Rebecca Kukla

Resituating the Principle of Equipoise: Justice and Access to Care in Non-Ideal Conditions

ABSTRACT. The principle of equipoise traditionally is grounded in the special obligations of physician-investigators to provide research participants with optimal care. This grounding makes the principle hard to apply in contexts with limited health resources, to research that is not directed by physicians, or to non-therapeutic research. I propose a different version of the principle of equipoise that does not depend upon an appeal to the Hippocratic duties of physicians and that is designed to be applicable within a wider range of research contexts and types—including health services research and research on social interventions. I consider three examples of ethically contentious research trials conducted in three different social settings. I argue that in each case my version of the principle of equipoise provides more plausible and helpful guidance than does the traditional version of the principle.

One is in a state of equipoise when one has no good basis for a choice between two or more options. According to the traditional version of the principle of equipoise, a state of equipoise or an "honest null hypothesis" (Levine 1986) with respect to the expected health outcomes of subjects in different trial arms is a condition upon the ethical acceptability of beginning—or continuing—a research trial. I will call this traditional version of the principle PE:

> PE: In order to begin or to continue an experiment on human subjects, one must be in a state of equipoise with respect to the relative expected health outcomes for participants in different trial arms.

Although fierce debates currently rage over the validity and interpretation of the principle,[1] PE "has become a widely accepted axiom governing the ethics of randomized controlled trials" (Miller and Brody 2002, p. 4) since it was introduced by Charles Fried in 1974.[2]

Kennedy Institute of Ethics Journal Vol. 17, No. 3, 171–202 © 2007 by The Johns Hopkins University Press

Overwhelmingly, discussions of PE have presupposed physician-investigators engaged in therapeutic research and have grounded the principle in physicians' Hippocratic duties to promote the welfare of their patients. Charles Weijer (2002, p. 116) points out that historically worries about equipoise arose directly out of questions concerning when physicians legitimately could offer trial enrollment to their own patients. The thinking is that it is unethical for a physician to offer a treatment option to a subject participant that reasonably can be expected to be inferior to the optimal care she would offer her own patients outside of the trial. PE—combined with the assumption that not all of the trial arms involve substandard treatment, for otherwise one could run a trial in which the arms were equal but uniformly substandard—precludes study designs with this result. As Paul Miller and Charles Weijer (2003, p. 93) put it, PE has been presented as a way of resolving "the moral tension between the physician's commitment to the personal care of her patients on the one hand, and her commitment to a program of research on the other." The ethical persona and the agent-relative duties of the physician herself have been used to defend PE in all of the classic and almost all of the recent discussions of PE.[3] Even the most powerful critics of PE have assumed that the principle relies upon the therapeutic duties of its physicians for its justification.[4]

Yet a great deal of research is not actually conducted by physicians, or by anyone with special Hippocratic duties. Health services research and public health research, for instance, may well be performed by program evaluators or social scientists. Furthermore, much health research, regardless of who performs it, is aimed at prevention, cost-cutting, or other goals besides therapy. To the extent that one grounds research ethics in the ethics of therapeutic clinical medicine, these other kinds of research will be left in an unconstrained ethical vacuum.

Against orthodoxy, I think that it is clear that the ethical concerns that inspire PE cannot be grounded solely in the dual role of physician-investigators as researchers and as caregivers. For surely one cannot turn unethical research projects into ethical projects just by putting a non-physician in charge. Consider any of the paradigmatic cases in which violations of PE are ethically troubling, which tend to be cases in which one trial arm receives an experimental treatment, while the other receives a placebo or a clearly substandard treatment even though a known effective treatment exists. I contend that in none of these cases would our ethical concerns be allayed by putting a non-physician at the helm. As far as our ethical

qualms are concerned, it does not seem to matter in the least whether such a trial is run by a physician or, say, an epidemiologist with no therapeutic obligations. Physicians who conduct therapeutic research may well face special ethical complications and duties, such as worries about preventing therapeutic misconceptions, or duties to provide ancillary care. Indeed, physicians may face special ethical pressures arising from conflicts between their therapeutic obligations and the goals of research. But violations of PE of the sort that have most interested bioethicists cause ethical discomfort that arises independently of the persona of the physician and her purported Hippocratic duties to her patients.

In this paper, I challenge the traditional grounding of PE in the therapeutic obligations of physician-investigators. I argue that this grounding requires one to assume an idealized research context of unlimited resources and access to care that rarely is incarnated, and that it thereby leaves one without helpful ethical intuitions or guidance in the non-ideal conditions in which research normally occurs. I claim that by attending instead to the general moral requirements of justice and respect for persons in non-ideal conditions, one can earn back the important moral intuitions that PE was designed to capture, without depending upon the image of the Hippocratically bound physician-investigator. I argue that there is nothing inherently ethically problematic about an expected inequity between trial arms. Rather, research ethicists' concern should be with the broader issue of expected inequities between the care offered to research participants and the care they deserve to receive as citizens. PE, with its narrow focus on trial arms, can address this larger moral concern only circuitously. I propose a different principle of equipoise that addresses it more directly, while being applicable within a wider range of actual research contexts. My version of the principle of equipoise does not focus on equipoise with respect to the relative expected outcomes of trial arms, but rather on equipoise concerning the social value of the intervention being tested. This version of the principle, I argue, provides a defeasible ethical constraint; although there are exceptional situations in which the principle can be waived, it provides important ethical guidance in typical research contexts. I consider three concrete examples of ethically contentious research projects conducted in three very different nations and social settings. Each of these research projects violates traditional equipoise. I try to show that in each case, my version of the principle of equipoise provides more plausible and helpful guidance than does PE.

Many defenses of PE are based on the idea that a trial is unethical if one or more of its arms receives an intervention that the investigators have good reason to believe is inferior to the *best* existing care; indeed, this principle is enshrined in the notorious Paragraph 29 of the Declaration of Helsinki, which states that "The benefits, risks, burdens and effectiveness of a new method should be tested against those of the best current prophylactic, diagnostic, and therapeutic methods."[5] The idea behind this principle is that physicians, who are presumed to be the ones doing the research, are duty-bound to provide their patients with optimal medical care, and hence should provide their research subjects with nothing less.

