

Reasonable risks in clinical research: A critique and a proposal for the Integrative Approach

Alex John London^{*,†,‡}

Department of Philosophy, Carnegie Mellon University, Pittsburgh, PA 15213-3890, U.S.A.

SUMMARY

Before participants can be enrolled in a clinical trial, an institutional review board (IRB) must determine that the risks that the research poses to participants are 'reasonable.' This paper examines the two dominant frameworks for assessing research risks and argues that each approach suffers from significant shortcomings. It then considers what issues must be addressed in order to construct a framework for risk assessment that (a) is grounded in a compelling normative foundation and (b) might provide more operationally precise guidance to the deliberations of various stakeholders. The paper concludes by sketching the outlines of what is referred to as the 'Integrative Approach' to risk assessment and by highlighting some of the ways in which this approach may be more promising than current alternatives. Copyright © 2006 John Wiley & Sons, Ltd.

KEY WORDS: research risk; risk assessment; reasonable risk; equipoise; Common Rule; research ethics

INTRODUCTION

A vital element in the ethical review of research involving human subjects is the assessment of the risks that such research poses to study participants. In the U.S., for example, the federal regulations that govern human-subjects research—known as the 'Common Rule'—charge institutional review boards (IRBs) with ensuring first that 'risks to subjects are minimized,' and second with ensuring that 'risks to subjects are reasonable in relation to anticipated benefits...' [1]. These requirements are prior to the other requirements listed, such as informed consent and fair subject selection, not simply in their order of exposition in the Common Rule, but in the sense that if an IRB finds that

*Correspondence to: Alex John London, Department of Philosophy, Carnegie Mellon University, Pittsburgh, PA 15213-3890, U.S.A.

†E-mail: ajlondon@andrew.cmu.edu

‡Associate Professor of Philosophy.

Contract/grant sponsor: Andrew W. Mellon Foundation

Copyright © 2006 John Wiley & Sons, Ltd.

Received 16 May 2006

Accepted 22 May 2006

the risks associated with a particular study are not reasonable then it is unethical even to offer participation to potential subjects.

Because ethical frameworks are needed most in precisely those cases where the intuitions of reasonable people conflict or diverge, it is imperative that IRB members, as well as other important stakeholders—such as researchers, research sponsors, and community members more generally—have a framework for evaluating research risks that will facilitate the principled resolution of difficult cases. In the first section below I briefly describe three criteria that should be used to evaluate different frameworks for assessing research risks. In the next two sections I argue that the two most prominent such frameworks suffer from fundamental deficiencies on one or more of these dimensions. Together, the arguments in these sections illustrate the need to move beyond, or to fundamentally rethink, the current orthodox approaches.

In an effort to facilitate this process, the final sections of this paper discuss the issues that must be addressed in order to construct a more adequate framework. In the hope of advancing the state of debate I then outline very briefly what I refer to as the Integrative Approach to clinical research. Although important details of this outline require greater elaboration and specification, one virtue of this framework is that it frames the central issues in ways that make it possible to further refine and specify their operational content through the use of methods from decision theory and the social sciences. Although the Integrative Approach is presented here in outline, I hope that the need for such an approach will be sufficiently clear, and the contours of this alternative sufficiently compelling, that statisticians, decision scientists, ethicists, and other stakeholders will be motivated to take up this project of refinement and extension in earnest.

CRITERIA FOR EVALUATING FRAMEWORKS FOR RISK ASSESSMENT IN CLINICAL RESEARCH

For the purposes of the present inquiry there are three main criteria that should be used to evaluate competing frameworks for risk assessment in clinical research. The first is *conceptual and operational clarity*. As a branch of practical ethics, the guiding ambition of research ethics is to provide a framework for evaluating clinical research that can help stakeholders resolve reasonable disagreements in a way that is publicly accessible and defensible. As such, an acceptable framework should (1) provide a clear definition of ‘reasonable risk’ and an account of the normative justification for this definition. In addition, it should (2) elucidate a set of operational criteria or markers that delineate in an operationally useful way the parameters or boundaries that separate reasonable from excessive risks. Finally, it should (3) articulate practical tests that deliberators can use in order to determine whether or not these operational criteria have been met in any particular case [2, 3].

The second criterion is *an appropriate balance of relevant moral concerns*. Assessments of the reasonableness of risks in clinical research are necessary to establish whether or not it is morally permissible to offer participation to potential participants. An acceptable framework should therefore draw an appropriate balance between concern for the welfare of individual trial participants and respect for the capacity of competent decision makers, when adequately free and informed, to determine what activities are worth pursuing in light of their own values and commitments. This balance between what are sometimes called protectionist concerns and respect for autonomy should in turn reflect an appropriate resolution to the tension inherent in clinical research between advancing science for the good of future patients and the interests of present research participants [4].

The third and final criterion is *theoretical unity in scope of applicability*. That is, an acceptable framework for risk assessment should be general enough to provide guidance about the full range of cases that occur in clinical research. In addition, when all else is equal, frameworks that achieve a greater range of applicability without recourse to *ad hoc* ancillary principles should be preferred to frameworks that require the addition of such principles.

SHORTCOMINGS OF THE COMMON RULE APPROACH

The first major framework for evaluating the risk profile of research involving human subjects is described by the National Commission for the Protection of Human Subjects of Biomedical and Behavioral Research in the *Belmont Report* and codified in the Common Rule [1, 5]. I will therefore refer to it as the 'Common Rule Approach.' Recently, however, this approach has been defended in the research ethics literature as an important element of what are referred to, respectively, as the 'Non-Exploitation' [6, 7] approach and 'Component Analysis' [8]. The Common Rule Approach therefore represents an area of overlap between what are in other respects competing and conflicting frameworks.

