

[[This paper appears in *Oxford Studies in Normative Ethics 1* (2011): 51-74.

Please cite the published version.]]

Treating Consenting Adults Merely as Means

According to the Mere Means Principle, it is wrong to treat others merely as means.¹ One specification of this principle, suggested by Onora O'Neill as an interpretation of Kant, relies on a notion of possible consent. It holds that an agent uses another merely as a means if she has rendered him unable to consent to her use of him (O'Neill 1989). A person cannot consent to being used if he has no opportunity to avert the use by withholding his agreement to it (110). Suppose that a mugger approaches you on a dark street, points a gun at your torso, and tells you with chilling candor that unless you give him your money, he will kill you. He leaves you no opportunity to avert his use of you by withholding your consent to it. So in O'Neill's sense, you cannot consent to his action. The mugger is treating you merely as a means.

This specification of the Mere Means Principle is vulnerable to serious objections (Kerstein 2009, Parfit 2008). In some cases it is morally permissible for a person to coerce or deceive another in the course of using him, or so many of us believe. But the account implies that in all such cases the person would be treating the other merely as a means. Here is an example of what many of us take to be permissible deception. In order to make your spouse's birthday party a surprise for her, you need to lie to her sister about your whereabouts during a certain afternoon. You use the sister to quell your spouse's suspicions regarding your

plans. As you realize, if you told the sister about the party, she would be unable to keep the secret from your spouse. According to the account we are considering, you treat the sister merely as a means. For your deception leaves her with no opportunity to avert your use of her by withholding her consent to it.

I have recently tried to show that such difficulties can be overcome. A plausible sufficient condition for treating others merely as means can incorporate a notion of possible consent. But it must also include other elements, for example, a notion of possible end-sharing (Kerstein 2009). The sister can presumably share the end you are pursuing in lying to her, namely that of quelling your spouse's suspicions regarding a surprise party. On the account I develop, that she can share this end prevents us from having to embrace the implausible conclusion that you are treating her merely as a means. However, my account is considerably more complex than the one O'Neill suggests.

One might wonder whether there is a simpler route to a plausible construal of the Mere Means Principle.² Perhaps instead of looking to Kant for inspiration regarding a notion of consent, we should look to Robert Nozick (1974, 31).³ He seems to advocate reliance on a notion of *actual* consent. One might, for example, propose the Actual Consent Account: when an agent uses another, she does *not* use him merely as a means if he has given his voluntary, informed consent to her using him.⁴ Of course, this account does not offer a complete understanding of the Mere Means Principle. But we would presumably take a step toward understanding when an agent uses another merely as a means by developing an account of when he uses another, but not merely as a means.

This paper explores the plausibility of the Actual Consent Account. It begins by sketching the account in more detail, setting forth briefly notions of *voluntary*

and *informed* consent. The paper then considers difficulties, suggested by O'Neill, that such an account faces. They turn out to be less serious than it might appear. Moreover, analogous difficulties apply to possible consent accounts of conditions under which an agent's use of another does not amount to her treating him merely as a means.

The remainder of the paper tries to establish that there are two further problems with the Actual Consent Account. One problem, from which possible consent accounts also suffer, is that of implying, implausibly, that an agent does not treat another merely as a means when she profits from vulnerability in another that she (the agent) has induced. Another shortcoming of the Actual Consent Account is not shared by possible consent accounts. The account has the unwelcome implication that certain ineffective or unnecessary attempts an agent makes at coercing or deceiving another to serve as a means to her ends do not amount to her just using the other. The Actual Consent Account thus fails to realize its promise of giving us a simple yet plausible way to capture jointly sufficient conditions for an agent's using another, but not merely as a means. The account can be altered so that it no longer has the unwelcome implication in question, but a Modified Actual Consent Account differs little, with respect to both complexity and content, from possible consent accounts.

Of course, accounts of treating others merely as means should, to a significant extent, square with the judgment of reflective common sense regarding circumstances in which an agent does (or does not) do so. But it would be unreasonable to expect them to accord entirely with ordinary thinking. For one thing, it is not clear that there is any one notion for them to accord with: there might not be a *univocal* ordinary concept of just using another. Accounts of

treating others merely as means should also minimize appeal to moral concepts. In particular, they should avoid relying on moral concepts that are just as elusive as our notion(s) of treating others merely as means. If they rely on such concepts, they are unlikely to be illuminating.

1: Voluntary, Informed Consent

According to the Actual Consent Account, when an agent uses another, she does *not* use the other merely as a means if the other has given his voluntary, informed consent to her using him. An agent uses another just in case she intentionally does something to the other in order to realize her end, and she intends the presence or participation of (some aspect of) the other to contribute to the end's realization.⁵ For example, a tourist uses a police officer when she asks him for directions to get to the train station. But a policeman in pursuit of a suspect does not use a bystander he pushes out of his way; for he does not intend any aspect of the bystander to contribute to his reaching his goal.

Treating another as a means is often morally permissible. Students treat professors as means to knowledge; professors use students to test their theories, and so forth. Granted, in saying "John used me," a person might mean that John treated her *merely* as a means and thus wrongly. But in our terminology, saying that someone used another does not itself imply any negative moral evaluation of his action.

For the sake of simplicity, let us assume throughout that the Actual Consent Account is meant to apply only to competent adults. An agent is competent only if he has the capacity to decide, based on his own values as well as information available to him, whether to use others and whether to consent to their using him.⁶

Of course, we need to get a reasonably clear idea of what *voluntary, informed* consent amounts to in the Actual Consent Account. Notions of such consent have a central role in discussions of the ethics of medical treatment and research, as well as in rules governing these enterprises (Brock 1993). But since the Actual Consent Account is supposed to apply to all cases of one agent (i.e., competent adult) using another, the notion of voluntary, informed consent we employ here needs to apply to all such cases.