In fact, of course, physicians rarely provide unrestricted optimal care to their patients, inside or outside of a study. In developing countries, people often have little or no access to decent care. Much of the American bioethics literature contrasts such conditions with the purported situation in developed countries, where access to therapeutically optimal care is supposedly the norm (see, e.g., London 2001, pp. 328–29). However, there is no nation on earth where more than a tiny minority of patients has unrestricted access to such care. In a primarily privatized health care system like that of the United States, whose performance on standard measures of health care access and health outcomes is mediocre, a physician-patient relationship structured by absolute fidelity without regard to economic circumstances can be enjoyed by the exceptionally wealthy, but not by the average patient, whether insured or uninsured. In the countries with the most successful health care systems, such as Canada and the Scandinavian countries, rationing is built into the system of public insurance. Hence, if one presupposes physician-patient relationships defined by unfettered fidelity and therapeutic obligations as the background against which to begin discussions of research ethics, one will be appealing to a decontextualized ideal that almost never is implemented in any domain of medical practice. In discussions of equipoise, the particular distorting effect of this ideal is to produce the working assumption that outside of a study, in the context of clinical care, the mere existence of a better treatment is sufficient to give patients a right to it, and physicians a duty to provide it. As a result, it appears that any research participant who does not have access to this better treatment is being denied access to something to which she is entitled.

For example, David Steinberg (2002, p. 27) claims, "people who agree to participate in clinical investigations do not relinquish the *right* to *optimal*

medical care." More strikingly, Robert Veatch (2002, p. 312) has defended the idea that "subject indifference" between trial arms, rather than traditional clinical equipoise, is a "morally necessary condition for bringing a subject into a randomization process." In other words, according to Veatch, an investigator can only enroll someone in a trial if the potential participant herself, once properly informed, does not care which trial arm she is assigned to. For "if a patient prefers one of the treatment arms in this sort of case, she should get it . . . Whether the subject gets access to the treatment inside the protocol or outside of it, she should have the right sort of access to it" (p. 312). Although Veatch later adds the qualification "unless the treatment is inherently scarce or very expensive" (p. 318), he clearly views this as the exception rather than the rule. In fact, however, almost all health care resources are scarce, although they may not appear so from the perspective of wealthy patients living in countries with privatized health care systems. Such a patient-centered version of equipoise can look appealing only if one assumes an idealized situation of absolute access to optimal interventions as the norm.[6]

Research ethics needs to begin from the recognition that conditions of scarcity are the norm in health care and that not all research is physician-run, therapeutic research. The accessibility of an intervention is governed by complex economic, cultural, and material conditions, and not by the mere march of scientific knowledge. When research ethicists treat limited access or scarce resources as an exception to the general rule—one that mostly concerns distant, underprivileged countries rather than "us"—we will end up with ethical principles that fit poorly with the reality of health care practice.

JUSTICE, RESPECT, AND STANDARDS OF CARE

The fundamental intuition from which research ethics stems—the intuition that, I would argue, underwrote the *Belmont Report*—is that, *regardless* of who happens to be doing the investigating, one should never let the ends of scientific investigation interrupt or trump one's basic moral treatment of those whose bodies are used in their pursuit. Other people must be treated respectfully, justly, and with concern for their well-being—hence the *Belmont Report*'s three principles (respect, justice, and beneficence).[7] Human-subjects research tempts researchers to use others instrumentally, as bodies that are means to an end, compromising their moral relation to them as persons. The principles of research ethics, in essence, should serve as tools that help one avoid this potential moral

wrong. Our moral obligations to persons as persons, which include obligations of justice, respect, and welfare protection, cannot be suspended or compromised for purposes of research.[8] I propose this as the fundamental, most general principle of research ethics, against the background of which any other principle, including any principle of equipoise, should be formulated. This principle has some direct corollaries. For instance, if researchers cannot compromise their moral treatment of persons as persons in the name of research, then they cannot demean, humiliate, or intimidate them. And, importantly for my purposes in this paper, researchers also cannot knowingly prevent research participants from receiving care that they are morally entitled to receive as ordinary citizens.

In other words, in the language of research ethics, one might say that no research participants should knowingly be given care that is inferior to the *standard of care*. However, as many have pointed out, the term "standard of care" is multivalent. Alex London (2000) distinguishes different senses of "standard of care" along two dimensions. The term "standard" is ambiguous between a descriptive or *de facto* meaning, namely that which is typical, average, or widespread, and a normative or *de jure* meaning, namely a standard to which particular cases *ought* to live up. Accordingly, one can distinguish between the *de facto* standard of care, which is the care that people actually can expect to receive, and the *de jure* standard of care, which is the care to which people *ought* to be entitled. At the same time, we might ask either type of standard might be indexed to a local or a global context. So for example, using London's divisions, one might ask about local, *de facto* standards of care: what kind of care can members of a given population, at a particular time and place, expect to receive? Or one might ask about global, *de jure* standards of care: to what sort of care ought people in general to be entitled?

Although London uses his two distinctions to demarcate four senses of "standard of care," notice that there is no such thing as a global *de facto* standard of care; different health care systems and social conditions are so diverse that surely there is no "typical" or "average" standard that helpfully applies to both citizens of wealthy nations with top-notch, universally accessible health care systems, and to citizens of disease-stricken developing nations lacking even basic social services. Furthermore, given the vast differences in *de facto* medical, economic, and cultural conditions between different localities, it seems naïve to think that there could be one *de jure* standard appropriate for all of them—even though there may well be a global *minimum* standard of the sort that might be written into

a declaration of human rights. The sort of care to which people *ought* to be entitled must depend on the local resources, the cultural facts about what sort of care people want, the medical facts about what they need, and so forth.

This leaves local *de facto* and local *de jure* standards as contenders for what one means when one says no one should knowingly be dropped below the standard in the name of research. Many people have worried that in moving to a localized standard of care, one risks complicity with societies that are (*de facto*) providing their citizens with inappropriately poor care, even given their relatively scarce resources. London (2000, p. 385) points out that often the response has been to insist upon an implausible and ill-defined global standard of care to which researchers should hold their trials—one that risks ruling out altogether swaths of important research designed to find interventions that will be effective in resource-starved locations with necessarily low local standards of care. However, as London argues, we can solve the problem more plausibly by moving to requiring the *de jure* standard of care, relativized to local circumstances, as the bar of acceptability for trial arms. A very poor country, stretching its resources between health care, education, infrastructure, and so forth, surely would offer less health care to its citizens than a wealthy nation, even if it used the resources it had in a maximally just and efficient fashion. Yet many countries, both poor and rich, fail to use their resources justly and efficiently, and offer poorer care than they ought. It is wrong for researchers to take advantage of *de facto* local standards that are unjustly low, given resources of the region, in designing their research trials. But it is not wrong for them to run a trial in which no one receives less than that to which they *ought* to be entitled, given the local cultural, material, and economic circumstances, even if participants thereby receive less than that which they would receive in a different social context. The policies and politics of a region, whether just or unjust, will, of course, help to determine the material and economic resources of that region, over time. However, at any time T, one can coherently ask to what kinds of health care services and protections the inhabitants of a region *ought* to have access, given that region's material conditions, but regardless of its actual policies and politics.[9] Such services and protections constitute the *de jure* local standard of care.

