Journal of Applied Philosophy

Journal of Applied Philosophy, Vol. 27, No. 3, 2010
doi: 10.1111/j.1468-5930.2010.00481.x

A Puzzle about Consent in Research and in Practice

ERIC CHWANG

ABSTRACT *In this paper, I will examine a puzzling discrepancy between the way clinicians are allowed to treat their patients and the way researchers are allowed to treat their subjects: in certain cases, researchers are legally required to disclose quite a bit more information when obtaining consent from prospective subjects than clinicians are when obtaining consent from prospective patients. I will argue that the proper resolution of this puzzling discrepancy must appeal to a pragmatic criterion of disclosure for informed consent: that what needs to be disclosed in order for consent to be valid depends on what the patient/subject needs to know in order to make a decision. I will then use this pragmatic criterion of disclosure to argue that when obtaining consent researchers should be permitted to omit the same information clinicians are, given certain qualifications. I will also examine how this puzzle forces us to confront some perhaps surprising truths about valid consent. My broader aim in this paper is to examine, not so much the puzzle itself, but rather what this particular puzzle can teach us about more theoretical issues surrounding informed consent.*

1. Introduction

In this paper, I will examine a puzzling discrepancy between the way clinicians are allowed to treat their patients and the way medical researchers are allowed to treat their human subjects: in certain cases, to be described below, researchers are legally required to disclose quite a bit more information when obtaining consent from prospective subjects than clinicians are when obtaining consent from prospective patients. I will argue that the proper resolution of this puzzling discrepancy must appeal to a pragmatic criterion of disclosure for informed consent, that what needs to be disclosed in order for consent to be valid depends on what the patient/subject needs to know in order to make a decision. I will then consider and reject a rival hypothesis, that what needs to be disclosed is what the patient/subject would want to know. I will argue that this rival does not distinguish adequately between standards of disclosure necessary for the validity of consent, which is my topic, and more general standards for when information should be conveyed, which is not my topic. Finally, I will use the pragmatic criterion of disclosure to argue that when obtaining consent researchers should be permitted to omit the same information clinicians are, given certain qualifications.

My broader concern in this paper is not that the admittedly esoteric puzzle I will examine has wide-ranging and devastating negative consequences and so must be resolved immediately; rather I am interested in what it can teach us about more theoretical issues surrounding the ethics of informed consent. Still, the puzzling issue is important, because the consent process in research is already quite complex and difficult to understand.[1] It could be shortened and simplified considerably if research obligations

more closely followed clinical obligations. Further, complexity in the consent process plausibly deters many potential subjects from enrolling. Needless complexity, then, should be eliminated.

The original seed for the position I defend was articulated most recently in Troug et al.'s 1999 paper, 'Is Informed Consent Always Necessary for Randomized Controlled Trials?,'[2] though the puzzle has been mentioned before.[3] Truog *et al.* conclude that, at least in some circumstances, researchers should be permitted to skip over the same details that clinicians are, but just about all commentators disagree with them.[4] I will defend Truog *et al.*'s conclusion by showing how it follows from a plausible pragmatic criterion of disclosure for informed consent.

2. Clinical Indifference

I begin by explaining the puzzling discrepancy of concern. The discrepancy turns on a concept which I will call *clinical indifference*, and which can be explained as follows. Suppose there are two drugs approved to treat your illness, tied for best for all we know, even after factoring in circumstances unique to you. Assume, if you like, that they have similar dosing requirements, side effect profiles, costs, and so on. We will say that in such situations the two drugs are in clinical indifference, because a clinician treating your illness could arbitrarily give you either without telling you first that there are two or that your treatment was chosen arbitrarily. For example, normal saline and lactated Ringer's solution are plausibly in clinical indifference for the intravenous treatment of mild dehydration.

Four points about clinical indifference deserve mention. First, clinical indifference occupies one extreme end of a spectrum which measures the superiority of one treatment over another at alleviating a given condition. At the other extreme end of the spectrum would be complete dominance of one treatment over the other in every important aspect, such as efficacy, cost, and convenience. Obviously, absolute clinical indifference is a theoretical limit, not practically achievable, but many treatment-pairs can get very close to perfect clinical indifference, with respect to a given condition they are employed to treat. Second, clinical indifference does not require that the two treatments in question really are tied for best. All it requires is that we are justified in believing that they are tied for best, given our current state of ignorance. Third, clinical indifference is not the same as equipoise. Equipoise refers to a situation in which neither treatment in question is known to be inferior, and it is constrained only by considerations of efficacy.[5] In contrast, clinical indifference requires that, based on all currently available evidence, the two drugs are considered equally good, and it encompasses all relevant factors, such as cost, convenience, and side effects, not just efficacy. Thus, many more treatment-pairs are in equipoise than are in clinical indifference. Finally, I do not mean to imply that clinicians should or in fact do choose arbitrarily when the best treatments are in clinical indifference. However, it would be in principle permissible, even if wrong,[6] for the clinician to choose arbitrarily in such a situation without disclosing the existence of the unchosen alternative.