A person's consent or agreement to an agent's using him is informed if he understands what the agent intends to do to him and to what purpose, it seems reasonable to say. For example, a person's consent to serve as a means for an agent to get over a fence might be informed as a result of her telling the person that she would like a lift over it so she can get the ball that fell on the other side. But *how much* does a person need to know about the ends an agent plans to pursue in using him in order to give his informed consent to this use? Does he need to have knowledge only of the agent's proximate end or of her ultimate end as well? For example, does the person give his informed consent to an agent's using him to get over the fence if she informs him that she aims to retrieve the ball, but not also that she aims to retrieve it in order to get revenge on her neighbor by throwing it through his window? I set aside this issue here, since I discuss it in detail in the next section.

A defender of the Actual Consent Account would presumably contend that a person's consent to another's use of him fails to be voluntary if it results from the agent's coercing her. An agent's coercing another involves his threatening her, that is, claiming to her that unless she does something, he will make her worse off than she would have been without his involvement, the defender might say. A

victim does not give her voluntary consent to a mugger's demand for her money when he threatens to injure her unless she gives it up. But the voluntariness of a person's consent to an agent's using him is not vitiated simply by its being the case that the person desperately wants or needs the agent to do something for him (Wertheimer 1996, 110) and will be worse off if the agent does not do it, the defender might conclude. For example, a motorist's being stranded on a sweltering desert highway does not render it impossible for him to agree voluntarily to a tow truck driver's picking him up and thereby using him to make a profit. Or a heart attack victim's being in dire need of medical help does not entail that he cannot give his voluntary consent to a surgeon's performing a life-saving operation for his usual fee. Of course, there are other ways in which consent might fail to be voluntary, the defender might conclude. Illness, injury, or brainwashing might prevent someone from doing anything voluntarily, including consenting to be used by another.

A defender of the Actual Consent Account might face the following objection to this view of voluntary consent. In the mugging case, the victim has a choice between losing his life and giving up some of his money. In the surgery case, the victim also faces a choice between these two outcomes. What basis do we have for agreeing with the defender's conclusion that what the victim does in the former case is involuntary but in the latter voluntary? The conclusion seems arbitrary.

Of course, if the defender altered his view and held that the consent of both the heart attack and mugging victims to be used for another's financial gain was voluntary, then he would no longer be vulnerable to the charge that he was making an arbitrary distinction between the two cases. But the defender would presumably not wish to do this. For, assuming that the mugger disclosed the nature of the use

he was going to make of the victim (e.g., taking his wallet to get some cash), an implication of the Actual Consent Account would then be that he was *not* treating the victim merely as a means. But that implication is simply not plausible. It is central to a common notion of treating others merely as means that the mugger is treating his victim in this way.

The defender might try to bolster his account against a charge of arbitrariness by maintaining that the action of neither victim was voluntary. But this response would also have problematic implications. First, the Actual Consent Account's force would be diminished. It would not have what many of us take to be a plausible implication, namely, that the surgeon was not treating the heart attack victim merely as a means. Second, suppose that the heart attack victim refuses the surgeon's offer of life-saving treatment for a fee. How could the defender avoid the conclusion that his *dissent* fails to be voluntary? It seems odd to insist that a victim cannot voluntarily consent to a procedure that would save his life but can voluntarily consent to forgo such a procedure. Perhaps the defender would say that he can voluntarily consent to *neither*. But then, contrary to an entrenched view, a doctor would never derive any justification for refraining from giving a patient (e.g., someone with terminal cancer and at most weeks to live) a life-saving procedure (e.g., yet another painful operation) from the patient's voluntarily dissenting from receiving it. Finally, the defender would be forced to answer a difficult question: when *does* one's allowing an agent to use him count as voluntary? Assuming that one believes that some actions are voluntary, it would be wildly implausible to insist that one never voluntarily consents to another's using him. Is one able to voluntarily consent only when failure to agree to be used would not be too costly in terms of one's well-being? What, then, is too costly? Is a person

able to give his voluntary consent to pay a surgeon to save his leg, but not to save his life?

Rather than embracing either the unappealing conclusion that both of our victims give their voluntary consent to being used or the equally uninviting one that neither do, the defender of the Actual Consent Account should, I believe, insist that it is not arbitrary to say that the one (the heart attack victim) gives his voluntary consent, while the other (the mugging victim) does not. The cases differ from one another with respect to the victims' bases for consenting to be used. In one, the victim consents as a result of a mugger's claim that unless she gives him her money, he will make her worse off than she would have been had she not interacted with him. In the other case, the victim's consent is based on no such threat. The voluntariness of a person's consent is not lost just because he is, like the heart attack victim, in a bad way—assuming, of course, that his condition is not so poor that he has lost his decision-making competence. But the voluntariness of a person's consent to an agent's use of him is lost if it is based on the agent's threat to make him worse-off.

2: Shortcomings of Actual Consent?

Onora O'Neill suggests reasons for rejecting the idea that we should construe the Mere Means Principle in terms of actual consent. First, she implies that doing so would render the principle too unwieldy to use in practice. It is difficult to determine when individuals give their informed consent to being used in a particular way, she asserts (O'Neill 1989, 108). Even if they are told in detail what will be done to them and to what purpose, have they really understood? In some cases it does indeed seem difficult to know whether they have.