However—to understate the case—it is by no means easy to determine what the local *de jure* standard is. As already indicated, the mere fact that something is medically beneficial does not prove that it ought to be

part of the local standard of care.[10] Likewise, I have argued that the local *de facto* standard of care cannot be used as a direct measure of the *de jure* standard; many communities distribute their resources unjustly, inefficiently, or both. Whether an intervention ought to be generally accessible depends on a complex combination of factors including economic factors, narrowly medical facts, social support systems, local preferences and values, and much more. Indeed, whether an intervention should be part of the local standard of care is, directly or indirectly, precisely what investigators are *trying to find out* in the course of most medical research. As John Arras and Robert Crouch (1998, p. 29) point out, we normally are not in a position to judge whether a population is entitled to access to an intervention until *after* trials are run to determine its safety and efficacy—and, Arras and Crouch surely would agree, its local ease of delivery, local acceptability, and so forth. London places it as a *condition* upon research that no trial arm fall below the local *de jure* standard of care. But most ethical health research is aimed at *determining* what this standard is, and hence investigators cannot begin by demanding that it be met. They can, however, insist that it not be *knowingly* violated.

RESITUATING THE PRINCIPLE OF EQUIPOISE: CHANGING THE QUESTION

At this point, I have defended and sharpened the key moral intuition that drove PE, but with no reference to the special duties of physicians and no assumption that physicians will be the ones conducting research: researchers should not run studies unless, *to the best of their knowledge*, every trial arm receives care that is at least as good as the local *de jure* standard of care. I dub this the "Minimum Standard" principle, or MSP. Determining the local *de jure* standard of care is a matter of settling normative questions about the just and appropriate way to set up social structures and services given a particular cultural and material context. It is not a matter of asking what physicians owe their patients in virtue of their special duties to them. I argued that MSP is a direct consequence of the foundational principle of research ethics, namely that investigators' moral obligations to persons as persons cannot be suspended or compromised for purposes of research. Although it may be acceptable to allow trivial exceptions to MSP when the costs are low, the potential benefits are high, and informed consent is responsibly obtained—for instance, asking participants in a trial arm to accept temporary nausea or similar discomforts for an important cause—such exceptions do not seem to affect the core ethical point, namely that participation in research should

not compromise participants' usual moral entitlements to justice, respect, and care in their particular local contexts.

So far, I have not used or proposed any principle of equipoise. A principle of equipoise, by definition, demands some specific form of uncertainty, and I have not demanded any such thing. MSP should accompany and complement any acceptable principle of equipoise, but many authors mistakenly have tried to fold it into their formulation of the principle of equipoise.[11] This conflation is not surprising, for PE is not very helpful unless it is conjoined with something like MSP: If one requires uncertainty as to which trial arm is likely to have better outcomes, but does not also require that no trial arm knowingly receive less than some minimum standard of care, then one would allow a trial in which all arms are expected to have terrible, but equal, outcomes. Clearly, anyone interested in defending PE would want to block such a trial, and hence would also want to defend some version of MSP—that is, some principle that sets a minimum standard for how research participants can be treated, regardless of which trial arm they are enrolled in. (One can read paragraph 29 of the Declaration of Helsinki as a maximally strong version of such a minimum standard, and the view that participants need to receive care no better than the local *de facto* standard as a very weak version of such a standard; MSP falls in between these two extremes.)[12] Nevertheless, MSP is not a principle that demands uncertainty, and hence it is not properly part of any principle of equipoise.

Remember that PE demands an honest null hypothesis with respect to which trial arm can be expected to have better outcomes:

> PE: In order to begin or to continue an experiment on human subjects, one must be in a state of equipoise with respect to the relative expected health outcomes for participants in different trial arms.

The traditional justification for PE concerns the relationship between subjects *inside and outside of the study*: the principle is designed to prevent (some) study participants from receiving care that is inferior to what their physicians would owe them as patients outside of the study. Yet PE focuses only on the *relative treatment of trial arms* and says nothing about the relative treatment of people inside and outside of the study. This makes it a rather roundabout ethical principle.

Arguably, if (informed, voluntary) participants in trial arms may be expected to have unequal outcomes, this is *in and of itself* no more ethically troubling than the fact that participants in a raffle for a day at the

spa or in drawing straws to determine who gets the corner office may be expected to have unequal outcomes. As a society, we accept and create all kinds of random inequalities as a matter of course. Our real concern is not with inequalities internal to a trial, but rather that in the ethically charged, high-stakes domain of health, no one should unfairly receive care that is inferior to that which he or she is entitled to receive. Inequities within the trial itself are only morally troubling if one is worried that those who are randomly assigned to a less advantageous trial arm are *also* suffering from unjust or insufficiently caring treatment on a larger scale—that is, lucky random allocation to an advantageous trial arm should not be the means for obtaining access to care than one should be entitled to receive anyhow.

I propose a new version of the principle of equipoise that does not concern differentials between trial arms. Instead, it demands genuine uncertainty, or an honest null hypothesis, with respect to whether or the extent to which the intervention being tested *ought* to be part of the local standard of care, and hence concretely accessible to the local population:

> PE*: In order to begin or to continue human subjects research, one must be in a state of equipoise with respect to whether or the extent to which the intervention being tested *should* be made accessible to the population that falls under the scope of the research.

In the previous section, I argued that normally the local *de jure* standard of care is what investigators are trying to discover during the course of research, rather than something that is known and can be insisted upon for all trial arms prior to beginning research. My PE* makes the stronger claim that such uncertainty as to the local *de jure* standard should be—in almost all circumstances—*required* as a condition of research.