Now I can explain the puzzling discrepancy that motivates this paper. What the clinician is permitted to do — namely, refrain from disclosing that there was an arbitrary choice between a chosen treatment and its clinically indifferent alternative(s) — the

researcher is widely thought to be forbidden from doing. That is, if a researcher is testing two treatments in clinical indifference, she is required to disclose quite a bit of information to her subjects when obtaining consent, including that there will be an arbitrary choice between two treatments in clinical indifference. For example, a clinician treating mild dehydration might just randomly choose whichever bag was closest to her arm, normal saline or lactated Ringer's, without bothering to inform her patient about the existence of the unchosen alternative. In contrast, a researcher who wants to run a head-to-head trial comparing normal saline to lactated Ringer's in the treatment of mild dehydration must inform her subject that there are two alternative treatments, and that his treatment will be chosen randomly.

Be careful not to frame the puzzling discrepancy as concerning whether consent (for care or research) must be obtained. The correct way to frame my question is as concerning what information must be disclosed (in clinical care or research) when obtaining consent. Even Truog *et al.*'s paper is misleading in this regard at times, as witness its title, 'Is Informed Consent Always Necessary for Randomized Controlled Trials?'. The proper answer to that title question is 'Yes, but what needs to be disclosed when obtaining informed consent might be less than one would expect'.

Now, some might already object that even the clinician should be forbidden to treat without first disclosing the existence of (all?) clinically indifferent alternatives and that the choice between them was arbitrary. This objection is bolstered by the idea that absolute clinical indifference is a theoretical limit which in reality is never reached, so that in practice there is always a (non-arbitrary) sufficient reason to choose one treatment over the other.

However, this objection misses an important reality of clinical life. To return to a prior example, if a clinician is treating mild dehydration, she might arbitrarily use either normal saline or lactated Ringer's, just grabbing whichever bag happens to be closer. When she does so, the clinician need not inform the patient about the existence of the unchosen alternative or that the choice between normal saline and lactated Ringer's was essentially arbitrary. Further, it is implausible to think that clinicians are permitted to skip these disclosures because they in fact do know of an important clinical difference between normal saline and lactated Ringer's (which they are then justified in refraining from disclosing to their patients), and it is likewise implausible to suppose that clinicians treating mild dehydration with whichever fluid happened to be closest are negligent for failing to inform their patients about other possible fluids that could have been chosen.

More realistic and honest would be to admit that, in at least some cases — perhaps such as the choice between normal saline and lactated Ringer's for mild dehydration — clinicians are permitted to choose between treatments arbitrarily, i.e. without good reason (where one bag being closer to your arm than another does not count as a good reason), and they need not inform their patients of this arbitrary choice. Indeed, such cases abound in clinical care. They include not only some drugs that are close analogues but also different brands of medical devices and instruments. Researchers may want to test, for example, which of two brands of pulse oximeter is more accurate, where both brands are already in use and accepted as reasonable choices. If so, they are legally required to disclose to their subjects that there are two brands and that the choice between them will be random. A clinician, in contrast, could arbitrarily grab whichever oximeter happened to be closest, and clip that one to her patient's finger.

3. A Pragmatic Criterion of Disclosure

At the heart of the puzzling discrepancy is the question of what information needs to be disclosed to someone prior to an intervention, whether therapeutic or experimental, in order to for her consent to that intervention to count as adequately informed and therefore valid. My concern in this paper is with a special case of this more general question about disclosure, namely the special case where the relevant information under consideration for disclosure focuses on two treatments in clinical indifference, so that experts have no good reason to think one treatment is any better than the other in any respect. Thinking about this special case will bring into the sharpest possible focus exactly what criteria should be implemented for disclosure standards in informed consent.

Research and clinical ethics documents often address the question of disclosure with an objective list. For example, the US federal regulations for research require mention of, among other things, the purpose of the research, its risks and benefits, and relevant alternatives.[7] Objective lists serve an important purpose, but my interest is in why any particular item, such as risks and benefits, deserves to be on the list. What criteria do we use to choose what makes the list? And could it be possible that the same set of criteria will require disclosure of clinically indifferent treatments in the context of research yet not in the context of clinical care?

The best way to answer these questions is to recall the purpose of disclosure: to enable the patient to make an autonomous decision. Therefore, the information that needs to be disclosed should be constrained by what the subject/patient would consider relevant in choosing whether to undergo the intervention. Thus, for example, most subjects and patients consider risks to be relevant to their decision, but most consider the history of their procedure to be irrelevant. Therefore, a procedure's risks may need to be disclosed, but not its history. Consider a more extreme example. Suppose someone were to insist that all clinicians and researchers must always disclose the fact *that Canada is north of the United States* to all patients and subjects, without exception. This is of course absurd, because the fact that Canada is north of the United States is not relevant to the decision whether to consent to treatment or participation.