But this point is not sufficient to undermine an actual consent construal of the Mere Means Principle. First, that in some cases we are unable to determine with confidence whether an agent has given his informed consent to be used in a certain way is consistent with our being able to do so in the vast majority of others. O'Neill gives us no reason for *general* skepticism regarding our ability to know whether someone has understood a proposal to use him. Second, the same sort of issue O'Neill points to regarding actual consent arises in the context of possible consent. Whether an individual really has an opportunity to avert someone's treatment of him as a means by withholding his agreement to it depends in part on whether he grasps what the proposed treatment is. But it can be difficult to discern whether he grasps this. Does the financial advisor's client understand the myriad ways she plans to use him to make a profit well enough to dissent from them? O'Neill says that patients "cannot easily understand complex medical procedures; yet if they consent only to a simplified account, they may not consent to the treatment proposed" (108). Indeed, it can be challenging to insure that patients give informed consent to complex medical procedures and, of course, challenging to determine whether they have done so. But analogous points apply to possible consent. If a doctor gives a patient only a simplified account of a complex medical procedure, then the patient might not have an opportunity to avert the doctor's performing *that procedure* by withholding his agreement to it. His having such an opportunity depends in part on his understanding what the procedure is. And it can be difficult to ascertain whether he understands this.

O'Neill suggests that there is a "deeper" problem with construing the Mere Means Principle in terms of actual consent. "When we consent to another's proposals," she says, "we consent . . . only to some specific formulation of what the

other has it in mind to do. We may remain ignorant of further, perhaps equally pertinent, accounts of what is proposed, including some to which we would not consent" (108). If we would not consent to these further accounts of some proposed treatment of us as a means, it is unclear whether our consent to the account we are familiar with suffices to make it such that we are not being used merely as a means.

This problem is familiar from our brief discussion above of *how much* a person needs to know about the ends an agent plans to pursue in using him in order to give his informed consent to this use. Imagine a conservationist who is trying to preserve tiger habitat from being taken over by nearby villagers. She enlists one of the village elders to act as a liaison. She would like his help, she tells him, in pinpointing what the villagers really need. He gives her information and she provides well-targeted aid. However, she does not tell him that her ultimate end is to preserve tigers and that she takes as a means to this end making the villagers prosperous enough not to need to appropriate more jungle.⁷ If the conservationist had proposed to the elder to use him to attain this ultimate end, he might not have consented to this use; for he might despise tigers on account of their deadly attacks on his people. If the elder would indeed have refrained from consenting, some of us might, intuitively speaking, conclude that the conservationist treated him merely as a means, even though he consented to be used as a source of information.

A defender of the Actual Consent Account might insist that the elder did not really give his *informed* consent. But then how much information does one need in order to be informed? Does one need to know not only the agent's proximate end,

but his further and ultimate ends as well? The danger here is that we set the bar for being informed so high that one almost never counts as such.

A different defense might be to alter the Actual Consent Account to maintain the following: An agent who uses another does not use her merely as a means if it is reasonable for her to believe that the other gives his voluntary, informed consent to her use of him. It is reasonable for an agent to believe that another gives his informed consent only if it is reasonable for her to believe both that the other is aware of her proximate end and that if he was aware of her further ends, he would still agree to the agent's using him. Here "reasonable" is being employed in a non-moral sense. Whether it is reasonable for an agent to believe something depends on his epistemic situation: his intelligence, upbringing, previous experience, and so forth. On this account, it might not be reasonable for the conservationist to believe that the village elder gives his voluntary, informed consent to her using him for information.

It is not obvious that a defender of the Actual Consent Account would do well to accept this revision; for it threatens to make the account anemic. In many instances, the account would fail to yield the verdict that an agent had not treated another merely as a means, even though many of us would embrace this verdict. Suppose, for example, that a marketing executive is using consumers to further her company's ends. She sends out discount coupons for a product in an effort to get consumers to try it, form a habit of purchasing it, and drive her closest competitor out of business, thereby dramatically increasing her profits. Since her competitor is an old, family run company for which many have warm feelings, it is reasonable for the executive to believe that some consumers would not take advantage of the coupon if they were aware of her end of driving her competition

out of business. Nevertheless, many of us believe that in sending out coupons to get consumers to purchase her product, the executive is not treating them merely as means. But the revised Actual Consent Account would fail to imply that she does not do so. Sometimes an agent's reasonable belief that someone she is using would not agree to be used if he was aware of some of her further ends seems compatible with her nevertheless not using this person merely as a means. But the revised account we are considering obviously fails to register this compatibility.

The revised account was designed to respond to a difficulty posed for the Actual Consent Account by cases like those of window breaking and tiger saving described above. But it is open to proponents of actual consent to respond in the same sort of way that we responded to the last difficulty. The problem applies just as much to possible consent as it does to actual consent accounts. Suppose that an agent has had an opportunity to consent to being used as a prop for someone to get over a fence to retrieve a ball. That does not mean that she has had an opportunity to consent to being used as a prop for someone to retrieve a ball so that she can break the neighbor's window with it. She might not have known that window breaking was in the cards. In that case, when she is used to get over the fence, is she treated merely as a means?

If one has the intuition that the answer is yes, then one might make a modification to a possible consent account analogous to the modification we suggested to the Actual Consent Account. One might say, for example, that an agent who uses another does not use her merely as a means if it is reasonable for her to believe that the other can consent to her use of him. Moreover, it is reasonable for an agent to believe this only if it is reasonable for her to believe

both that the other is aware of her proximate end and that if he was aware of her further ends, he would not have chosen to avert the agent's use of him.

For the sake of simplicity, let us assume that in order to count as giving his informed consent to be used in a certain way, a person need have knowledge of the agent's proximate end in using him, but not of his further ends. This assumption makes no substantive difference to the discussion that follows.