The population falling under the scope of the research, as I am using the phrase, is the population over which the research results are designed to generalize. If reasonable certainty already exists that an intervention should be made generally accessible to a particular population, then this is just to say that it is already established as part of the *de jure* standard of care for that group. In this case, investigators ought not to be conducting research that denies that intervention to members of that population, including research participants. Researchers cannot get a pass for violating the moral entitlements of citizens, and this is so even if the participants are citizens of an unjust local regime that already violates these entitlements. Subjects who do not have access to the goods and care to which

they are morally entitled may well be willing to participate in studies that are complicit with these low standards, but this makes the studies no less exploitative. On the other hand, if reasonable certainty already exists that the intervention to be tested should *not* be made available to the population on whom it is to be tested—because it is too expensive, too risky, too hard to administer, or whatever—then one ought to be *prima facie* suspicious of the reasons for conducting the testing on that population. Although I consider some exceptions to this rule later, health research generally should be directed at improving the health and well-being of the population falling under the scope of the research.[13] Therefore, tests on an intervention typically are ethically appropriate only when the question of the place of this intervention in the local *de jure* standard of care is *not yet settled*. And this is just what PE* demands.

Some (hopefully) uncontentious examples should help clarify the point. In many developing nations, access to safe drinking water is not standard. Almost all of us would agree that even in the poorest nations, safe drinking water ought to be a right of all citizens and a top priority in resource allocation. It would be unethical for researchers to run a study that compared health outcomes in a trial arm receiving safe drinking water to those in an arm denied safe drinking water, even if, *de facto*, access to safe drinking water was not the local standard. To run such a study would involve unethical complicity in unjustly low local standards of care, and hence it would fail both PE* and MSP. On the other hand, consider a potentially risky, astronomically expensive new cancer treatment. Imagine that there already exist cheaper and reasonably effective treatments for this type of cancer, but that the new treatment holds the promise of a quicker recovery time. Presumably we all would agree that it would be inappropriate to test the new treatment in a poverty-stricken community where this treatment would never be made generally accessible, regardless of how well it turns out to perform on tests. Such a study would fail PE*, even though it does not fail MSP. Both these trials are unethical because in neither case is there equipoise concerning appropriate access to the intervention in the local context in which the research is being carried out.

PE* makes room for research that is designed to determine not just whether an intervention should be accessible, but what kind of accessibility is appropriate. Those who are in a position to distribute resources within a community may decide, after studying an intervention, that it is worth making it accessible to some subset of a population but not to everyone—for example, they may decide that only people with certain health risks or

in certain professions should receive flu shots.[14] As long as such questions of appropriate access are genuinely open, research designed to settle them is allowed by PE*. Of course, physician-investigators legitimately may be interested only in the narrowly medical effects of an intervention. Such information is a crucial piece of the ultimate resource allocation puzzle. But such research ought to be helping to answer an *open question* about whether or the extent to which this intervention ought to be accessible to the population falling under the scope of the research.

Many commentators have pointed out that those who assess the ethical acceptability of a research project need to be concerned not only with the internal features of its methodology, but also with how well-suited the research is to the local context in which it will be applied. For instance, Scott Halpern (2006, p. 3) suggests that "research responsiveness" ought to be an ethical condition upon research, along side PE: "Merely addressing a prominent problem is insufficient if the study cannot reasonably be expected to further efforts to solve that problem . . . [research] ought to reflect thoughtful analysis of the potential for interventions in the local setting." PE only asks one to consider the relative outcomes of the trial arms, and hence any concerns about the relationship between the research project and its application in a local context must be addressed in a separate principle. PE*, in demanding equipoise concerning whether it is appropriate to make an intervention generally available in this local setting, builds in at least some of the responsiveness that Halpern seeks.

PE* is substantially different from PE, because it changes the type of question that ought to be open from a question about the relative expected outcomes for individual study participants, to a quite different question about the social value of an intervention. In the next section, I show that PE and PE* often yield very different ethical verdicts. And yet, there is an interesting convergence between PE and PE* as the context of research approaches ideal conditions of unlimited resources. In a community with unlimited economic resources, unlimited health professionals with time to devote to patients, no real barriers to health care access, and so forth, the *only* factor that would determine whether an intervention should be accessible would be how beneficial it was to the individuals that received it, compared to alternative interventions. In such a case, PE would demand that investigators try to figure out whether the new intervention provides genuine benefits compared to the existing alternatives, and stop testing once they know.[15] And in this ideal context, PE* would demand *just the same thing*, because the question of whether the new intervention is *worth*

making accessible *would just be* the question of whether it would provide people with genuine benefits compared to the alternatives. Given that both PE and PE* must be accompanied by MSP, in neither case could investigators test the new intervention against anything other than the best known alternative, since in a mythical situation of unlimited resources, everyone would be entitled to care as effective as the best known alternative. Hence in such a context, given MSP, both PE and PE* would recommend and rule out the same trials. The real world, of course, rarely approximates such an idealized context.

PE* has potentially surprising results when it comes to cases in which one believes that the current *de facto* standard of care might be *higher* than it should be. Imagine that an intervention has become standard, and although it has some genuine benefits, one might suspect that its benefits are not large enough to justify its costs.[16] In this case, a randomized controlled trial that tested the *de facto* current standard of care against a *lower* standard, which is perhaps the *de jure* standard, would satisfy PE* (but not PE). But one might object that there is an ethical problem with dropping some participants below the current *de facto* standard, even if that standard might be higher than it ought to be, on the grounds that it is unfair for subject participants to get less than do their peers outside the study. For instance, consider a trial designed to test the effects of reducing hospital stays after a particular procedure. It might be a genuinely open question whether such a reduction in hospital stays is a good way to free up resources, all things considered. Yet one might think it wrong to take away some participants' right to the length of stay they would have access to outside of the study. However, I think that good-quality informed consent can alleviate these worries. Although it is certainly wrong to deny people something that they have a *de facto* right to receive, I cannot see any ethical objection to letting people make an informed, autonomous choice to volunteer to forgo it, perhaps out of a larger commitment to the social cause of the just distribution of resources (although it might be hard to find volunteers for such a study). In contrast—except in the case of minor discomforts—informed consent is not sufficient to override the injustice of knowingly giving some people care that is inferior to that to which they are *morally* entitled. PE* combined with MSP rules out this latter possibility but not the first.

I have proposed two (defeasible) conditions upon the ethical acceptability of a study:

MSP: Investigators should not knowingly offer any research participants care that is inferior to that which they would be morally entitled to receive outside of the study, given their local material, economic, and cultural context.

PE*: In order to begin or to continue human subjects research, one must be in a state of equipoise with respect to whether or the extent to which the intervention being tested should be made accessible to the population that falls under the scope of the research.

