biases to render an ethical judgment about a client or a fellow social worker.
 - (18) If it were true horses have wings and winged creatures fly, what conclusion should we draw? Is the conclusion true? Are the premises true? Can we draw the conclusion

- even though the premises are false? Ought we to believe the conclusion? If so, why? If not, why not?
- (19) If we know that someone concluded that they had seen a male goldfinch, and when asked they said they saw a small bright yellow bird with black wings, what premise are they assuming?
 - (20) Give an example of a valid argument.
 - (21) What are the two ways in which we can assess an argument?
 - (22) Why do we need to provide reasons for ethical judgments?
 - (23) Give an example of tracking harms in a difficult case in your practice with clients or relations with colleagues. How did you decide which course of action to take?
 - (24) Someone can harm me and my interests by breaking a pencil I am about to use. They can also harm me by shooting me. List five different harms for each situation and rank them by the degree of severity. Make sure that you have listed harms that differ sufficiently in severity that no disagreement seems possible about the ranking you propose.
 - (25) Pick one of the cases we have examined or one of your own and lay out the ethical principles involved.
 - (26) Explain why it is a mistake to rely solely on ethical principles to resolve ethical problems.
 - (27) Pick one of the cases or a case of your own and examine how principles serve as signposts.
 - (28) What are the competing visions of deontological theory, utilitarian theory and virtue theory? That is, what kind of life does each think we ought to lead?
 - (29) What tests do the competing theories provide to determine what, according to them, is the right thing to do?
 - (30) Suppose that a colleague of yours lies to you regarding something about your job. The colleague thinks that if you knew the truth, you would be upset and thus harmed. Explain how a utilitarian would look at the situation and compare that to how a deontologist and to how a virtue theorist looks at it.
 - (31) In regard to your answer to (30), explain how a deontologist thinks that one problem with utilitarianism is that it condones using people for ends they may not share.
 - (32) If you were a slave, would you prefer to live in a society of utilitarians or deontologists? Why?