O'Neill points to some genuine, but not irremediable, difficulties with formulating and applying an actual consent account. But since analogous difficulties apply with equal force to possible consent accounts, the difficulties do not give us good reason to move from actual to possible consent in formulating conditions under which using another does not amount to using him merely as a means.⁸

3: Just Using and Inducing Vulnerability

Another difficulty confronts both actual and possible consent accounts. Since this difficulty has not been appreciated, it makes sense to investigate it in detail, beginning with a case: A customer has called a technician to his house in order to ready his new computer for use. But the technician intentionally leaves his machine vulnerable to malware in an effort to get business repairing it in the future. A few weeks later, the technician returns to the customer's house. She has learned from him that he aims today to email a document, the only copy of which is on his machine. The document must arrive at its destination today, or the customer will be laid off of his job. But his computer is frozen and she is the only one in position to fix it. The technician uses him to make a profit by getting him to authorize her to do the repair at her usual fee. I think we would say that in getting

the customer to authorize the repair, the technician is treating him merely as a means.

However, the Actual Consent Account seems to imply otherwise. We can easily imagine the customer giving his voluntary, informed consent to the technician on her second visit to repair his computer for her usual fee. The Actual Consent Account requires that the consent be based on a reasonably accurate understanding of which aspect of the agent will be used (namely, his capacity to authorize the repair) and to what proximate end (namely, for profit). But the customer's consent might, of course, fulfill these criteria without it being given against the background of information regarding how he came to need to have his computer repaired in the first place. Of course, on the previous service call the customer presumably did not give his informed consent to paying the technician for her work (e.g., because she misled him about what she did). So the Actual Consent Account allows us to conclude that the technician treated him merely as a means *on that occasion*. Yet, implausibly, it seems to imply that she does not do so on the second occasion. This implication is implausible because it seems to be a hallmark of one's "just using" another that one behave as the technician does, that is, that one try to profit from a vulnerability in another for which one bears responsibility.

Defenders of the Actual Consent Account might respond to this criticism in various ways. First, they might claim that, on her second visit, the technician is coercing the customer into paying for a repair. Therefore, he does not voluntarily consent to pay for it. But this claim is implausible. A person counts as coercing another to do something only if he threatens to make the other worse off if he fails to do it (e.g., 'If you don't give me your wallet, I'm going to shoot you'). But the

technician does not threaten to make the customer worse off if he fails to go through with the transaction; she does not threaten him at all.

Second, a defender of the Actual Consent Account might claim that the customer does not really give his informed consent to the technician when she visits him for the second time. The customer would be in position to give informed consent only if he understood the technician's role in bringing about his need for service. In general, the defender might hold, a person can give his informed consent to an agent's using him only if he is informed about her role, if any, in putting him in a situation in which it might further his interests to consent to the agent's using him.

This notion of informed consent seems too exigent. Suppose that, under an unremovable veil of anonymity, a benefactor has promoted a young pianist's career, putting her in a situation in which it will be to her advantage to agree to the benefactor's hiring her to perform in an upcoming concert. On this notion, the pianist could not give her informed consent to the benefactor's use of her—even if she had knowledge not only of all the details of the proposed performance but also of the benefactor's ends in booking her for it.

In any case, it remains implausible to hold that an agent avoids treating another merely as a means if the other gives his informed consent to be used by her, even on this notion of informed consent. Consider again the example of the computer repair. Suppose that the customer finds out that, on the initial service call, the technician had intentionally left his machine vulnerable to malware. Aware of this, the technician nevertheless demands on her second service call that the customer pay her the usual fee to do a repair. The customer might nevertheless consent to this, that is, to be treated as a means to the technician's

making a profit. If he does, he gives his informed consent, even on the expanded notion of such consent. However, it still seems plain that the technician is “just using” the customer and thereby acting wrongly.

A defender of the Actual Consent Account needs to find a path to the conclusion that, when the technician repairs the customer’s computer on her second visit, she does not escape the charge of treating him merely as a means. The defender might insist that, despite what it appears, the customer does not on that visit consent to the technician’s use of him. Our original description of the case implies that the technician uses the customer to make money twice: once on the initial service call and again when she repairs his machine, giving him access to his document. But the defender might insist that the technician’s use of the customer actually constitutes one complex action that has two parts. One part occurs on her initial visit when she attempts to make it the case that the customer will need her in the future; the other part occurs on her second visit when she does the repair and profits from this need. By describing the technician’s use of the customer in this way, the defender wants to be in position to say that the customer never gives his informed, voluntary consent to the technician’s use of him; for he never gives such consent to being put by her in a position such that he will likely need her services in the future. It is evident that the customer never gives his informed, voluntary consent to this even if he finds out what the technician did to him on her first visit and, on her second visit, nonetheless agrees to her repairing his machine. So, concludes the defender, the Actual Consent Account avoids the implausible implication that the technician is not on her second visit treating the customer merely as a means.

A couple of responses to the defense seem to be in order. First, it makes the Actual Consent Account much more complicated to employ in practice than one might have suspected. In order to determine whether someone has given his informed, voluntary consent to being used by an agent, we need to know what this use is. And we cannot determine what it is simply by looking at what the agent is doing to the other at one particular time in pursuit of one particular proximate end. So, for example, suppose that a government establishes a regulated market in human kidneys and sets itself up as the sole buyer. In purchasing a kidney from someone, that is, a vendor, the government is using him—for example, in order to save the life of another citizen. But, according to the defender, we cannot conclude from the fact that the vendor understands all the potential risks (e.g., surgical complications) and benefits (e.g. money) of selling his kidney and sells it voluntarily that the government has not treated him merely as a means. In order to reach this conclusion we would need to know whether the government's use of the vendor was a complex action, one part of which was its buying the vendor's kidney, but another, earlier part of which was, say, an effort to convince citizens through a public relations campaign that selling a kidney is a noble thing to do. Suppose that the government's action was complex in the way described and that the vendor's agreement to sell his kidney was conditional on his having been convinced by the government that doing so was noble. We would have license to conclude, based on the Actual Consent Account, that the government did not treat the vendor merely as a means when it purchased his kidney only if we had good reason to believe that the vendor gave his informed, voluntary consent to being subject to the public relations campaign. In sum, the proposed defense of the

Actual Consent Account introduces a way, in addition to those suggested by O'Neill, in which the Actual Consent Account is hard to apply.