This is not a new insight; it was articulated in the federal courts as far back as the landmark *Canterbury v. Spence* opinion in 1972, regarding a case where a young man consented to back surgery without being informed of some of its most salient risks:

> In our view, the patient's right to self-decision shapes the boundaries of the duty to reveal. That right can be effectively exercised only if the patient possesses enough information to enable an intelligent choice. The scope of the physician's communications to a patient, then, must be measured by the patient's need, and that need is information material to the decision. Thus the test for determining whether a particular peril must be divulged is its materiality to the patient's decision: all risks potentially affecting the decision must be unmasked.[8]

Two points about this opinion deserve mention. First, the decision bucked the then-dominant tradition of tying standards for disclosure to the judgment of medical professionals. In contrast, this ruling's new paradigm was *patient-centred* rather than doctor-centred: what needs to be disclosed is determined by the patient (or perhaps a hypothetical reasonable patient — more on this shortly), not by what the doctor thinks

is appropriate. Indeed, the case is considered a landmark for precisely this reason.[9] Today about half the states in the US have adopted a patient-centred approach to disclosure,[10] as has Canada's Supreme Court.[11] Courts in England and Australia are now following suit as well.[12]

The appropriate standard for disclosure relies on the patient's perspective, then, not on the perspective of medical professionals. But this point alone still underdetermines what needs to be disclosed, for, even if we appeal to a hypothetical reasonable patient, a patient may want to know things that are irrelevant, and she may neglect to request information which would be pertinent to her decision. Thus, the second point worth noting about the *Canterbury v. Spence* decision is that it gives a further criterion of disclosure, on top of the urge to re-orient towards patients and away from medical professionals. In particular, it articulates a *pragmatic* approach: the way in which a particular patient's opinion matters is determined by what the patient needs, where that means what the patient needs to know in order to make a decision, or in other words what is material to the patient's decision.

Thus, for example, even if a (hypothetical reasonable) patient might like to know about the history of cholecystectomies, her consent to a cholecystectomy would still count as informed even if that history remains undisclosed, on the assumption that the history would have no impact on the patient's decision whether to undergo the procedure. Contrariwise, even if a (hypothetical reasonable) patient has no particular interest in finding about the risks involved with cholecystectomy before her doctor brings the subject up, the doctor is still obligated to disclose them, because information about risks might have an impact on the patient's decision whether to undergo the cholecystectomy. While *Canterbury v. Spence* has been influential in its patient-centred turn, my interest in the decision is focused on the second, under-appreciated criterion it articulates, the pragmatic criterion that, *in order for consent to be valid, whether some piece of information needs to be disclosed depends on whether it might affect the patient's (subject's) decision whether to consent*.

I will end this section with two points about the pragmatic criterion italicized above. First, note that all I require is that information must be divulged *only if* it is material, not '*if and only if* it is material'. For example, an obstetrician is arguably not required to disclose horrific pictures of aborted foetuses to an impressionable pregnant girl seeking an abortion, even if that disclosure would change the girl's decision from 'yes' to 'no'. All that is necessary for this paper is the principle that, for the purposes of ensuring valid consent, any information that is *not* material thereby need *not* be disclosed. Thus, I need not take a stand on certain controversial issues where a converse materiality criterion — 'if material, then must disclose' — has been wielded, such as whether surgeons should be required to disclose their success and complication rates, or the extent to which one must disclose (e.g. genetic) information to potential insurers.[13]

Second, I make no claims as to whether the pragmatic criterion is easy to apply, or whether it is easier to apply than rival criteria. One might think that the pragmatic criterion will be more illuminating than, say, an objective list criterion precisely because at least the pragmatic criterion gives us a principle to determine just how much detail on, say, risks and benefits it is necessary to divulge, while the objective list might merely specify 'risks and benefits'. This strikes me as in general correct, though it is always open to the defender of an objective list to specify her list in more detail. Still, the point of emphasis for this paper is not that the pragmatic criterion is easy to apply. After all, it will

very likely turn out to be a difficult question, subject to intricate conceptual analysis and tortured jurisprudence, to determine just what counts as material. The point is that the pragmatic analysis is correct, not that its application is always straightforward. Indeed, as we will see shortly, the application of the pragmatic criterion to our puzzle requires some finesse.

4. The 'Reasonable Person' Standard

Before applying the pragmatic criterion to our initial puzzle, I consider an objection to the criterion. One way to encapsulate the conclusion of *Canterbury v. Spence* is as advocating a 'reasonable person' standard for disclosure. A common way to interpret this 'reasonable person' standard is that whether information needs to be disclosed depends on whether a reasonable person would want to know it. For example, Dennis Mazur writes:

> Under a reasonable person standard, the decision about whether a patient should have been informed of a risk is based on whether a reasonable person in that patient's position would want to be informed.
>
> The reasonable person standard was established by Judge Robinson in 1972 in a landmark US Federal case Canterbury v Spence . . .[14]

We can distinguish two ideas from this way of expressing the reasonable person standard. The first idea is that we should focus on a hypothetical reasonable person rather than on the particular patient (subject) in question.[15] After all, particular patients may be inscrutable or capricious. Whether the move from actual patients to hypothetical reasonable ones is appropriate is outside the scope of this paper. Fortunately, the move isn't threatening to the pragmatic criterion, which itself remains silent on who the decision-maker is. Thus, to insist that the relevant standard for disclosure should refer to hypothetical reasonable people rather than actual people is not to disagree with the pragmatic criterion but rather to take a stand on how it should be spelled out. My conclusions will apply regardless of which stand is taken.[16]

The second idea that can be gleaned from the expression, 'what a reasonable person would want to know,' is potentially more threatening to the pragmatic criterion. This is the idea that the appropriate standard for disclosure refers, not to what is material to a patient's decision, but rather to what she 'would want to know'. The first point worth making about this idea is that it is clearly not an accurate summary of *Canterbury v. Spence*, as it is sometimes taken to be. Recall the lynchpin portion of that text, here re-quoted with my emphasis:

> The scope of the physician's communications to a patient, then, must be measured by the patient's need, and that need is *information material to the decision*. Thus the test for determining whether a particular peril must be divulged is its materiality to the patient's decision: all risks *potentially affecting the decision* must be unmasked. (My emphasis.)[17]

The opinion does not say that the doctor must divulge anything that the patient wants to know. Rather, it says that the doctor must divulge anything that is material to the patient's decision.[18]

The second point to make about the 'whatever the patient would want to know, regardless of materiality' criterion is that it is incorrect. The pragmatic criterion links the requirement for disclosure to what the patient might find material, whereas the reasonable person criterion links that requirement to what the patient might want to know, regardless of materiality. There are two categories of information, then, where the appropriate verdict, '*must disclose*' versus '*need not disclose*', is contested. The less interesting category is that group of information which might be material, but which the patient does not want to know. For example, a potential cholecystectomy patient might not want to know about the procedure's risks. This category is less interesting to me because, if it embarrasses any criterion at all, it embarrasses the 'what the patient would want to know (regardless of materiality)' criterion. I take it as so obvious that this sort of information should be disclosed that I am confident a staunch defender of the 'what the patient would want to know' criterion would come up with some epicycle to save her theory.[19]

The second, potentially threatening, category of information is information that a patient would want to know but which would not affect her decision. Sometimes we want to know things even if they will not change our decision whether to consent, as for example when I want to know whether a necessary and life-saving intervention will be painful. If I have had an accident, and my leg needs to be amputated in order to save my life, I expect the doctor to tell me whether the procedure will be painful. Surely it would be grossly inappropriate for my doctor to omit that information, even if it would have no effect on my decision whether to consent.

The pragmatic criterion — that the standard for disclosure should be linked to materiality of the consent-decision — respects the intuition that my doctor should answer my question about amputation (honestly), in two ways. First, the most viable version of the pragmatic criterion will not claim that the relevant bit of information must by itself change the patient's mind all the way from 'no' to 'yes' (or vice versa) in order to be required. Rather, it will claim only that the bit of information might move the patient along a continuum away from one decision (e.g. 'no') and towards the other ('yes'). What I mean when referring to a continuum between 'yes' and 'no' is illustrated by the fact that there might be multiple pieces of information, each of which alone is insufficient to change a patient's mind from 'yes' to 'no' (or vice-versa), but which taken together do suffice to change her mind. For example, I might consent to (elective) cholecystectomy if the side effects were disclosed and the cost was hidden, and I might consent if the cost was disclosed and the side effects hidden, even though I would dissent if both its side effects and costs were disclosed. In such a case, the most viable version of the pragmatic criterion would require disclosure of both bits of information.

The second way in which the pragmatic criterion respects the intuition that we sometimes want to know things even if they have no effect on our decisions is with a reminder: there are reasons to disclose information other than the obligation to obtain valid consent. My concern in this paper is with what information needs to be disclosed in order for consent to be valid. The pragmatic criterion addresses that concern, but there is more to the ethics of communicating information than the pursuit of valid consent. To take a simple example, if the clinician's patient really wants to know who won the baseball game last night, it would be (morally) good for the clinician to disclose that information if she knows it. Likewise, if the patient wants to know what her prognosis is,

as most patients do, there is an especially powerful reason, probably powerful enough to count as a role-related obligation, for the clinician to disclose this information too.

Of particular interest is the potential reason that the information might make a big difference in some aspects of the patient's life other than her decision whether to consent. For example, knowing that I will share my hospital room with three other patients might not make any difference to whether I consent to an elective cholecystectomy, but it might lead me to bring some earplugs. Then if no one discloses that I will share my room with three others, and I consent to the cholecystectomy regardless, I cannot legitimately argue that the clinician is not permitted to proceed with the operation. I can, however, legitimately argue that I should have been informed about the existence of roommates, not because it would make a difference to my decision whether to consent but rather because it would have influenced my preparation for my associated hospital stay. In that case, there are good reasons to tell me about the noisy room, and indeed perhaps good reasons to disclose this information when obtaining consent. But the important point is that the reason the clinician should disclose this information has nothing to do with the goal of ensuring that consent is valid; we must distinguish the content of what should be disclosed from the reasons for disclosure, and the goal of ensuring valid consent is served adequately by the pragmatic criterion of disclosure.[20]

Consider how these two responses work together in the provocative case I briefly raised earlier. Suppose that the only way to save my life is by amputation. To get consent for 'life-saving treatment', but without also mentioning that the treatment in question is amputation, would be grossly inappropriate. One reason it would be inappropriate is that even though knowledge that the treatment in question is amputation would not change my decision all the way from consent to dissent — it is life-saving after all — amputation is still a significant setback. Then knowing that the treatment is amputation (rather than, say, massage therapy) pushes me closer to dissent, even if other factors in favour of consent swamp that effect.