The Common Rule Approach holds that risks are reasonable when they are offset or outweighed by the anticipated benefits of the research. It thus straightforwardly embraces the language of risk-benefit analysis. However, the Common Rule enumerates two types of benefit that might offset risks to study participants. It states that IRBs must be satisfied that 'risks are reasonable (a) in relation to anticipated benefits, if any, to subjects, and (b) the importance of the knowledge that may reasonably be expected to result' (bracketed letters added). The first scenario reflects a decision problem that is common in the clinical context. That is, patients and their physicians must often decide whether the risks associated with a particular treatment are outweighed by the possible benefits to that same individual. As difficult as these emotionally charged decisions may be in practice, however, there are well-known techniques that can be used to support such decisions, at least in principle. According to the Common Rule Approach, however, risks to individual research subjects may also be outweighed or justified by the potential benefits that may accrue to other individuals or to the community in general through the advancement of science.

The normative foundation for this view is stated succinctly in the *Belmont Report* [5]. It is that the moral value of beneficence gives rise to a general duty that can be expressed by the rule 'maximize possible benefits and minimize possible harms' [5, at B.2]. However, relevant to this consideration are risks and benefits that accrue to individual research subjects as well as risks and benefits that accrue to society more broadly. As such, 'beneficence thus requires that we protect against risk of harm to subjects and also that we be concerned about the loss of the substantial benefits that might be gained from research' [5, at C.2]. This approach therefore has a clear normative foundation. Furthermore, it is perfectly general in its scope of applicability since for every case of clinical research the salient question is whether the risks that are not outweighed by benefits to the individual participant are outweighed by the value of the information to be generated.

There are, however, two fundamental and interrelated deficiencies with this approach. The first relates to its failure to provide any substantive guidance about how to balance relevant moral concerns in this domain. This is closely related to the second deficiency, namely, its lack conceptual and operational clarity.

To begin with, the Common Rule Approach sets up a decision problem in which the interests of individual research subjects have to be weighed or traded off against benefits to science without

providing any practical guidance about how such trade-offs ought to be made. We are told clearly that the interests of individual participants can be outweighed by the prospect of sufficiently important benefits to science. But there is no guidance about (1) how to represent or quantify the ‘importance of the knowledge that may reasonably be expected to result,’ (2) how to measure that value, so defined, in concrete cases or (3) how to weigh or trade-off the value of that information against the various interests of individual trial participants [2, 3, 9]. Moreover, it is not that there is no way to specify these different variables. It is, rather, that the ways of doing so are as boundless as the imaginations of different decision makers [10, 11].

Proponents of this approach have dismissed these worries as a misguided desire for a kind of precision that cannot be had. Such issues ultimately boil down to matters of ‘judgment,’ they claim, and cannot be quantified [6, 7, 12]. But if risks to individual welfare really are *justified* by offsetting gains in social value, then we need some general criteria for establishing when this condition has been met, as well as some interpersonally useful decision rule or heuristic for helping members of a pluralistic, liberal democratic community to arrive at reasoned judgments about when these criteria have been met. We therefore need a reasoned account of the issues outlined in the previous paragraph in order to know how to apply the proposed framework.

If the central variables in this framework really cannot be measured and the associated trade-offs really cannot be specified in a meaningful way, this approach amounts to setting up a decision problem and then exhorting deliberators to use their intuitions in order to decide whether the risks to subjects seem ‘worth it.’ This, however, leaves research ethics without a means of providing reasoned guidance to deliberators in exactly those cases where it is needed most, cases where intuitions are unclear or conflict. It also permits a wide degree of variability in judgments about the reasonableness of risk across different IRBs. Together, these factors may generate the appearance, if not the reality, of a degree of arbitrariness that may compromise the social trust of the various stakeholders whose interests are supposed to be served and whose judgments are supposed to be assisted by a framework for moral decision making in this domain.

SHORTCOMINGS OF THE DUTY OF PERSONAL CARE APPROACH

The second major approach attempts to provide greater conceptual and operational clarity to the evaluation of risks in clinical research by using the professional obligations of physicians to define the boundaries or parameters on the risks to which it is permissible to expose research participants. This approach is defined by a core set of normative assumptions that mark out as salient a set of variables that variants of this approach, such as frameworks that embrace either the ‘uncertainty principle’ or ‘clinical equipoise,’ then deal with in slightly different ways [11].

The first normative assumption that defines this approach is that researchers, like clinicians, owe a Duty of Personal Care [13], sometimes referred to as a therapeutic obligation [14] or a fiduciary obligation [15], to their patients and that the content of this duty or obligation provides the standard that should be used to determine whether it is permissible to enter a participant into a particular clinical trial. Second, the Duty of Personal Care is an affirmative duty that requires physicians to advance the interests of individual patients to the best of their ability. Consider, then, a set of individuals $I = \{1, \dots, i, \dots, n\}$ with a particular medical condition for which there is a set of available treatment options $S = \{s_1, \dots, s_m\}$, some subset U of which may be offered within a particular clinical trial. On this view, it is permissible to allow a patient i to enter a clinical trial

in which she will be assigned to an element of U at random only if the relative therapeutic value of the elements of S for i are uncertain.

Variants of this approach can then be distinguished by their treatment of three variables that I purposefully glossed over in the previous paragraph: (1) whose uncertainty is morally relevant, (2) how to specify the epistemic threshold for uncertainty and (3) whether the relative therapeutic value of an intervention should be treated as a one-dimensional measure of a single outcome, or as a multi-dimensional measure that takes account of several possible outcomes [11]. For the sake of brevity I briefly describe two from among the many ways of dealing with the first two of these variables and then raise some objections to this general approach.

For example, as formulated by Peto and Baigent, the uncertainty principle requires that:

A patient can be entered [into a trial] if, and only if, the responsible clinician is substantially uncertain which of the trial treatments would be most appropriate for that particular patient. A patient should not be entered if the responsible clinician or the patient are for any medical or non-medical reasons reasonably certain that one of the treatments that might be allocated would be inappropriate for this particular individual (in comparison with either no treatment or some other treatment that could be offered to the patient in or outside the trial) [16].

They therefore require both that the 'responsible clinician' and the individual patient be uncertain about the appropriateness of each treatments in S for that patient.