In any case, I believe that the Actual Consent Account also implausibly implies that the technician does not treat her customer merely as a means if we alter the example: On her initial service call, the technician is hurried and cuts corners in readying the customer's new machine for use. She foresees that her cutting corners will leave his machine vulnerable to malware and that she might be needed to repair it in the future. Although she says nothing about this to the customer, it is not her aim to set the customer up so that she can profit from him later. After a week, events transpire as they do in the initial version of the example. The technician learns that the customer is desperate to email a document and that she is needed to make this possible. She then uses him to make a profit by getting him to authorize her to do the repair at her usual fee.

In this second case, it seems that she treats the customer merely as a means by profiting from a vulnerability in him that she has foreseeably contributed to bringing about.⁹ Yet the Actual Consent Account implies, implausibly, that she does no such thing; for the customer gives his voluntary, informed consent to serve as a means for her to make a profit. Here a defender of the Actual Consent Account has no grounds to describe the technician's use of the customer in terms of a complex action composed of an attempt to set the customer up to need her services in the future and an attempt to profit from this set-up. The technician simply never attempts to set the customer up.

It is important to observe that a simple possible consent account has the same implausible implication. Such an account might hold, for example, that an agent does not treat another merely as a means if it is reasonable for the agent to

believe that the other can consent to her use of him. Moreover, the other can consent to her use of him if he can avert this use by withholding her consent. On the technician's second visit, it is reasonable for her to believe that the customer can avert her use of him in this way. For, as she knows, all he needs to do is to tell her 'No thank you' and she will leave.

We might be tempted to maintain that it is actually *not* reasonable for the technician to believe that the customer is able, by withholding his consent, to avert her use of him to make a profit. After all, as the technician is aware, the customer needs to get his computer repaired to keep his job. However, this reasoning conflates the impossibility of someone's averting an agent's use of him with the undesirability from his perspective of his averting this use. The customer has the power to prevent the technician's use of him, albeit at the risk of his job, in a way that victims of coercion and deceit do not. As the technician knows, the customer can stop her from using him simply by not consenting to hire her. The victim of a mugging can withhold his consent to being used by the mugger all he wants. But the mugger is going to use him in order to get money, regardless of whether he does it the "easy way or the hard way." So, in sum, it is reasonable for the technician to believe on her second visit that the customer can consent to his use of her. Of course, we might respond by modifying the possible consent account so that a person can avert an agent's use of him only if his disallowing this use would not result in a decrease in his well-being. But this modification would have a high price. For example, it would imply that someone who needs heart surgery in order to survive and enjoy a good life cannot consent to the operation.

The plausibility of a possible consent or an actual consent account of an agent's using another, but not merely as a means, increases if it avoids implying

that in cases where an agent foreseeably induces vulnerability in another and profits from it, she is not treating the other merely as a means. Here is a Possible Consent Account that tries to accomplish this:

Suppose an agent uses another. She does not use him merely as a means if it is reasonable for the agent to believe:

that the other can consent to her use of him

and

that nothing she has done to the other has contributed to his being in the position that unless the agent herself uses him, he will undergo a significant loss of well-being.

Of course, one might modify the Actual Consent Account in an analogous way. People have various views on what constitutes well-being. They disagree about whether it consists in experiencing pleasure, satisfying one's desires, realizing a set of goods (e.g., including friendship and knowledge), or some combination of these. The Possible Consent Account invokes what it is reasonable for *the agent* to believe regarding the other's well-being. It evaluates the agent's action in terms of her own views regarding the other's well-being, including what his well-being consists in, as long as these views are reasonable.

4: From Actual toward Possible Consent

So far we have found no reason for moving away from actual consent and toward possible consent in an account of conditions under which an agent uses another, but not merely as a means. But I think there is such a reason, which the following case helps to illustrate.

A 65 year-old salesperson believes that her company is trying to get her to retire by keeping its latest sales leads from her. Desperate to make a sale, she intends to use the office manager to get the leads. They are kept in a database, to

which he has the password. With no moral qualms, she tells him that she really needs to close some deals, and unless he gets the leads for her, she is going to reveal to everyone in the office that he is gay. She believes reasonably, given her limited understanding of him and of the attitudes of her other co-workers, that this revelation would be damaging to his reputation. But the office manager takes the salesperson, whom he thinks of as a friendly colleague, to be making a misguided joke. Everyone in the office who cares to know is already aware of his sexual preference. And, he believes, she knows full well that it is company policy that all salespeople are to be granted access to the latest leads upon request. He gives her a puzzled look and agrees to get her the leads right away.

The salesperson receives from the office manager his voluntary, informed consent to her use of him to obtain the leads. He understands that she intends to use him for this purpose. Although she threatens to make him worse off if he does not give her the leads, it is not the threat, which he does not even perceive as such, that prompts him to agree to her use of him. He does so voluntarily. Yet, despite getting his consent for her use of him, many of us believe that the salesperson treats the office manager merely as a means and acts wrongly.