But another reason it is also inappropriate to withhold the information — that the life-saving treatment will be amputation — has nothing to do with valid consent. Amputation is a serious life-changing event, and there are good reasons to tell people when they will undergo serious life-changing events, for example to help prepare them. Then just as there are reasons to tell me I will share a hospital room with three other patients, so that I can prepare by bringing earplugs, so also there are similar though obviously much more powerful reasons to tell me that I will be undergoing an amputation, so that I can begin the myriad preparations necessary to live as an amputee.[21]

Before applying the pragmatic criterion to the puzzling discrepancy with which I began this paper, allow me to summarize my most pertinent findings about principles of disclosure so far. We can distinguish three distinct debates in this area. First, there is the debate between patient-centred and doctor-centred approaches. Abstracting from clinical care, this is essentially a debate between subject-centred and intervener-centred approaches to what an intervener must disclose before acting in some way that might affect the subject. Second, there is the choice between hypothetical reasonable subjects and actual subjects. Third, there is a potential debate between a materiality or pragmatic constraint and the broader notion of 'what the subject would want to know' independent of materiality. All that is necessary for the purposes of this paper is that I take a stand in favour of the materiality constraint, so I have simply skirted the first two debates just mentioned. Indeed, each of the three debates seems orthogonal to the other two, so that

there are potentially eight possible positions, corresponding to which of the two options we choose for each of the three debates.[22] Though the first two debates have received extensive discussion in the legal and ethical literature, my sense is that the third debate, between the pragmatic criterion and 'what the patient would want to know' is largely the product of an over-simplified misinterpretation of *Canterbury v. Spence*, when referred to as the so-called 'reasonable person' standard.

5. Applying the Pragmatic Criterion

Now I move on to applying the pragmatic criterion to our initial puzzle. Notice that the pragmatic criterion is silent on whether the intervention in question is clinical care or medical research — it applies to all interventions where valid consent is necessary in order to proceed. Now, of course, one and the same principle might dictate different results in different situations. In particular, my single pragmatic criterion might dictate that clinicians need not disclose certain facts even while also dictating that researchers must disclose them. Thus, we must examine what falls out when we apply the pragmatic criterion to the problem at hand.

A clinician can gloss over quite a bit of information when obtaining consent in cases of clinical indifference, because presumably her patients trust her judgment enough that adding the extra information will not make any difference in their decision whether to consent. However, things are not quite so easy for the researcher. I will consider two arguments for the position that the extra information would make a difference to a (reasonable) subject's decision whether to consent to clinical indifference research. The first argument fails, but the second one enjoys a qualified, temporary success that is susceptible to being blocked, as I will explain.

The first argument I will address claims that the extra information might make a difference because no two treatments are exactly on a par in every respect, including treatment efficacy, side effect profiles, and cost. And surely some (reasonable) people might consider information about such differences to be relevant to their decision whether to consent. This observation is certainly right, and the pragmatic criterion explains why it is relevant and why it is tempting to use this observation to impose high standards on researchers. Two comments are in order, however.

First, though it is surely right that no two (distinct) interventions are exactly alike along every conceivable dimension, including side effect profile and cost, surely sufficiently many are close enough for all practical purposes that the issue is still relevant and important. As mentioned earlier, absolute clinical indifference is the extreme, perhaps impossible-to-instantiate endpoint on a continuum. But the issue still arises for treatment pairs that are very close to that endpoint, if not right on it. Recall the example of normal saline versus lactated Ringer's solution already mentioned. Then the objection that no two treatments are exactly similar in every respect misses an interesting and important issue, which arises when two treatments are very similar in most respects, even if no two distinct treatments are exactly the same in every respect.

Second, to the extent that this observation — that no two treatments are exactly alike — can be used to argue for high standards for researchers, so too can it be used to argue for high standards for clinicians. In other words, the observation that no two treatments are exactly alike in every respect cannot be used to argue that researchers should disclose

more information that clinicians, which is an essential feature of our puzzle. We are assuming that clinicians are permitted to suppress details about alternatives in clinical indifference (even if doing so is wrong), and we ask why researchers are not so permitted.

I now turn to a second materiality-based argument for the position that researchers should disclose more than clinicians. This second argument begins by acknowledging that, in a limited sense, the same thing might happen to a potential subject if she dissented from participating in a clinical indifference trial as would happen if she consented. In particular, she would get one of the two treatments in clinical indifference, and the choice between them would be arbitrary. Return to the dehydration example. Suppose a researcher approaches a potential subject with mild dehydration in the emergency room and asks for consent to participate in a study comparing normal saline to lactated Ringer's. If the potential subject consents, she will receive one of those two fluids, at random. However, even if she dissents, we can suppose that her attending physician (who may be identical to the researcher) would still just choose one of the two fluids arbitrarily. In effect, whether she decides to participate in the study has no effect on her clinical care — either way, she receives one of the two fluids, at random.[23] In a sense, then, this subject will receive the same care regardless of whether she participates in research. Why, then, must her researcher disclose more than a clinician would have in the same situation?