In contrast, what Benjamin Freedman calls clinical equipoise locates the relevant uncertainty, not in the mind of the individual clinician, but in the larger medical community [17, 18]. In fact, Freedman explicitly states that clinical equipoise can exist when there is 'a split in the clinical community, with some clinicians favouring A and others favouring B' and that clinical equipoise is 'consistent with a decided treatment preference on the part of the investigators. They simply recognize that their less favoured treatment is preferred by colleagues whom they consider to be responsible and competent' [17, at 144].

Similarly, different versions of this approach can be distinguished by the threshold for uncertainty that they adopt. For example, one proponent of the uncertainty principle argues that:

An ethical physician must do what is best for his or her patients. She cannot participate in a controlled trial if she is certain that one arm is superior to the others and that some of her patients will receive an inferior treatment by participating in the trial. It does not matter whether her certainty is based on formal scientific studies, on personal experience, on anecdote, on tacit understanding, or rules of thumb. Whether her certainty is in accord with or diverges from the view of the medical community is irrelevant. Uncertainty is a moral prerequisite for a controlled study. If we know what we should do, we should do it, not study it [19, at 758].

Here, the focus on uncertainty in the mind of the individual clinician is combined with a relatively fragile epistemic threshold according to which uncertainty is the absence of even anecdotal reason to expect that at least one the elements of U is inferior to some other member of S .

In contrast, Freedman claimed that 'clinical equipoise' embodies a more robust epistemic threshold according to which the relevant uncertainty persists until evidence for the superiority of one intervention emerges that would be sufficient to forge a consensus in the relevant expert clinical community [17, 18]. This threshold requires that the evidence supporting a claim to superiority on behalf of one element of S must be sufficiently compelling that it will influence the practice behaviour, not just of one physician, but of the community of physicians [15].

Although the Duty of Personal Care Approach admits of more varieties than I have described here, the above illustrations are sufficient to highlight three important shortcomings of this general approach. The first relates to the criterion of theoretical unity in scope of applicability. That is, it is unclear how to apply this framework to the full range of research cases since this approach is grounded in the physician's Duty of Personal Care. However, some research involves healthy volunteers who are asked to accept affirmative risks to their health without the prospect of individual therapeutic benefit. There are two ways in which this problem might be addressed. One gives rise to the second shortcoming of this approach, namely, a dilemma for the Duty of Personal Care Approach regarding the way it balances concern for participant welfare against respect for participant autonomy. The other highlights the third shortcoming regarding the conceptual unity of this approach in its scope of applicability.

One way to ensure that the Duty of Personal Care Approach has an adequately general scope of applicability is to stick with the foundational idea that the parameters on the risks to which it is permissible to expose participants in clinical research should be set by the content of the Duty of Personal Care and to claim that, since studies involving healthy volunteers represent research, they fall under the ambit of this duty. We then need to know how this duty should be understood in this context.

If we assume that the Duty of Personal Care binds physicians to maximize the health interests of their patients, independent of the way in which those interests fit into the patient's larger plans and commitments, then this would rule out such research as unethical. For healthy participants, that is, the option of participating in research is always dominated by option of not participating. This interpretation of the Duty of Personal Care provides clear guidance about these cases, but the guidance seems to be objectionably paternalistic. In particular, it would prohibit participants from foregoing even minor health benefits for the sake of advancing science. It is also at odds with the way the Duty of Personal Care is understood in clinical practice.

If the Duty of Personal Care is understood less paternalistically, as a duty to safeguard and advance the medical best interests of patients in a way that is mediated by the patient's reflective understanding of those interests in light of his or her broader goals and commitments, then the substantive content of this framework becomes much less clear. In clinical practice physicians routinely assist their patients in medical procedures that pose affirmative risks to the patient that are not compensated for by direct benefits to that patient, in order to benefit others. For example, large teams of physicians routinely assist otherwise healthy patients in donating a lobe of their liver or one of their kidneys to other individuals. If these activities are consistent with the Duty of Personal Care, then perhaps the participation of individuals in clinical research should be understood analogously. After all, research participants should be free to knowingly and voluntarily forego some health benefits for the sake of scientific advancement. But now the key question is whether the Duty of Personal Care provides independent criteria for determining what kinds of health benefits it is permissible to ask subjects to voluntarily forego. As I suggest below, the question of what burdens professionals, be they physicians or researchers, can assist individuals in accepting in order to help others is itself a broader social question about the proper role of medical professionals, and not necessarily one that can be resolved by appeal to norms that are internal to those professions.

Another possibility, therefore, is to claim that the Duty of Personal Care Approach does not apply to all forms of clinical research since activities involving healthy volunteers fall outside the scope of the therapeutic obligation. This approach thus has a restricted scope of applicability and needs to be augmented by ancillary principles. This appears to be the view that is adopted by the proponents of 'Component Analysis' [8]. On their view, IRBs should first distinguish elements of a clinical

REASONABLE RISKS IN CLINICAL RESEARCH

trial that are administered to participants because of their possible therapeutic value from elements whose warrant for inclusion derives from purely research-related grounds. IRBs must then ensure that equipoise exists between the former, therapeutic elements, and any other relevant alternative therapeutic modalities. They then apply the Common Rule Approach to determine whether the risks associated with the remaining, purely research-related elements of the trial are reasonable in relationship to the value of the information the trial is designed to generate.

Although this option broadens the scope of applicability of what I am calling the Duty of Personal Care Approach, it does so at the expense of the conceptual and operational clarity of the resulting framework. In particular, this option inherits the shortcomings of the Common Rule Approach discussed in the previous section. In addition, however, restricting the applicability of the Duty of Personal Care does not obviate the need to address the dilemma that I just sketched in those cases where the duty applies.

This discussion of what I am calling the two dominant frameworks for evaluating the reasonableness of risks in clinical research has been highly condensed. I hope, nevertheless, that it is sufficient to illustrate the seriousness of the problems they face and to motivate the search for an alternative.