Notice that neither of these conditions makes any reference to physician-investigators or to therapeutic obligations, nor did I assume a physician-investigator in their defense. Instead, I claim that these principles follow from the more basic principle that the research enterprise gives investigators no license to compromise citizens' moral entitlements to justice, respect, and welfare protection. How helpful and convincing these conditions are depends on how compellingly they illuminate actual hard cases that trouble those of us in research ethics. In the next section, I examine three such hard cases.

APPLICATIONS

Preschool in an Impoverished East London Borough

The Mapledean Early Years Centre provides a preschool program for the children of Hackney, a low-income borough of East London in which the demand for preschool spaces far outstrips supply. Spaces in the program are allocated randomly to families that request them. Although nobody doubts that the preschool provides a real benefit to the families it serves, it would be helpful to know the exact size and character of these benefits: Whom does preschool benefit most? What is the impact of the program on health outcomes, socioeconomic outcomes, and academic outcomes? To test some of these outcomes, researchers ran a randomized controlled trial in which some participants were randomized into a trial arm in which they received a space in the Mapledean program and others into a trial arm in which they did not. Tami Toroyan, Ian Roberts, and Anne Oakley (2000) point out that this trial clearly violates PE; no one seriously doubts that the families receiving spots in the program can be expected to have better outcomes than those that do not. However, they argue that the trial was ethical, because of the scarcity of preschool spots in the area; since most local families seeking such spots will be denied them anyhow, they argue, no research participants are denied anything they could otherwise expect to receive, and hence the randomization is acceptable. S. J. Edwards and

S. Kirchin (2002, pp. 20, 21), in a reply, argue that even in a situation where resources are scarce, randomization is ethically acceptable only when no one yet knows who those resources would help most—that is, they argue, investigators must be in a state of equipoise, not with respect to which trial arm will do better, but with respect to "how to allocate resources fairly among the current population," or "who out of a given set of individuals should get the treatment." They agree that the Mapledean trial is ethical, but only because prior to the study there existed genuinely uncertainty, for instance, as to whether the preschool spots would do the most good for the children of single parents, or the poorest children, or the children with learning delays, and so forth. Once such information becomes available, they claim, we ought to distribute the preschool spaces accordingly rather than randomly.

This trial involves social research that is not likely to be run by physicians,[17] and hence the traditional grounds for PE do not apply. Yet Edwards and Kirchin make a compelling case that some equipoise constraints ought to govern such research, although they do not offer an alternative foundation to justify such constraints. Although the trial violates PE, it appears to meet the requirements of PE*: although it seems clear that preschool provides a real benefit to low-income children, it is unknown just how much or what kind of a benefit it provides for the population in this neighborhood, and hence, presumably, we do not know how we ought to allocate resources to the development of more preschool spaces, all things considered. The purpose of the research is to determine the relative importance of preschool in improving the health and well-being of children in this type of community; as Edwards and Kirchin (2002, p. 22) put it, "the idea behind a trial is to establish that certain resources should be made more widely available." Furthermore, no one in the study receives a level of services below what we know to be their entitlement, since their appropriate entitlement is just what the investigators are trying to determine.

Edwards and Kirchin's proposed principle of equipoise is similar to mine in that it focuses on a question about the just allocation of resources rather than the relative outcomes of the trial arms, and in that it does not appeal to physicians' therapeutic obligations. However, I suggest that PE* is superior to their version of the principle of equipoise, because it takes into account a richer set of resource allocation issues. They call for equipoise concerning who ought to get a scarce intervention, given a particular level of supply. PE* requires us to ask larger questions about how the

intervention ought to be prioritized relative to other interventions, rather than just how to allocate the intervention itself. Since Edwards and Kirchin themselves point out that the ultimate goal of the Mapledean study is to decide whether the city ought to be *increasing* the availability of preschool slots, the broader focus of PE* is better suited to the case at hand.

The Special Supplemental Nutrition Program for Women, Infants, and Children (WIC) in the United States

The United States' Special Supplemental Nutrition Program for Women, Infants, and Children, commonly known as WIC, provides nutritious foods, health and social service referrals, and breast-feeding support to low-income women who are pregnant or have children under five years old. Although WIC is not an entitlement program, coverage is widespread and approaches 100 percent of eligible users. Although there have been several observational studies showing moderate health benefits from the program, including reduced rates of low birth weight, fewer preterm births, and better iron levels in children, there have been virtually no randomized controlled trials testing WIC's efficacy.[18] Recent critics have made vocal demands for randomized controlled testing of the WIC program, citing pervasive risks of selection bias in observational studies of WIC health outcomes (Besharov and Germanis 2001; Rossi 1998; Rossi and Hamilton 2002). (Selection bias is a risk because it is plausible that there are relevant systematic differences between low-income mothers who do and do not enroll in WIC, including overall concern with health, level of need, ability to utilize social services and support networks, and so forth.) David Besharov and Peter Germanis (2001, p. 75) write, "Randomized experiments may be the only way to develop valid estimates of WIC's impact, because they ordinarily do not require uncertain statistical adjustments to eliminate differences between treatment and control groups. . . . The need for randomized experiments is becoming increasingly clear to policy makers." Likewise, William Hamilton and Peter Rossi (2002, p. 30) insist upon "the importance of using *all possible efforts* to use randomized experiments" to test the effects of WIC.

But because slots in the WIC program are widely available, almost all commentators also have agreed that trials comparing a trial arm receiving WIC benefits to a no-benefits trial arm are ethically impossible. Citing participants' access to WIC benefits outside of the study, they argue that denying benefits to a control group inside of the study would be ethically unacceptable (see, e.g., Rossi 1998). The food nutrition policy commu-

nity has arrived at a bit of an impasse: the sentiment is that randomized tests of WIC's efficacy against a no-benefit control are both necessary in order to justify continued expenditure on the program, and impossible in practice.

Here then is another case of research that is not likely to be run by physicians, and hence an ethical analysis of its limits and possibilities must be grounded in something other than physicians' Hippocratic duties. As in the Mapledean case, a trial comparing an arm receiving the WIC package to an arm receiving no benefits would violate PE: However imperfect the observational data may be, it is hard to believe, under the circumstances, that providing nutritious food and social support to low-income pregnant women and mothers has *no* net positive effects.[19] Since PE looks only at the expected relative outcomes of trial arms, it offers no help in ethically differentiating between the WIC case and the preschool case, but simply rules out both research projects.