This question is answerable. In a very important sense, being in research does change what happens to patients/subjects, and this is the crux of the second argument, which I will call the *Mere Fact of Research Argument*. The mere fact that research is being done, with its necessary data collection, potential publication, and the intention to help future patients by collecting generalizable data, might plausibly influence whether a (reasonable) potential subject wishes to consent to participate, even if the research is so innocuous that the randomization information would not make any difference at all in an analogous clinical scenario. There are at least two reasons why the mere fact of research might be morally significant. First, we might imagine a reasonable person saying 'I know it won't make any difference to the nature or quality of my care at all, but I just don't want data collected on me, and I don't want to participate in research'. And second, the mere fact of research might matter because research has different goals than clinical care. The goal of clinical care is to benefit the particular patient, whereas the goal of research is to benefit future patients.[24]

6. Preemptive, Umbrella Consent

The Mere Fact of Research Argument seems compelling, and I am willing to grant its cogency for the sake of argument. In the remainder of this paper, I want to discuss some perhaps surprising implications of a brief remark Truog *et al.* made, which threatens to undermine the cogency of the Mere Fact of Research Argument. Although Truog *et al.* think that researchers need not disclose any more information than clinicians in cases of clinical indifference, they soften their conclusion in various ways, such as by insisting that the study be minimal-risk and that the subjects not prefer one treatment to the other. Most relevant to my discussion in particular is their fifth suggestion:

> Fifth, patients should be informed that the institution or clinical setting in which they are being treated uses the standards that we have described as

> guidelines for determining the need for specific instead of general informed consent. Thus, patients would have the opportunity to obtain additional information about the policy or to seek care elsewhere.[25]

The strategy here is to block the Mere Fact of Research Argument by requiring both that consent be preemptive, because it occurs when the patient first agrees to receive health care from their particular chosen institution, and that it be umbrella, because it covers a wide range of (clinical indifference) research all at once.

This strategy blocks the Mere Fact of Research Argument by stipulating that researchers are allowed to omit information about clinically indifferent alternatives only if they have already obtained preemptive umbrella consent to participate in (future, as-yet-unspecified) clinical indifference trials.[26] That way, each potential subject already knows that her medical centre enrols patients in such trials, so that a potential subject whose decision might be altered, if information about research with clinically indifferent alternatives were disclosed, would have been given ample earlier opportunity to withdraw from treatment/research at that particular centre.

What should we think about the moral status of preemptive and umbrella consent? There is some literature on this, and within the confines of this paper I will not attempt to say anything novel, but I want to draw attention to four points relevant to our discussion. First, the extent to which consent is preemptive and umbrella comes in degrees. The only sort of consent that is completely non-preemptive is consent that comes simultaneously with its intervention, and the only sort of consent that is completely non-umbrella is consent for an intervention where that intervention is spelled out in excruciatingly, impossibly fine detail.[27]

A second point follows immediately from the first: because absolutely non-preemptive, non-umbrella consent is at an impossible-to-obtain extreme end of a spectrum, it is clear that ordinary instances of morally binding consent are always, to some extent, preemptive and umbrella.[28] If I today consent to your crossing through my land tomorrow, then that consent is preemptive, because a day early, and also umbrella, because I have not bothered to specify the precise path you are required to take.

Now, even if one admits that all (real-world) consent is to some degree preemptive and umbrella, surely there is some limit to the moral power of such consent.[29] This brings me to my third point. The mere insistence that clinicians and researchers in clinical indifference scenarios should be required to disclose the same information does not entail an endorsement of unrestricted preemptive umbrella consent. First, the domain of its umbrella is restricted to clinical indifference research. It need not require, for example, the permissibility of using preemptive umbrella consent for placebo-controlled research. Second, it does not require that preemptive consent remain valid throughout all eternity: if there are good reasons for thinking that preemptive consent should carry with it a statute of limitations, then those limitations should apply to clinical indifference research as well.

The fourth and final point I wish to make about preemptive umbrella consent is that, while it may carry with it a host of contingent problems, those problems speak, not against the theoretical permissibility of using such consent, but rather against its cost-effectiveness. I will give two examples of this, though others also exist. First, if consent is too preemptive — i.e. it occurs too far in advance of its intervention — the chance of forgetting and of changing one's mind increases. This does not imply that, therefore,

preemptive consent is always illegitimate. Rather, it implies that if preemptive consent is used, sometimes measures should be taken to remind and to double-check for constancy of opinion. Of course, it might turn out that in cases of clinical indifference research it is more costly to implement such reminders and double-checks than it is merely to disclose everything at the time of each particular instance of clinical indifference research. But, obviously, that would not be to argue that it would be morally impermissible to conduct clinical indifference research on the basis of preemptive umbrella consent, only that it would be more costly to do so.[30]

Second, it is possible that some subjects in a clinical indifference study will have to be withdrawn from the study, perhaps because unforeseen side effects dictate a change to the other treatment, or perhaps because as the study evolves evidence begins to accrue that suggests that one treatment is superior to the other. But, as before, that does not imply that preemptive, umbrella consent is never legitimate, even if we assume that the researcher then acquires an obligation to debrief the subject when that subject is withdrawn. Rather, it implies only that, when using preemptive umbrella consent, the researcher makes a cost-benefit gamble between the efficiency gains of employing preemptive umbrella consent and the risk of incurring the extra cost involved in debriefing unwitting subjects who need to be withdrawn from the study.

These remarks have been brief, and the extent to which preemptive, umbrella consent can be valid deserves a fuller treatment than I can give it here. Still, I conclude *pro tanto* that it is morally permissible to enrol subjects into clinical indifference research on the basis of preemptive, umbrella consent. First, we do it all the time anyway; second, it is restricted to clinical indifference research; and third, the problems incurred by preemptive, umbrella consent are contingent and speak not to whether such consent can ever be valid but rather to whether in fact its use, when coupled with necessary safeguards, would be cost-effective.