MOVING FORWARD: SOME MAJOR ISSUES

Before a single participant can enrol in a clinical trial an IRB, and sometimes several IRBs, must find that the risks associated with that trial are reasonable. I have argued above that the clearest aspect of the Common Rule Approach is that it permits the interests of present research participants to be traded off for gains in information and understanding. I have also argued that this approach gives IRB members, researchers, and other stakeholders no substantive guidance about how to: (a) distinguish the relative value of the various interests of research participants that may be at stake in a particular trial, (b) distinguish the relative value of the information or understanding that might be generated by trials that study different questions and employ different methods and (c) distinguish permissible from impermissible trade-offs between these variables.

At some point every framework for decision making must rely on the common sense and good judgment of those who apply it; no framework can apply itself, after all. For some approaches, however, this point comes sooner than others. The challenge for any framework in this area, therefore, is to facilitate reasoned decision making by first setting out a decision problem that has a determinate and interpersonally articulable structure and then by articulating a shared set of criteria or standards that deliberators can use to demarcate reasonable from unreasonable risks.

In this regard, the Duty of Personal Care Approach does a better job of giving a determinate structure to the problem that IRB members and other stakeholders face in assessing the risks associated with clinical research. That is, it directs deliberators to ascertain what the available treatment options are, to assess their relative therapeutic merits, and then to ensure that research participants are only allocated to treatment options that do not violate the Duty of Personal Care. Although this gives the decision problem a clearer structure, I argued above that this approach faces important difficulties in formulating the Duty of Personal Care in a way that preserves the determinacy of its practical guidance without falling into unjustified paternalism.

Neither the Common Rule Approach, nor the Duty of Personal Care Approach *directly* distinguishes the relative importance of the various interests of research participants that may be implicated in a particular trial or explains directly how those differences in importance should

be reflected in the evaluation of risk. Rather, they address this issue at best only *indirectly* by relying on the judgments of IRB members or of expert clinicians.

I believe that in order to overcome the deficiencies of the above frameworks it will be necessary to address directly fundamental questions concerning the relative importance of the various interests of research participants and how such differences ought to be reflected in the evaluation of risk in clinical research. Providing a clear and normatively defensible account of this issue is necessary in order to ensure that a framework for risk assessment adequately balances the relevant moral concerns in this domain. Any such distinction, therefore, must be drawn in such a way as to facilitate the conduct of sound science without allowing respect for the fundamental interests of individual research participants to be outweighed or overridden by the potentially unlimited number of future patients who stand to benefit from successful clinical research.

Similarly, any such distinction should help to explain, in publicly accessible terms, the most basic intuitions of IRB members or expert clinicians that the Common Rule Approach and the Duty of Personal Care Approach rely on so heavily. Finally, if possible, these distinctions should cohere as much as possible with policies in other social arenas where the autonomy of individuals is limited or restricted out of a concern for their basic welfare.

Because the above frameworks deal at best only indirectly with these issues, they may appear to be less controversial than any approach that attempts to address them directly. For this reason, no matter how these distinctions are drawn, they will need to emanate from a clear normative foundation that is capable of garnering widespread public support within a pluralistic, liberal democratic community.

In the following section, I outline in very general terms a distinction between two sets of interests that all individuals share. I also describe briefly how this distinction can be derived in a way that makes it particularly well suited for making social evaluations of risk within a pluralistic, liberal democratic community. To be clear, because I believe that some distinction of this kind is necessary in order to advance the debate in this area, I think that this proposal has important merits in its own right.

For my present purposes, however, I am interested in illustrating how this distinction might provide a more attractive normative foundation for a framework of risk assessment in clinical research that has significant practical content. I therefore outline very briefly what I refer to as the Integrative Approach, and I argue that it is able to capture what is most compelling about the Duty of Personal Care Approach in the area of clinical research to which the latter applies most directly, namely, those cases in which research participants suffer from a significant sickness, injury, or disease. Although this discussion will of necessity be fairly brief, I hope that it will be sufficient to motivate interest in clarifying and refining the Integrative Approach.

I conclude by making a tentative proposal about how this framework might be extended to cover a set of issues that constitute an open problem for research ethics. This problem is how to evaluate the risks associated with purely research-related elements of a clinical trial, where this includes the risks to healthy individuals who participate in clinical research.

THE INTEGRATIVE APPROACH: A TENTATIVE PROPOSAL

At a very basic level, assessments of risk involve two components. One is an estimation of the likelihood or probability that a set of events will occur. The second is a valuation of the outcomes that are associated with those events. How an individual values a set of outcomes depends a great deal on which interests of that individual are at stake and how the agent values those interests.

When individuals value most highly different of their interests they may have very different attitudes toward the same action or activity. In light of these differences, some individuals may consent to participate in a particular trial while others do not.

IRBs and other stakeholders, however, must evaluate the risks associated with a clinical trial before *any* individual has been offered participation. They are asked, in effect, to make a social valuation of the risks associated with a particular clinical trial as part of an oversight process that operates as a social filter on which trials are allowed to recruit participants. A central challenge, therefore, is to articulate a *social perspective* from which the relative value of the various interests of agents can be assessed in order to give more determinate content to the social evaluation of research risk.

For this purpose I want to distinguish two sets of interests that all persons possess [20]. What I will call ‘personal interests’ are interests that individuals have in light of their personal conception of the good and the particular life plans and projects that they adopt in pursuit of this. Some individuals, for example, value intense physical challenge, and therefore pursue activities such as mountain climbing or skydiving or undertake careers in which they routinely encounter personal danger. Others, in contrast, view these activities as having little value, and even as a little crazy, and incline instead toward intellectual stimulation or find meaning in intense personal relationships and commitments.

In liberal democratic communities individuals may differ radically in their personal interests and these differences are a common source of conflict in social decision making and public policy. Insofar as they disagree in their personal interests, individuals may have little common ground for making social decisions because they value very different things—where one values the thrill of physical challenge, for example, the other values the virtues of contemplation and intellectual achievement. Nevertheless, this conflict masks a higher-order level of agreement. That is, each of these individuals shares what the philosopher John Rawls calls a common, ‘higher-order’ interest in being able to pursue their individual personal interests, whatever they may be, without unwarranted interference [21]. This shared higher-order interest provides a kind of common ground that members of a pluralistic, liberal democratic community can use for making social judgments about the shape, scope, and limitations of the social structures that will impact their rights or welfare.