Unlike in the preschool case, WIC is an intervention that is widely available outside of the research context. On my account, in order to proceed with such research on WIC, PE* requires that a genuine state of equipoise over whether the program as a whole is worth keeping. We might agree that the program provides *some* benefit, but still claim equipoise with respect to whether this benefit is substantial enough to justify the cost of the program. If doubt really exists about this, then research such as that proposed would satisfy PE*, even though it would involve placing some subjects below the current *de facto* level of care they could receive outside of the study. As I argued earlier, investigators legitimately can drop people below the local *de facto* standard of care if they are in a genuine state of equipoise over whether this *de facto* standard is higher than it ought to be, although doing so requires exceptionally diligent informed consent procedures.[20] One might well wonder, however, who would give genuinely free informed consent to be in such a test, and whether it is ethical to deny benefits to families with children under five years old, given that these children could not consent to their own participation.

In any case, I do not believe that we can plausibly claim to be in such a state of equipoise. In 1992, the U.S. General Accounting Office estimated that the WIC program saved $3 on health care and related costs for every dollar spent, which would make it exceptionally cost effective (GAO 1992). Regardless of whether this figure has been distorted by its dependence upon observational studies, it seems unlikely that it is off by a factor of three. Economics aside, I am willing to stand by the claim that

in any nation, and certainly in any wealthy one, all pregnant women and young children are morally entitled to secure access to nutritious food and basic social support. Hence a trial that denied benefits altogether to some WIC-eligible families would violate PE*, since it is not really an open question whether some such benefits are part of the *de jure* standard of care for this population. Thus, unlike PE, PE* yields different verdicts on such a trial and on the Mapledean daycare trial: Whereas neither trial meets the requirements of PE, the Mapledean trial meets the requirements of PE*, but the (imagined) WIC trial does not.

On the other hand, it is not at all obvious that the current WIC program, in all of its details, is the best benefit package the government could provide to low-income families in the United States, all things considered. It may well be that WIC could be made more cost effective, or that the government ought to be providing a higher level of benefits or somewhat different benefits, or that the priorities of the program or its criteria for eligibility ought to be adjusted. Hence although we are not in a state of equipoise with respect to whether the program should be cut altogether, we are in a relevant state of equipoise with respect to which of various versions of the program would best further the ends of social justice and serve the needs of poor families. For example, reasonable experts disagree over whether the WIC food package overemphasizes carbohydrates such as cereal and juice, whether the provision of infant formula does more harm than good by discouraging breast-feeding, and whether the income cut-off for WIC is unnecessarily high. Thus there is plenty of room for randomized tests of WIC that satisfy PE* and my other criteria for the ethical acceptability of a study: studies that compared a trial arm receiving the standard WIC package to an arm receiving some plausibly worthwhile variation on the package may well be both ethical and useful in furthering the just and effective distribution of resources.

One can imagine many additions to the WIC benefit package that likely would yield substantial benefits for the individual recipients—unlimited fresh fruits and vegetables, for instance. Counterintuitively, PE tells us that a trial that compared a group receiving regular WIC benefits to a group receiving the regular benefits *plus* unlimited fresh fruits and vegetables would be unethical, simply in virtue of the fact that one would expect differential outcomes between the trial arms. Such an odd result is a product of the fact that PE originally was formulated against the background of an idealized notion of accessibility, according to which anything that clearly benefits an individual is automatically accessible to that individual

outside of the context of research. On such a model, a study in which participants were provided with less than "the best" care would be *hampering* participants' access to this care. But research on WIC is far removed from any such background. Low-income families do not have unlimited access to fresh fruits and vegetables outside of the research context; their limited access to high-quality food is the reason the WIC program exists in the first place. No loyal physician ensures that these families receive "the best" or "optimal" food. For these reasons, PE provides an inappropriate and unhelpful verdict on such research. In contrast, according to PE*, the ethical acceptability of such a trial depends upon whether it is a genuinely open question whether such a fruit and vegetable benefit ought to be added to the WIC package, all things considered. The mere expected differential between the trial arms does not tell against the ethical acceptability of the test.

Short-Course AZT Treatments for Pregnant HIV+ Women in Developing Countries

In the mid- and late-1990s, researchers ran trials testing a short course of AZT against a placebo in HIV+ pregnant women in Uganda, Thailand, and elsewhere. Although the longer 076 protocol, which was standard in wealthy nations, had been quite successful at preventing vertical HIV transmission, it was well out of reach as a standard treatment in developing countries, both because of its expense and because of the level of contact with the health care system that it required. The question was whether a more affordable and deliverable course of AZT, which no one thought would be as effective as the long course, would still have substantial benefits. The trials and the ethical dilemmas they pose have become well known and have fueled much of the current vigorous interest in PE, partly in virtue of an article in the *New England Journal of Medicine* by Marcia Angell (1997) in which she accused such trials of violating PE. The use of a placebo in such trials was unethical, she charged, because there was good reason to think that the short course of AZT would provide substantial benefits over a placebo. She appealed to the therapeutic obligations of physician-investigators to argue that researchers should test the short course against the full 076 protocol, rather than against a placebo. London (2000) pointed out a few years later that Angell's suggested alternative also violates PE, since there was good reason to think that the 076 protocol would be more effective than the short course. Hence PE does not seem to be a useful tool for thinking through the ethical contours of such

trials. Yet the case still dominates contemporary discussions of PE,[21] as it highlights the difficulty of applying the principle in any context in which resources are very scarce and the most effective treatments simply cannot be made available to the general population. In order to avoid ruling out research that enables investigators to discover locally realistic treatments for underserved populations, it seems, one needs to suspend PE at least temporarily, and acknowledge limits on physicians' therapeutic obligations to their research subjects. But one ought to be uncomfortable indeed with the conclusion that researchers should put the usual principles of research ethics on hold when conducting research in poverty-stricken areas.

On the other hand, PE*, by design, applies equally well in conditions of scarcity and of bounty. PE* was satisfied as long as researchers were in equipoise with respect to whether the short-course AZT treatment would provide enough benefit to be worth making it generally available to HIV+ pregnant women in the countries where the research was conducted, given local conditions and resources. And indeed, this is presumably the very question that was driving the research. Hence on this parameter, the research seems to have been ethically acceptable.[22] Furthermore, since it was not clear that there was *any* intervention that would be both effective in preventing vertical HIV transmission and realistically affordable for the local population, the use of a placebo control did not knowingly drop anyone below the *de jure* local standard of care—although now that the sizeable benefits of the short course and its relative ease of administration have been fairly well-established, the use of a placebo in such a trial would be unacceptable. Hence the trials met both of the criteria for ethical acceptability that I have proposed. (Of course, this is not to say that they did not perhaps suffer from other ethical problems that lie beyond the scope of this paper, such as problems with obtaining good quality informed consent across a major culture divide, or unmet duties of ancillary care.)