7. Conclusion

I began by describing a puzzling discrepancy between what clinicians are obligated to disclose in order for consent to be valid, when they treat patients with one of two arbitrarily chosen treatments in clinical indifference, and what researchers are obligated to disclose, when they test subjects with one of the same two clinically indifferent treatments, again arbitrarily chosen. I argued that the best way to resolve this puzzle is by appealing to a pragmatic criterion for disclosure — only information which is material to the subject's decision must be disclosed, in order that the subject's consent be valid. I then applied that pragmatic criterion to the puzzling discrepancy. I concluded that, though prospective subjects may legitimately want to dissent from the mere fact of research, it is morally permissible to use preemptive umbrella consent to give them this opportunity.

As I said at the outset, the point of my detailed examination of the puzzling discrepancy is not that I think it is a pervasive and widespread problem afflicting society today. Rather, I think we can use the puzzle to examine and get clear on more theoretical issues surrounding informed consent. In particular, I have argued that the resolution of the puzzle requires adopting a pragmatic criterion for disclosure; it forces us to face up to the fact that valid consent is always, at least to some degree, preemptive and umbrella; and

it raises the further question, which I have tried to articulate and sharpen but not answer, of exactly how preemptive and how umbrella morally valid consent can be.

Eric Chwang, Department of Philosophy, Department of Philosophy, University of Colorado at Boulder, Boulder CO 80309-0232, USA. eric.chwang@colorado.edu

Acknowledgments

Thanks to audiences at the NIH Department of Bioethics and the University of Colorado at Boulder's Center for Values for their helpful feedback on presentations of this material. Special thanks to Dave Wendler and Frank G. Miller for helpful discussion and written comments on multiple earlier drafts of this essay.

NOTES

1 S. Joffe, E. Cook, P. Cleary, J. Clark & J. Weeks, 'Quality of informed consent in cancer clinical trials: A cross-sectional survey', *Lancet* 358 (2001): 1772–1777; J. Griffin, J. Struve, D. Collins, An Liu, D. Nelson & H. Bloomfield, 'Long term clinical trials: How much information do participants retain from the informed consent process?', *Contemporary Clinical Trials* 27 (2006): 441–448.

2 R. D. Truog, W. Robinson & A. Randolph, 'Is informed consent always necessary for randomized controlled trials?', *New England Journal of Medicine* 340 (1999): 804–807.

3 R. W. Smithells, 'Iatrogenic hazards and their effects', *Postgraduate Medical Journal* 51 (1975): 39–52. N. Fost, 'Ethical dilemmas in medical innovation and research: Distinguishing experimentation from practice', *Seminars in Perinatology* 22 (1998): 223–232.

4 Cf. the correspondence on the topic published not five months later, in *New England Journal of Medicine* 341 (1999): 448–450.

5 F. G. Miller & R. M. Veatch, 'Symposium on equipoise and the ethics of clinical trials', *Journal of Medicine and Philosophy* 32 (2007): 77–78. Indeed, that entire issue of the *Journal of Medicine and Philosophy* is devoted to philosophical problems surrounding the (mis-)use of equipoise in medical research. Note that clinical indifference is also not the same as normative equipoise, as described in B. A. Brody, L. B. McCullough & R. R. Sharp, 'Consensus and controversy in clinical research ethics', *JAMA* 294 (2005): 1411–1414. The distinction between equipoise and clinical indifference does not reduce to the distinction between clinical practices and available evidence. Even if equipoise is defined in terms of available evidence, the point I raise in the text remains.

6 For more on the distinction between something being (morally) wrong and its being (morally) impermissible, see J. Waldron, 'A right to do wrong', *Ethics* 92 (1981): 21–39.

7 Code of Federal Regulations Title 45 Part 46.116; 2003.

8 *Canterbury v. Spence* 464 F.2d 772 (D.C. Cir. 1972).

9 See, for example, R. R. Faden & T. L. Beauchamp, *A History and Theory of Informed Consent* (New York: Oxford University Press, 1986), pp. 133–138.

10 J. W. Berg, P. Applebaum, C. Lidz & L. Parker, *Informed Consent: Legal Theory and Clinical Practice*. 2nd edn. (New York: Oxford University Press, 2001), pp. 48–51. There is a nice appendix classifying each state in J. S. King & B. W. Moulton, 'Rethinking informed consent: The case for shared medical decision-making', *American Journal of Law and Medicine* 32 (2006): 429–501. (King and Moulton's article classifies 23 states and the District of Columbia as adopting a patient-centered standard, 25 as adopting a physician-centered standard, and New Mexico and Minnesota as having hybrids.)

11 J. R. Dillon, 'Informed consent and the disclosure of risks of treatment: the Supreme Court of Canada decides', *Bioethics Quarterly* 3 (1981): 156–62.

12 L. Skene, 'Informed consent: Lessons from Australia', *British Medical Journal* 324 (2002): 39–41.

13 On surgeon disclosure see, for example, S. Clarke & J. Oakley, 'Informed consent and surgeons' performance', *Journal of Medicine and Philosophy* 29 (2004): 11–35 and I. Burger, K. Schill & S. Goodman,

'Disclosure of individual surgeon's performance rates during informed consent: Ethical and epistemological considerations', *Annals of Surgery* 245 (2007): 507–513.