In particular, this shared higher-order interest in being free to advance their personal interests provides the basis for distinguishing and giving evaluative priority to what I will call ‘basic interests.’ These are interests that each individual has solely in virtue of having this higher-order interest in being able to pursue some determinate set of personal interests. Basic interests, therefore, are interests that agents have in being able to cultivate and to exercise those fundamental human capacities that are constitutive of what Rawls refers to as our two moral powers: the capacity to formulate and to pursue a life plan based on a conception of the good and to regulate our conduct with others on the basis of principles of right [21]. Such basic interests include developing and exercising one’s capacities for reflective thought and practical decision making, developing and cultivating one’s basic affective or emotional capacities and having the effective freedom to exercise those capacities in the pursuit of particular projects and meaningful social relationships.

How might such a distinction be relevant to the assessment of risks in clinical research? What I refer to as the Integrative Approach uses this set of basic interests to define the ‘space of equality,’ the domain over which community members have a just claim to equal treatment. This is reflected in the following definition of reasonable risk:

Definition of reasonable risk. Risks to individual research participants are reasonable just in case they (1) require the least amount of intrusion into the interests of participants that is necessary

in order to facilitate sound scientific inquiry and (2) are consistent with an equal regard for the basic interests of study participants and the members of the larger community whose interests that research is intended to serve.

The first part of this definition quantifies over all of the interests of participants, both basic and personal. Under this requirement fall two important considerations. First, in order to constitute sound scientific inquiry, a particular trial must address a research question that is appropriate in light of the current state of scientific inquiry. Second, once this condition has been met, in order to count as reasonable, risks must be reduced to those that are necessary to address this question in a rigorous manner. These requirements simply reflect the idea that it is unreasonable to subject research participants to risks that are gratuitous or unnecessary for the conduct sound scientific inquiry.

The second part of this definition quantifies over the basic interests of participants. This focus on basic interests represents the normative claim that it is permissible to ask individual community members to alter, risk, or even to sacrifice some of their *personal interests* as part of an effort to advance or secure the *basic interests* of others. This means that it is permissible, on this view, to ask individuals whether they are willing to take up, as a personal project, participation in a research initiative that might involve the dedication of time and energy that might otherwise have been dedicated to other projects or goals. It is also permissible to ask them to endure experiences that are unpleasant, painful, or otherwise difficult. The constraint, however, is that such participation must not pose risks to the basic interests of participants that are inconsistent with the same degree of concern that is shown for the basic interests of other community members. This last claim will be explicated by the two operational criteria below.

This social judgment thus reflects the way that the Integrative Approach reconciles the potentially competing interests of future patients and current research participants. In particular, the Integrative Approach recognizes an affirmative obligation to ensure that clinical research is organized in such a way as to generate the means to help safeguard and advance the basic interests of future patients. But it does not aggregate the basic interests of future persons and then weigh that aggregate against the basic interests of present research participants. Instead, this approach seeks to create the social conditions within which community members can take on, as a personal project, the goal of assisting future patients, while being assured that their basic interests are treated with the same moral concern as that which provides the moral motivation for the research endeavour itself.

Operational criteria for the application of this definition are generated by considering what conditions have to be met for clinical research to show equal regard for the basic interests of participants and non-participants. Like the Duty of Personal Care Approach, the Integrative Approach will deal differently with cases in which research participants themselves suffer from sicknesses, injuries, or diseases that threaten or restrict their basic interests. Here, however, the Integrative Approach makes clear what is addressed only indirectly by the Duty of Personal Care Approach. Namely, if researchers have a special obligation to provide a certain level of care to participants of this kind it is not because of a feature of the physician-researcher *per se*, such as his or her Hippocratic duty, but because restrictions on the basic interests of individuals compromise their ability to function in ways that all community members can recognize as fundamental to their ability to form or to pursue some meaningful set of personal interests. These restrictions place those who are afflicted with them at a disadvantage, not with respect to goals or ends that they happen to have insofar as they, for example, value being a member of a particular club or looking like a particular celebrity, but with respect to their basic ability to pursue their personal interests, whatever they are, on an

equal footing with others. As a result, the Integrative Approach makes clear the reasons why researchers who are not physicians, and who may not have the physician's Hippocratic duties, should nevertheless be bound by this requirement.

The Integrative Approach, therefore, adopts the following as the first of two criteria for equal regard:

First operational criterion. Equal regard for the basic interests of research participants and non-participants requires that when the basic interests of an individual participant are threatened or compromised by sickness, injury or disease, the basic interests of that individual must be protected and advanced in a way that does not fall below the threshold of competent medical care.

This operational criterion invokes the threshold of competent medical care, not as a source for the normativity of this framework, but as a standard for determining what is required in order to show equal regard for the basic interests of those whose basic interests are themselves threatened or restricted by sickness, injury, or disease. As the proponents of equipoise often note, competent medical care represents the socially enforceable standard of professional knowledge, skill, and ability that community members have a legitimate claim to receive when they access the medical system [15]. Although some clinicians may rise above the rest in terms of various professional excellences, competent medical care denotes the level of care that the medical profession is accountable for providing on a uniform basis. Competent medical care thus refers to the use of practices, procedures, and methods that have a reasonable likelihood of success. In this respect, competent medical care serves as an indicator of what expert medical professionals believe is a causally efficacious means of effectuating desired clinical goals.

The following practical test can then be used to determine whether or not a particular clinical trial satisfies this operational criterion. Let $I = \{1, \dots, i, \dots, n\}$ be the set of individuals with a particular medical condition for which there is a set of available treatment options $S = \{s_1, \dots, s_m\}$. Let U_i be the set of interventions from S to which individual i might be allocated within a particular clinical trial and let U_i^* be the set of interventions from S that are admissible treatment options for the individual i .