It would be easy to multiply cases like this in order to illustrate that a person's giving her voluntary, informed consent to being used is compatible with her being used merely as a means. The cases would include ones in which an agent attempts to coerce or deceive another into allowing her to use him, but in which the other allows himself to be used by her for reasons that have nothing to do with this attempt.

One might object to the notion that such cases illustrate this point. Inspired by Thomas Scanlon, one might claim, for example, that the salesperson does not

act wrongly, that is, what she does is not morally impermissible. After all, her coercive threat is impotent. Nevertheless, her use of the office manager does manifest a moral fault. In light of her belief in the force of her threat, she ought to, but fails to, *see* the way she uses the manager as morally impermissible.

In order to assess this objection, let us return to its source. Scanlon introduces a distinction between factors that make an action wrong and factors that make it an action that the agent should see as wrong:

This distinction is clearest in cases in which the action an agent proposes to take is in fact utterly harmless. A person who believes in voodoo, for example, may think that by sticking pins in a doll he is bringing about the agonizing death of his former girlfriend's new lover. But there is no reason to think that sticking pins into a doll is in fact harmful. So how could it be impermissible? There does seem to be something wrong with the action I have described. What is wrong with it, however, is not that it is impermissible but rather that the agent should (given his beliefs) see it as impermissible. (46)

Let us call the person practicing voodoo Kulev (which apparently means "harmless snake" in voodoo terminology). Scanlon seems correct to hold that Kulev should, morally speaking, see his action as morally impermissible. But, at least to many of us, he seems to be on shaky ground in suggesting that Kulev's action is actually morally permissible. After all, he is trying to bring about the agonizing death of an innocent person.

Scanlon supports his suggestion that Kulev's action is morally permissible with the assertion that there is "no reason to think that sticking pins into a doll is in fact harmful." But this assertion prompts the question: no reason *for whom* to think that? Granted, there is no reason for someone with knowledge of the context, including knowledge of the actual effects of Kulev's sticking pins into this wax figure, to believe that his doing so will kill the new lover. But that no person with

such knowledge would believe his action to be harmful is compatible with its being wrong.

In May 2009, James Cromitie rigged two vehicles parked in front of a Bronx synagogue with powerful bombs, or so he believed. But, as someone with knowledge of the context would realize, the bombs were actually fake: Cromitie had purchased them from an FBI informant (Sataline et al 2009). That there was no reason for a person with knowledge of the context to think that the devices were harmful does not entail that it was morally permissible for Cromitie to use them in an effort to maim and kill innocent people. There was presumably reason for Cromitie himself to believe that his manipulating the devices would harm others. Given his epistemic situation, it was reasonable for him to believe that they would explode.

In a similar way, given Kulev's epistemic situation, which is a function of his native abilities, his upbringing, and his circumstances, perhaps *he* has reason to believe that he can kill his ex-girlfriend's new lover by sticking pins in a wax representation of him. The strangeness of this belief might prompt the reaction that Kulev simply cannot have reason to embrace it. But is this belief really any less strange, from a scientific perspective, than the belief, apparently held by almost eight-in-ten Americans, that "miracles still occur today as in ancient times?" (Pew Forum on Religion and Public Life 2008, 11). In light of their upbringing, experience, and so forth, might not some of these Americans have reason to believe this? Of course, an affirmative answer to this question in no way implies that we ought to embrace the belief as true.

Let us now return to the case of the salesperson. Someone with knowledge of the context in which she tries to get the latest leads from the office manager

would have no reason to believe her threat to have any causal power.

Nevertheless, the salesperson has reason to believe that it does. Kulev and Cromitie's having reason to believe in the efficaciousness of their attempts to bring about harm supports the conclusion that the attempts are themselves wrong, not merely actions that each ought to see as wrong. In a similar vein, the salesperson's having reason to believe in the efficaciousness of her threat supports the notion that the way she uses the office manager *is* wrong, not merely something she should see as such. So the Scanlon-inspired objection does not invalidate the claim that an agent can use another merely as means (and thus wrongly) even if the other gives his voluntary, informed consent to being used by her. The Actual Consent Account is vulnerable to a serious objection.

Of course, the salesperson case differs from the two others mentioned. Kulev is presumably using the new boyfriend in an effort to hurt his ex-girlfriend, and Cromitie is presumably using the killing and maiming of Bronx Jews in order to promote his religious ideology. In those cases, the envisaged use of a person would itself be harmful to the person. But using the office manager as a means to get leads would not itself be harmful to the manager. What, in the misguided view of the salesperson, would harm him is the revelation of his sexual preference. This disanalogy between the cases would be important if the notion of using others merely as means included by its very concept only the harmful using of them. (There would then be little reason to believe that the salesperson was treating the office manager merely as a means.) But there is no such conceptual connection. Whether an agent treats another merely as a means is not a function of whether her use of the other actually harms him or is intended to harm him, but of *how* she uses him, which includes what she does or tries to do in order to make this usage

possible. What the salesperson does in an effort to be able to use the office manager as a means makes it particularly plausible to conclude that she treats him merely as a means.

In light of cases such as that of the salesperson, one might, of course, amend the Actual Consent Account. Let us call the following account, which incorporates not only the changes we have just found necessary, but also the ones we found we needed at the end of section 3, the Modified Actual Consent Account:

An agent uses another, but does not use him merely as a means, if it is reasonable for her to believe that

the other gives his voluntary, informed consent to her using him,

she has made no attempt to get him to agree to her using him other than that of trying to gain his voluntary, informed consent to her using him, and

nothing she has done to the other has contributed to his being in the position that unless the agent herself uses him, he will undergo a significant loss of well-being.