14 D. J. Mazur, 'Influence of the law on risk and informed consent', *British Medical Journal* 327 (2003): 732.

15 A 'hypothetical reasonable person' standard is often called *objective*, and an 'actual patient' standard is often called *subjective*. (Cf. King & Moulton op cit.) I think that terminology harbors the potential for confusion, so I avoid it.

16 The decision between actual patients and hypothetical reasonable ones is tricky and often conflated with the distinct issue of whether disclosure should be linked to materiality (the pragmatic criterion). For example, King and Moulton's critiques of patient-centered standards, of which the pragmatic criterion is but one variant, all reduce to the question of whether it is preferable to use actual patients or hypothetical reasonable ones, but they go on to abandon all patient-centered standards in general, including the pragmatic variant which emphasizes the role of materiality. (See King & Moulton op. cit., pp. 458–459.) Unfortunately I do not have the space to examine their arguments in detail here.

17 *Canterbury v. Spence* 464 F.2d 772 (D.C. Cir. 1972).

18 Of course, the two standards will coincide when the patient is the sort of person who would dissent whenever she is not told everything she wants to know. But, first, not all patients are like this, and, second, this seems to point in favor of a hypothetical, as opposed to actual, patient standard. Also, not all commentators make the mistake I attributed to Mazur above. For example, in a section of their classic *A History and Theory of Informed Consent* entitled 'The Reasonable Person Standard', Faden and Beauchamp write, 'The legal litmus test under this [reasonable person] standard for determining the extent of disclosure is the "materiality," or significance of information to the decision-making process of the patient . . . The reasonable person standard requires a physician to divulge any fact that is material to a reasonable person's decision . . . ': R. R. Faden & T. L. Beauchamp, *A History and Theory of Informed Consent* (New York: Oxford University Press, 1986), p. 32.

19 Two candidate epicycles are that what was really meant all along was either (1) 'what a hypothetical reasonable, *knowledgeable, and prudent* person would want to know' or (2) 'what the patient would have agreed in *hindsight* that she wanted to know, after it was disclosed.'

20 Of course, the clinician in charge of my care might be particularly well placed to deliver various pieces of information, and, because she is the leader of my health-care team, she might even have a special (role-related) obligation not shared by others to be the conduit for that information. Indeed, others in the know may have a duty to refrain from disclosing certain information to me, such as a grim diagnosis.

21 Indeed, these reasons to disclose have nothing to do with the fact that the amputation is a clinical intervention. If I faced an incurable medical condition where my leg would fall off on its own, the clinician treating me would be seriously wrong to refrain from telling me this.

22 It is plausible, though, that the third debate — between materiality and 'what the patient would want to know' — depends on a patient-centered approach for its legitimacy. Also keep in mind that the debate between actual versus hypothetical patients may admit of an intermediate position, 'what actual patients would want if they are reasonable', and that the debate can have different answers depending on whether we are discussing legal issues or moral issues. Plausibly, we have moral obligations to respect what actual patients want (or perhaps 'actual patients when they are being reasonable'), though for reasons of verifiability and the minimization of frivolous lawsuits we prefer hypothetical patients for legal standards.

23 Those who object to the claim that the clinician's arbitrary choice is thereby random are invited to imagine the scenario unfolding as follows. There are two emergency rooms in the hospital, equivalent for all practical purposes. In preparation for the study, the researcher has stocked one emergency room only with normal saline, the other only with lactated ringer's. (If necessary, the reader is invited to invent a back-story according to which it was already accepted practice at this hospital to stock the two emergency rooms selectively in this way, before this particular study was even conceived.) New patients are assigned to the two emergency rooms at random, say by alternating. Then whether the patient dissents to participate in the study makes no difference to what fluid the patient gets, and dissenters are still assigned their re-hydration fluid at random.

24 Cf. F. G. Miller, 'Revisiting the Belmont Report: The ethical significance of the distinction between clinical research and medical care', American Philosophical Association, *Newsletter on Medicine and Philosophy* 5 (2006): 10–14.

25 Truog *et al.* op. cit., p. 805.

26 Careful readers will note that I said 'only if', not 'if and only if'.

27 No matter how finely one describes a cholecystectomy procedure, for example, the description still leaves open an infinite number of ways it might be performed, including different locations, speeds, and approach angles of the surgeon's hand when making the first incision. Presumably the surgeon gets consent for all these sorts of details merely by asking for (preemptive) umbrella consent to '(laparoscopic) cholecystectomy'.

28 This point strikes me as obvious, and I am not the first person to make it. See, for example, N. C. Manson & O. O'Neill, *Rethinking Informed Consent in Bioethics* (Cambridge: Cambridge University Press, 2007), p. 12 and D. Beyleveld & R. Brownsword, *Consent in the Law* (Portland, OR: Hart Publishing, 2007), pp. 212–221 for more discussion on umbrella consent.

29 See, for example, D. Shickle, 'The consent problem within DNA biobanks', *Studies in History and Philosophy of Biological and Biomedical Sciences* 37 (2006): 503–519.

30 Of course, I acknowledge the indirect, consequentialist argument that we ought to allocate our resources as efficiently as possible.