Practical test for first criterion. A treatment s_j is admissible for individual i just in case there is either uncertainty among, or conflict between, expert clinicians about whether s_j is dominated by any other members of S as a treatment for individual i . For each individual in I , the care and protection afforded to that individual's basic interests falls within the threshold of competent medical care just in case each intervention in U_i is a member of U_i^* .

This condition is similar in important respects to the criteria of admissibility recommended by Kadane and colleagues [22]. In that trial, a treatment was deemed to be admissible for a participant just in case it was judged to be the best treatment option for that individual by at least one from among a number of expert clinicians. In this case, however, the expert clinicians were actually computer models that had been constructed out of a careful elicitation process involving real clinicians. The most salient difference between these two concepts of admissibility is that the Integrative Approach limits the scope of these judgments to the basic interests of study participants. Nevertheless, when this condition is met, each individual who participates in a clinical trial is assured of receiving a package of medical care that would be recommended for them by at least a reasonable minority of expert clinicians.

This practical test is also similar to Freedman's clinical equipoise, but here too there are some important differences. First, the moral force of this requirement is grounded, not in the individual

physician's therapeutic obligation, but in broader claims about the need for basic social structures to function so as to preserve and to advance the basic interests of all community members and about the role of clinical medicine as a social structure that fills this mission when the basic interests of community members are restricted by sickness, injury, or disease.

Second, this practical test explicitly distinguishes uncertainty in the mind of expert clinicians from a state of clinical conflict between such clinicians [11]. The former obtains when individual clinicians have no ground for preferring one treatment from S over any others as a treatment for i . This state might occur, for example, when a novel intervention begins to show sufficient promise in animal models and in early trials in humans that clinicians become uncertain about its net therapeutic advantage relative to existing interventions for some set of individuals. When this occurs it may be permissible to initiate a clinical trial in which individuals for whom both of these interventions are admissible are randomized to one of them.

Clinical conflict exists when expert clinicians have conflicting opinions about the superiority of one intervention over others. So, for example, imagine that one or more clinicians recommend treatment s_1 over the other available option s_2 for a particular patient i while one or more clinicians recommends treatment s_2 over s_1 for i . In this case, it would be permissible to offer individual i the option of being randomized to either s_1 or s_2 or to continue a trial in which this individual has been so randomized, because no matter what the result of the randomization, this patient is guaranteed to receive an intervention that would be recommended for her by at least a reasonable minority of expert clinicians.

As the critics of equipoise frequently point out, the state of uncertainty will likely be quite fragile and evanescent. As data emerge from a fixed size clinical trial, for example, the uncertainty of clinicians will likely be disturbed well before the trial reaches the predetermined level of statistical significance. Since this would frustrate the conduct of sound science, critics argue that the equipoise requirement must be rejected [14, 23, 24].

The Integrative Approach responds to this objection as follows [11, 25]. If the evidence generated by such a trial is sufficient to create consensus about the superiority of one option over the other, then the trial will have served its social function and ought to be terminated. However, if only some clinicians are convinced and a reasonable minority remain uncertain or would continue to make conflicting clinical recommendations, then we have moved from a state of uncertainty to one of clinical conflict. It is therefore permissible to continue the trial.

It is worth noting that the admissibility of allowing an individual to participate in a clinical trial in the face of definite but conflicting medical opinion cannot be accounted for by any version of the Duty of Personal Care Approach that takes the Duty of Personal Care to bind or vest in the individual physician. That is, when physicians have definite preferences about how to treat a particular patient, but these preferences conflict, no physician is uncertain about what is best for the individual in question. But, on at least most standard interpretations of the Duty of Personal Care, if the individual physician is not uncertain, then that physician's duty to the individual patient prohibits him from recommending trial participation as an admissible option. However, it is difficult to see the wisdom in forcing individual patients to choose at random between the conflicting recommendations of expert clinicians while prohibiting those same patients from being randomized to the interventions recommended by those clinicians.

Freedman himself seemed to recognize the possibility of clinical conflict when he claimed that clinical equipoise can exist when there is 'a split in the clinical community, with some clinicians favouring A and others favouring B' and that clinical equipoise is 'consistent with a decided treatment preference on the part of the investigators. They simply recognize that their less favoured

treatment is preferred by colleagues whom they consider to be responsible and competent' [17]. But this brief comment is not further elaborated and I have argued above that it is in tension with Freedman's claim that the equipoise requirement is grounded in the individual physician's Duty of Personal Care. In this respect, the Integrative Approach is able to capture the most appealing aspects of the Duty of Personal Care Approach while avoiding some of the fundamental shortcomings of that approach. This alone should be sufficient to motivate further interest in the Integrative Approach.

However, the first operational criterion and its associated practical test are not sufficient to cover all the relevant cases in clinical research. In particular, these requirements ensure that when individuals whose basic interests are threatened or restricted by sickness or disease participate in research they receive a level of care that is consistent with an equal regard for the basic interests of participants and non-participants alike. But the conduct of sound clinical research sometimes requires research participants to undergo tests or diagnostic interventions to which they would not be exposed if they simply sought treatment in a clinical setting. The most immediate effects of such procedures often fall on the personal interests of individuals in that they involve discomfort or some degree of bodily or personal intrusion. But even when this is the case, such procedures also pose some incremental risk to the basic interests of research participants. When this is the case, additional guidance is necessary in order to determine when such incremental risks to the basic interests of participants are consistent with an equal regard for the basic interests of participants and non-participants.

The same is true for healthy subjects who participate in clinical research. Here subjects expose themselves to research procedures that often most directly impact their personal interests. But these procedures also pose some additional degree of risk to their basic interests. The above principles allow subjects to be offered research participation that involves incurring some costs in terms of their personal interests. It remains to establish, however, what kind of incremental risks to the basic interests of subjects it is permissible to ask potential subjects to take on in order to facilitate scientific inquiry.

It should be emphasized that this remains an open problem in research ethics since the only approach to this question that has been articulated to date of which this author is aware is the Common Rule Approach. In conclusion, therefore, I make a tentative and largely schematic proposal about how we might make operational in these cases the underlying moral commitment to showing equal regard to the basic interests of participants and non-participants:

Second operational criterion. In all cases, the cumulative incremental risks to the basic interests of individuals that derive from research activities that are not offset by the prospect of direct benefit to the individual must not be greater than the risks to the basic interests of individuals that are permitted in the context of other socially sanctioned activities that are similar in structure to the research enterprise.