This account would avoid the implausible implication that the salesperson does not treat the office manager merely as a means. She has threatened him in order to get him to yield to her using him.

Notice how far we have come. We started out with the notion that an agent does not use another merely as a means if the other gives his voluntary, informed consent to her using him. We started out, we might say, with a *patient-focused* account. According to the Actual Consent Account, whether an agent treated another merely as a means depended primarily on the state of the other (the patient), that is, on whether he understood the use he was going to be put to and voluntarily consented to being put to that use. But we have found that, intuitively speaking, a patient-focused actual consent account is inadequate. Whether an agent uses another merely as a means can be a function not only of the other's

state, but also of what the agent tries to do, regardless of whether it affects the other's state. What an agent attempts to do to another in order to make practically possible her use of him can influence whether she treats the other merely as a means.

Our modifications to the Actual Consent Account have, for practical purposes, moved it far closer than it originally was to the Possible Consent Account. An agent's attempt to get another to agree to her using him other than by trying to gain his voluntary, informed consent would likely amount to an attempt to deprive the other of an opportunity to avert her use of him by withholding his consent to it. For example, the salesperson tries to get the office manager to give her the leads with the help of a threat, thereby trying to deprive him of the opportunity to prevent her use of him simply by dissenting from it. Conversely, central cases of an agent's preventing another from having an opportunity to consent to her use of him, that is, cases of deceit and coercion, involve attempts by the agent to get another to submit to her using him other than by gaining his voluntary, informed consent. If the agent successfully employs deception, then the other's consent is not informed. If she successfully employs coercion, then the other's consent is not voluntary. A defensible version of an actual consent account veers close to a possible consent account.

It would be remiss not to acknowledge that both accounts are subject to an objection.¹⁰ Echoing a case discussed earlier, suppose that a patient with a serious heart condition needs and wants an operation immediately in order to survive. A surgeon, who has never before interacted with the patient and who is the only one who can perform the surgery, proposes to operate not at her typical, already considerable fee, but for an even higher one. She informs him of the risks of the

operation and so forth, and makes it clear to him that unless he pays her fee, which she acknowledges to be higher than usual, she will not operate. The patient agrees, and the surgeon uses him to make a profit, saving his life in the process. The Modified Actual Consent Account implies that the surgeon does not treat the patient merely as a means. For it is reasonable for her to believe that the patient gives his voluntary, informed consent to her using him; she has not tried to get him to agree to the operation through means other than asking for this consent (e.g., she has not employed deception); and nothing she has done to him has contributed to his being in a situation in which her use of him is necessary to prevent a significant loss of well-being. The Possible Consent Account also implies that the surgeon does not treat the patient merely as a means. For the patient can avert the surgeon's use of him by withholding his consent to it.

How implausible are these implications? Of course, the accounts do not aim to capture the whole of morality. They try to specify conditions under which someone uses another, but does not treat him merely as a means. But someone can act wrongly with respect to another without using him at all, for example, by behaving toward him as one would behave toward an inanimate obstacle to be kicked away. Moreover, to say that, if certain conditions are met, someone is not using another merely as a means fails to imply that the person's *use* of the other is morally permissible. The accounts leave open the possibility that the surgeon's use of the patient is wrong because it fails to accord with her professional obligations as a licensed physician, or because it discriminates unfairly between this patient and others whom she charges her normal fee, or, perhaps most obviously, because it exploits the patient in the sense that it takes unfair advantage of him.

Both the Modified Actual Consent and Possible Consent accounts imply that some cases of what we might intuitively label as an agent's taking unfair advantage of another are not also cases of treating the other merely as a means. They might, therefore, fail to capture the full range of cases in which some find it natural to say that someone is "just using" another.

But I doubt whether any philosophically satisfying account of not treating others merely as means is going to square entirely with ordinary judgment. We might try to build into an account a condition such that if an agent takes unfair advantage of another, we must refrain from concluding that she is not treating the other merely as a means. But such an account would not be illuminating. The notion of taking unfair advantage of another is closely related to that of just using another. (We might describe the technician's behavior towards her customer either as her just using him or as her taking unfair advantage of him, for example.) Moreover, the notion of taking unfair advantage of another is no less in need of clarification than that of treating another merely as a means. For example, suppose that the patient is impoverished. Does the surgeon take unfair advantage of him if she charges her usual fee, or half of her usual fee, or less than that, but more than he can easily afford? It would be easy to multiply such questions. In any case, it is hardly surprising that rendering precise a commonsense idea of treating others merely as means involves some narrowing of the idea's scope. The accounts we have ended up with might not accord fully with ordinary thinking. But they nevertheless promote understanding of the notion of treating others merely as means.

Returning to where we began, the Actual Consent Account has strong initial appeal. If someone gives his informed, voluntary consent to serve as a means to some end, then how could it be that he is being treated merely as a means? After all, he agrees to be used. I hope to have shown the account's appeal to be questionable. Someone's agreeing to be used by an agent is compatible with her using him merely as a means. This compatibility can be manifest when the agent's use of another involves her profiting from some vulnerability in him that she has helped to induce. It can also be evident when the agent's use of the other involves an otiose attempt at deceiving or coercing him. We can amend the Actual Consent Account so that it is sensitive to the possibility of an agent's treating another merely as a means in these ways. But the necessary changes yield an account that is both considerably more complex than the Actual Consent Account and that, for practical purposes, differs little from possible consent accounts.¹¹