London (2001; 2006a) tries to salvage an ethical role for a version of traditional PE in the context of these trials. According to his "integrative approach," a principle of equipoise directs one to compare the relative expected benefits for individuals in different trial arms. However, he rejects a narrow, medical conception of "benefit," in favor of a contextual conception that takes into account not only the biological effect of an intervention, but how easy it is to administer, how burdensome it its for patients to follow through with the treatment, whether the treatment is culturally appropriate, and so forth. I applaud this expanded notion of

benefit, and I think that London shares much of my motivation to build the local context of a research project directly into our principles for determining its ethical acceptability. However, even while appealing to this richer notion of benefit, London's version of the principle of equipoise still follows PE in demanding uncertainty—or "credible doubt," in his words—with respect to the relative benefits *to the research participants* of the various trial arms (London 2001, p. 324). In other words, London's principle still would prohibit any trials in which there is good reason to expect participants in one trial arm to fare better than those in another, despite his broader and more nuanced view of what counts as faring better. But I already have argued that expected equality of trial arm outcomes is not, in and of itself, and important ethical consideration. PE* asks us instead to consider the larger question of the value of making an intervention accessible to a *population*, where such value certainly should not be measured in the narrowly medical or biological sense that London rightly rejects. Hence it seems that London still insists upon uncertainty in the wrong location.

This is especially striking given that London's long-term goal, in insisting upon any principle of equipoise, is deeply similar to mine. London (2001, p. 328) hopes that "although [he] cannot justify this assertion here, a possible implication" of his account is that "to pass the equipoise requirement international research initiatives will have to be coordinated with, or at least responsive to, a nation's larger public health initiatives and political needs." If this is his goal, then his route to achieving it is indirect at best, and—as his own language here insinuates—opaque at worst. In a later paper that functions in some ways as a sequel to the 2001 equipoise paper, London (2005, p. 33) writes:

> If clinical research is to be permissible, it must function as part of a division of labour in which the distinctive scientific and statistical methods of the research enterprise target and investigate the means of filling the gaps between the most important health needs in a community and the capacity of its social structures to meet them.

PE* accommodates this imperative more directly than London's version of PE, by insisting that research be designed to address uncertainty concerning how to address most appropriately the contextual health needs and entitlements of a specific population. If the purpose of research is to discover how best to fill the gaps London describes, then why not simply demand equipoise concerning how best to do this? Why insist upon

uncertainty with respect to the relative outcomes of trial arms at all?[23] It seems that the only reason to worry about the expected outcomes of trial arms is to make sure that no trial arm receives care that falls below the local *de jure* entitlement of the participants. But this is secured by the acceptance of MSP, which, as I have argued, is both an ethically necessary accompaniment to any principle of equipoise, and cannot be considered part of any such principle.

POTENTIAL EXCEPTIONS TO PE*

Imagine a case in which researchers wish to test the impact of an intervention, *T*, that poses minimal risk and offers the possibility of substantial benefits to those who receive it. Imagine also that *T* clearly is not worth making generally accessible—because it is too expensive, perhaps, or too difficult to administer. Testing *T* would violate PE*, as it is not an open question whether *T* ought to be accessible to the local population. However, it seems odd to rule out such a test as unethical. There is no harm done, after all, and even if the population at large will not benefit from *T*, at least some of the research participants will, along with whoever does end up having access to T. Why, the objection goes, should they be denied this benefit simply because it cannot be universalized? In response to this challenge, one must distinguish among reasons for running a test on *T* in the first place, given the impracticality of its general use. There are two types of reasons that come to mind. One might want to run the study in order to determine whether *T* is worth making available to some population *other* than that from which the subjects are drawn, or one might have a theoretical interest in the effects of *T*, even though it is known that it would not be appropriate to make it generally accessible in any community. I will treat these two possibilities in turn.

First, consider a case in which *T* is an intervention that might become part of the standard of care for *some* people, but not the people participating in the trials. Imagine, perhaps, that a resource-rich community is considering whether *T* ought to become the *de facto* standard of care, but researchers cannot drum up enough research volunteers in that community to run the study properly, so they recruit volunteers from another community with different local conditions, in which it is not plausible to think that *T* should become a standard intervention.[24] For instance, consider again the hospital stay example. It may be that investigators are unable to find volunteers in community *A* for a study in which one trial arm receives a shorter hospital stay, given that all members of that com-

munity are entitled to a perhaps unnecessary or cost-ineffective longer hospital stay. They might instead conduct the research in community *B*, where the typical hospital stays are shorter, comparing a trial arm receiving the longer stay typical of community *A* with the proposed shorter stay. It may appear that there is no harm in running such a test in community *B*, since the research will harm no one and perhaps benefit some, and the researchers will not knowingly be denying anyone care to which he or she is morally entitled, given a just distribution of local resources.

The first question to be asked about such a trial is whether it is scientifically sound, given that information about the outcome of an intervention in one set of local circumstances does not routinely generalize to knowledge about what the outcomes will be in another. In the hospital stay example, for instance, it is unlikely that the hospitals themselves, and the post-hospital home support available, would be sufficiently similar for this study to provide any helpfully transferable information. However, if the research really does generalize from the test group to the group who may have access to *T*, then *both* populations fall under the scope of the research. In this case, the experiment would not actually violate PE*, because the experiment indeed would be designed to answer an open question about the extent to which *T* should be made accessible to the population falling under the scope of the research. Thus the case would not serve as an exception to the principle.

Such research is surrounded with ethical pitfalls, and most of the reasons for running tests on a population other than the one that stands to benefit from the research are suspect. If the second group is willing to volunteer for the study only because their local *de facto* standard of care is so low that the study is their only means of access to services to which they should have access anyhow, such as basic health monitoring, then the research trial may well be exploiting participants, dropping them below their *de jure* standard of care and thereby violating MSP. Or, the differences between the two groups' willingness to volunteer might indicate problems with the informed consent process, or an offer of compensation that is problematically attractive to one group and not the other. But the fact that a trial may face such ethical problems while also satisfying PE* does not indicate a problem with PE* itself; it merely indicates that satisfying PE* should not be researchers' only ethical concern.