The guiding intuition behind this second operational criterion is that in order to pursue their particular projects and goals, individuals routinely take on some degree of risk to their basic interests. Respect for the moral equality of individuals, therefore, cannot require that individuals be prohibited from voluntarily assuming some risk to their basic interests, since such a standard simply cannot be achieved. The question, then, is how to establish when incremental risks to an individual's basic interests violate the underlying commitment to moral equality.

The second operational criterion addresses this question by requiring stakeholders to identify social activities that are structurally similar to the research enterprise and to ensure that the

incremental risks to the basic interests of participants associated with purely research-related activities do not exceed the incremental risks to the basic interests of individuals associated with those structurally similar social activities. The central challenge then lies in delineating criteria for structural similarity that can be used to locate relevant comparison classes of activities and then in determining how to make these comparisons in practice. Here I can only explain what I mean by criteria for structural similarity and outline the kind of considerations that stakeholders would have to take into account when constructing relevant comparison classes of activities.

The requirement of structural similarity is meant to capture the idea that it may not be appropriate to use just any social activity to determine what kind of incremental risks to the basic interests of participants are morally permissible. For example, some people enjoy mountain climbing or auto racing at least partly because of the thrill that comes from the associated risks. In clinical research, however, increased risks should not add to the attractiveness of the activity. Rather, the goal of clinical research is to benefit future patients and the associated risks are necessary evils that should be reduced as far as possible.

This feature of clinical research should therefore be used as a criterion of structural similarity. That is, appropriate comparison classes of activities should be ones whose primary social purpose is to benefit others, where the associated risks are viewed as necessary evils that should be reduced or eliminated where possible.

Additionally, when individuals perceive themselves as having control over salient features of an activity they are often willing to accept greater risks to themselves than in similar activities where they lack such control. This may help to explain why people are willing to tolerate greater risks from driving than from airline flight or other forms of public transportation. Such asymmetries matter in the research context because research participants put their interests in the hands of identifiable parties who possess a particular expertise that creates limited but important responsibilities on the part of those experts to protect and to safeguard the interests of those participants. This may militate in favour of comparing the risks to the basic interests of research participants that emanate from purely research-related procedures to risks to the basic interests of community members that are associated with social activities that involve this kind of principle-agent relationship.

Finally, it is imperative to avoid using as comparators activities where oversight mechanisms or safety regulations are poorly enforced or are widely recognized to be inadequate. For instance, coal mining has become a safer occupation than it was several decades ago because of tougher safety regulations. However, at the time of this writing there have been several high-profile accidents at mines that were repeatedly cited for safety violations. As a result, the risks that coal miners face in actual practice may be higher than what is judged to be socially acceptable, at least if existing health and safety regulations are an accurate guide to the latter issue. Appropriate comparator activities should therefore be the subject of active public oversight so that the risk-profile associated with the activity can be seen, at least *prima facie*, as representing a level of risk that is deemed socially acceptable after due reflection.

To be clear, the goal here is to find reasonable criteria of similarity that can be used to identify appropriate comparison classes of activities and then to examine the risk-profiles associated with activities that satisfy some or all of these criteria. This process itself may require careful adjustments in the criteria of similarity as well as discerning judgments about whether the activities that meet these criteria ought to be endorsed as appropriate comparators for clinical research. This is therefore an inherently normative or evaluative process. The objective is not to avoid making such normative judgments. It is, rather, to find a reasonable set of criteria that can be used to facilitate this process so that data that exist about the risks associated with socially important activities in one sphere

can be used to assess the incremental risks to the basic interests of research participants that come from purely research-related elements of a particular study.

One place to look for appropriate comparison classes of activities might be to public service professions, such as volunteer fire departments or paramedic services. The volunteer nature of these activities, combined with their orientation to serving the public interest represent important structural similarities to the research enterprise. Similarly, these occupations are often subject to varying degrees of public oversight. However, because there is no principal-agent relationship in these activities it may not be appropriate to permit in clinical research activities that have a risk profile that is similar to the most dangerous activities that individuals in these roles sometimes undertake.

The idea I am proposing is to use these comparison classes of activities to construct practical tests for this second operational criterion. Such a practical test would require, for example, that the incremental risks to the basic interests of research participants associated with the purely research-related elements of a clinical trial should not be greater than the incremental risks to the basic interests of fire fighters or paramedics that members of these professions face on a routine basis. If a phase I clinical trial involving healthy participants failed to meet such a test, then it would have to be redesigned, delayed until further work in animal models could be completed, or the data would have to be generated in some other fashion.

CONCLUSION

I have argued that the two most prominent approaches to evaluating the risks associated with clinical research suffer from significant problems. In response I have outlined some of the issues that I believe must be addressed in order to advance the state of the debate in this area of research ethics. I have also tried to illustrate how some of these issues might be addressed by outlining an alternative approach that provides a compelling normative foundation for a framework for assessing risk in clinical research that has significant operational content. In particular, I argued that the first operational criterion of this new approach and its associated practical test capture important aspects of the Duty of Personal Care Approach while avoiding some of the shortcomings that attend that view. The second operational criterion that I sketched above represents a tentative proposal for making headway on what remains an open question.

As Aristotle wisely cautioned, it is imperative not to seek in any domain a greater degree of precision than the subject matter permits [26]. Although readers will undoubtedly find different aspects of my positive proposal more or less compelling, it represents an attempt to illustrate some areas where a framework for decision making in research ethics might be more precise than the current alternatives. In this regard, there are several respects in which the Integrative Approach might be further refined. For example, the condition of admissibility that is outlined in the first practical test above could be formulated more precisely from a decision theoretic standpoint. Important work along these lines has been carried out by Kadane and colleagues and a general decision theory for dealing with decisions of this type has been articulated by Isaac Levi [22, 27]. As Seidenfeld and colleagues note, however, Levi's program faces a significant problem in situations where decision makers disagree about both the likelihood that an event will occur and about the utility or value of that event [28]. Refinements in these decision theoretic elements of the Integrative Approach will enhance its ability to provide practical guidance in a larger number of cases.