Samuel J. Kerstein

University of Maryland, College Park

References

- Brock, Dan W. 1993. "Informed Consent." In *Life and Death: Philosophical Essays in Biomedical Ethics*. Cambridge: Cambridge University Press, 21-54.
- Cherry, Mark. 2005. *Kidney for Sale by Owner*. Washington D.C.: Georgetown University Press.
- "Dr. Jane Goodall." 2009. *The Animal House*. Interviewed by Gary Weitzman. WAMU FM: American University Radio. 19 September. http://wamu.org/programs/the_animal_house/2009/09/19/ (accessed 26 November 2010).
- Kerstein, Samuel. 2009. "Treating Others Merely as Means." *Utilitas* 21: 163-180.
- Nozick Robert. 1974. *Anarchy, State, and Utopia*. New York: Basic Books.
- O'Neill, Onora. 1989. "Between Consenting Adults." In *Constructions of Reason*. Cambridge: Cambridge University Press, 105-125.
- Parfit, Derek. 2008. *On What Matters* [manuscript].
- Pew Forum on Religion and Public Life. 2008. *U.S. Religious Landscape Survey: Religious Beliefs and Practices: Diverse and Politically Relevant*. Washington, D.C. June. <http://religions.pewforum.org/pdf/report2-religious-landscape-study-full.pdf> (accessed 26 November 2010).
- Sataline, S., Bray, C., and Fields, G. 2009. "Four Men Held in U.S. Terror Case" *The Wall Street Journal*. 22 May, A3.
- Scanlon, Thomas. 2008. *Moral Dimensions: Permissibility, Meaning, Blame*. Cambridge, MA: Harvard University Press.
- Schaber, Peter. 2009. "Using People Merely as a Means" [manuscript].
- Wertheimer, Alan. 1996. *Exploitation*. Princeton: Princeton University Press.

Notes

¹ The understanding of the Mere Means Principle I employ here is narrower than the understanding employed by some (e.g., Parfit 2008), according to which the principle specifies as wrong treating anyone, presumably including oneself, merely as a means.

² As I argue in Kerstein 2009, in order to be plausible, the Mere Means Principle must be modified to assert that it is wrong *pro tanto* to treat others merely as means. In extreme situations, for example, when the only way to save many thousands of innocent people is to treat one such person merely as a means, then doing so might not be wrong *all things considered*. But I set aside that complication in this paper.

³ I do not wish to imply that Nozick actually embraces this view. His remarks are brief and leave it unclear whether he does.

⁴ Some philosophers (in my view incorrectly) attribute this view to Kant himself (e.g., Cherry 2005, 98).

⁵ This account is based on, but not identical with, an account Thomas Scanlon offers (Scanlon 2008, 106-7). Derek Parfit (2008, 183, 191) suggests some forceful, but in my view ultimately inadequate, reasons for rejecting it. But I cannot give them the attention they deserve here.

⁶ For discussion of competence as it applies in healthcare settings, see Brock 1993, 36-43.

⁷ This sort of example is not fanciful, as an interview with Jane Goodall reveals (“Dr Jane Goodall” 2009).

⁸ Both actual and possible consent accounts of when an agent uses another but not merely as a means imply that there are circumstances in which an agent might kill another for his body parts without “just using” the other. We can, for example, imagine that a competent adult gives his informed, voluntary consent to be killed by a transplant surgeon so she can use his organs to save the lives of five other people, and that the surgeon gives him ample opportunity to avert his use of him for this purpose. Some (Schaber 2009) hold that, if the surgeon goes ahead with the operation in order to save the five, she treats the person

merely as a means. But while we might conclude that the surgeon is acting wrongly, I do not believe we would say that she is treating him merely as a means. The surgeon would be doing that, for example, if she did not give him the opportunity to prevent her from harvesting her organs.

⁹ Reflections such as this might lead us to develop an account of jointly sufficient conditions for an agent's just using another. An Induced Vulnerability Account might take the following form:

Suppose an agent uses another. She uses him merely as a means if:

1. It is reasonable for the agent to believe that something she has done to the other has contributed to his being in the position that unless the agent herself uses him, his well-being will diminish significantly.
2. It was foreseeable to the agent that she would contribute to his being in this position as well as avoidable that she would do so; and
3. The agent was, as a practical matter, able to but did not give the other an opportunity to dissent from taking the risk that she (the agent) would contribute to the other's being in this position.
4. The end of the agent's use of the other is not limited to her discharging what she reasonably believes to be a moral obligation towards him.

¹⁰ There is an orthodox Kantian view according to which even the Modified Actual Consent Account is inadequate. As a variation on our previous example, imagine that the salesperson suspects, but is far from sure, that the office manager will refuse to give her the latest leads. She acts on the following plan: she will ask him for the leads, and if he does not give them to her immediately, she will threaten to reveal his sexual preference. When she asks him for the leads, he gives them to her right away. So in the end she makes no threat. According even to the Modified Actual Consent Account, it turns out that she is not using him merely as a means. Not only does he give his voluntary, informed consent to her use of him, but she makes no attempt to get him to agree to her usage of him through

other means. Those who adopt this orthodox Kantian perspective might champion a possible consent account of conditions under which an agent *does treat* another merely as a means that entails that the salesperson, in the latest variation of the example, is doing just that. According to these Kantians, an agent does not use another merely as a means unless she uses him, of course. But the action of using itself, defined roughly as doing something to another in order to attain one's end, does not exhaust what is relevant to consider in deciding whether one is using another merely as a means. Also relevant is *how* one uses the other, including, according to them, which options one *intends* to leave for the other as one employs him as a means. I do not have space to explore the plausibility of this view here.

¹¹ For very helpful comments on this paper I would like to thank Steve Emet, Ryan Faselow, two referees for Oxford University Press, and participants at the 2010 the Arizona Workshop in Normative Ethics, especially Tom Christiano, Thomas Hill, Jr., and Mark Timmons.