Similarly, there is in political philosophy and developmental economics a rich discussion about how to define and measure what have been called 'basic capabilities' [29–32]. Systematic treatment

of this issue in the medical domain is necessary in order to further refine the operational content of the distinction between personal and basic interests that is fundamental to the Integrative Approach. When it comes to constructing practical tests for the second operational criterion outlined above, however, there may be significant limits on the degree of precision that can be achieved from the construction of relevant comparison classes.

Regardless of the merits of this particular proposal, however, it is imperative to recognize the need to articulate a compelling normative foundation for the assessment of risks in clinical research, to search for ways of formulating the central questions in this domain that capture the fundamental moral issues and that use where possible methods from statistics, decision theory, and the social sciences to sharpen the operational guidance the framework provides to decision makers.

ACKNOWLEDGEMENTS

This paper was written with the generous support of a New Directions Fellowship from the Andrew W. Mellon Foundation.

REFERENCES

1. U.S. Department of Health and Human Services. Protections of Human Subjects. 56 Federal Register 28012, 45 Code of Federal Regulations 46.
2. London AJ, Kadane JB. Sham surgery and genuine standards of care: can the two be reconciled? *The American Journal of Bioethics* 2003; **3**(4):61–64.
3. London AJ. Sham surgery and reasonable risks. In *Cutting to the Core: Exploring the Ethics of Contested Surgeries*, Benatar D (ed.). Rowman & Littlefield: New York, 2006; 211–228.
4. Jonas H. Philosophical reflections on experimenting with human subjects. *Daedalus* 1969; **98**(2):219–247.
5. National Commission for the Protection of Human Subjects of Biomedical and Behavioral Research. *The Belmont Report*. US Government Printing Office: Washington, DC, 1979.
6. Miller FG, Brody H. A critique of clinical equipoise: therapeutic misconception in the ethics of clinical trials. *The Hastings Center Report* 2003; **33**(3):19–28.
7. Miller FG, Brody H. What makes placebo-controlled trials unethical? *American Journal of Bioethics* 2002; **2**(2):3–9.
8. Weijer C, Miller PB. When are research risks reasonable in relation to anticipated benefits? *Nature Medicine* 2004; **10**(6):570–573.
9. London AJ, Kadane JB. Placebos that harm: Sham surgery controls in clinical trials. *Statistical Methods in Medical Research* 2002; **11**(5):413–427.
10. London AJ. Two dogmas of research ethics and the integrative approach to human-subjects research. *Journal of Medicine and Philosophy* 2006, in press.
11. London AJ. Clinical equipoise: foundational requirement or fundamental error? In *The Oxford Handbook of Bioethics*, Steinbock B (ed.). Oxford University Press: New York, 2006, in press.
12. Miller FG. Sham surgery: an ethical analysis. *American Journal of Bioethics* 2003; **3**(4):41–48.
13. Fried C. *Medical Experimentation: Personal Integrity and Social Policy*. North-Holland: Amsterdam, 1974.
14. Marquis D. Leaving therapy to chance. *Hastings Center Report* 1983; **13**(4):40–47.
15. Miller PB, Weijer C. Rehabilitating equipoise. *Kennedy Institute of Ethics Journal* 2003; **13**(2):93–118.
16. Peto R, Baigent C. Trials: the next 50 years. *British Medical Journal* 1998; **317**:1170–1171.
17. Freedman B. Equipoise and the ethics of clinical research. *The New England Journal of Medicine* 1987; **317**:141–145.
18. Freedman B. Placebo controlled trials and the logic of clinical purpose. *IRB: A Review of Human Subjects Research* 1990; **12**(6):1–6.
19. Enkin MW. Clinical equipoise and not the uncertainty principle is the moral underpinning of the randomized controlled trial: against. *British Medical Journal* 2000; **312**:757–758.
20. London AJ. Threats to the common good: biochemical weapons and human subjects research. *Hastings Center Report* 2003; **33**(5):17–25.

REASONABLE RISKS IN CLINICAL RESEARCH

21. Rawls J. Social unity and the primary goods. In *Utilitarianism and Beyond*, Sen A, Williams B (eds). Cambridge University Press: New York, 1982; 159–185.
22. Kadane JB (ed.). *Bayesian Methods and Ethics in a Clinical Trial Design*. Wiley: New York, 1996.
23. Gifford F. The conflict between randomized clinical trials and the therapeutic obligation. *Journal of Medicine and Philosophy* 1986; **11**(4):347–366.
24. Hellman D. Evidence, belief, and action: the failure of equipoise to resolve the ethical tension in the randomized clinical trial. *Journal of Law, Medicine and Ethics* 2002; **30**:375–380.
25. Evans EL, London AJ. Equipoise and the criteria for reasonable action. *Journal of Law, Medicine and Ethics* 2006; **34**(2):441–450.
26. Aristotle. Nicomachean ethics in Barnes. In *The Complete Works of Aristotle*, Jonathan (ed.). Princeton University Press: Princeton, 1984; 1729–1868.
27. Levi I. *Hard Choices: Decision Making Under Unresolved Conflict*. Cambridge University Press: New York, 1986.
28. Seidenfeld T, Kadane JB, Schervish MJ. On the shared preferences of two Bayesian decision makers. *The Journal of Philosophy* 1989; **86**(5):225–244.
29. Sen A. Equality of what? In *Choice, Welfare, and Measurement*, Sen A (ed.). Harvard University Press: Cambridge, MA, 1982; 353–369.
30. Sen A. *Commodities and Capabilities*. Oxford, New York, 1999.
31. Nussbaum M. *Women and Human Development: The Capabilities Approach*. Cambridge University Press: New York, 2000.
32. Anderson ES. What is the point of equality? *Ethics* 1999; **109**:287–337.