Second, consider a case in which the research on *T* is not directed at the question of appropriate access to *T* at all. Researchers may have a purely theoretical interest in *T*, or they might hope that understanding *T* will

lead to other advances concerning different, more useful treatments in the future. Such research clearly would violate PE*. Research that uses other people's bodies with no foreseeable benefit to the larger community is cause for caution, and investigators should take extra care to make sure that volunteers truly understand the speculative nature of the research. However, as long as there are no other ethical problems with such a study—e.g., nefarious motives for collecting the data, poor informed consent, drawing resources away from more pressing research, failing to satisfy MSP, and so forth—the mere fact that it violates PE* does not seem to render it unethical. The usual motive in placing constraints on human subjects research is the fear that participants will be harmed in some way in the name of that research—that they will be exposed to possible physical or psychological harm, treated disrespectfully or unjustly, denied care to which they ought to be entitled, and so forth. However, by stipulation, there is no risk of harm in this case. In the minority of cases in which the moral risks that PE* is designed to avoid are not present, the principle need not apply. In other words, in the fairly unusual circumstance in which an experiment offers the prospect of benefit to participants and poses no significant risk of harm—whether physical or moral—we can admit of exceptions to the usual ethical constraints on enrolling subjects in research trials, including PE*. In such a case, enrolling in the trial is morally similar to entering a raffle. Such a case seems to constitute a harmless exception to PE* that does not impugn its value as a general principle providing research participants with protection from moral harm.

WHOSE UNCERTAINTY?

Whereas Charles Fried (1974), in his original formulation of PE, requires *individual physician-researchers* to be in a state of equipoise concerning the expected outcomes of the trial arms in which their patients were participating, later writers, such as Eugene Passamani (1991) and Benjamin Freedman (1987), argue instead for the importance of what is usually called "clinical equipoise," or uncertainty within an expert community of clinicians. According to Passamani (1991, p. 1591), equipoise requires "a community of competent physicians who would be content to have their patients pursue any of the treatment strategies being tested in a randomized trial, since none of them has been clearly established as preferable." More recently, some authors have pointed out that there is no reason to think that physicians are the only ones who ought to have a say about which treatment is "preferable." Veatch (2002) argues persuasively that

particular patients' preferences and values are essential to determining the "preferability" of a treatment. Jason Karlawish and John Lantos (1997) point out that community values and cultural traditions help to determine which interventions are preferable in a particular local context. The upshot of both articles is that equipoise must emerge out of a conversation that includes laypeople and not just clinicians.

Given that I am not presupposing that research is led by physicians and geared toward clinical treatments, it would be inappropriate for me to index PE* to a community of competent physicians. More fundamentally, the question for which PE* demands equipoise—that is, whether a given intervention is worth making accessible to the population falling under the scope of the research—is not one that we would expect physicians to have the authority to answer on their own, even in the case of therapeutic research. Indeed, this is an inherently interdisciplinary question, and answering it may require the expertise of physicians, policymakers, economists, epidemiologists, anthropologists, geographers, patients, family caregivers, community leaders, program administrators, or social workers, among others. The narrowly medical risks and benefits of a particular intervention are only pieces of the complex picture that determines the overall value of making that intervention accessible to a community. Hence, like Karlawish and Lantos, I promote a form of community equipoise: the "we" who must be in a state of equipoise is neither an individual nor an insulated group of experts, but a complex mix of stakeholders with a variety of expertise.

Furthermore, these various parties are not routinely in conversation with one another—they form no ready-made "we." Thus, establishing the presence of the relevant kind of equipoise often will require a concerted effort to *forge* a conversation and a consensus among different kinds of experts, concerning whether there really is an open question to be answered about the appropriateness of making an intervention available to a community.[25] I suggest that some such interdisciplinary conversations normally should occur prior to embarking on research. This would help investigators to design research projects that are both ethical and maximally useful, capable of answering pressing questions about how a society should deliver care in the most direct manner possible.

Basic moral principles of justice, respect, and concern for welfare are sufficient to ground researchers' moral obligation not to pursue research at the cost of providing people the care to which they are entitled. Yet the principle of equipoise traditionally has been grounded in the special

obligations of physician-investigators to provide research participants with optimal care. This grounding has made the principle hard to apply in contexts with limited health resources, or to research that is not directed by physicians, or to nontherapeutic research. I have used socially situated, contextually variant notions of accessibility and entitlements to care, and widened the lens to include health services research, program evaluation, and other types of nontherapeutic, non-physician-driven research. These changes in starting point led me to shift the required location of equipoise, from the relative expected outcomes for individuals in different trial arms, to the appropriate, contextually situated standard of care. My version of the principle of equipoise does not depend upon an appeal to the Hippocratic duties of physicians, and it is designed to be applicable within a wide range of research contexts and types, including health services research and research on social interventions. Although PE and PE* converge as a society approaches ideal conditions in which resources are plentiful and access to care is nearly unlimited, PE* gives better guidance under the conditions of limited access to health care and resources that are typical for all but a privileged few.

This research for this project was supported by the Social Sciences and Humanities Research Council of Canada, the Greenwall Foundation, and the Food and Nutrition Assistance Research Division of the United States Department of Agriculture. I thank Richard Manning, Elizabeth Frazao, and Holly Taylor, as well as two anonymous referees, for helpful comments and suggestions.

NOTES

1. Prominent examples of this debate over just the last few years include attacks on the principle of equipoise in Chiong (2006) and Miller and Brody (2002), both of which were *American Journal of Bioethics* target articles followed by multiple responses by leading bioethicists, as well as major articles about the principle in each of the top bioethics journals, including Jansen (2005), London (2001; 2000), Menikoff (2003), and Miller and Weijer (2003). In this issue of the *Kennedy Institute of Ethics Journal*, Gifford (2007) responds to Miller and Weijer's (2003) defense of equipoise, and Miller and Weijer (2007) reply.
2. Those who propose any principle of equipoise, and thus demand some kind of uncertainty as an ethical condition upon research, must specify with whose uncertainty they are concerned. I discuss this issue in some detail toward the end of this paper. They also must specify what they mean by uncertainty. *How* open must a question be in order for those who pose it to

count as being in a state of equipoise? Answers can range from accepting nothing short of proof by randomized controlled trial as ending uncertainty (in which case randomized controlled trials would never be ruled out unless they were purely duplicative), to allowing any sort of preliminary data or anecdotal evidence that tilts the scales to count as ending equipoise (in which case almost no research would ever pass the equipoise test). There have been many discussions of the appropriate criteria for uncertainty; for very recent examples see Hansson (2006) and Halpern (2006). I will leave this debate to others. I simply assume a rough, intuitive middle path: one counts as knowing something when one has sufficient reason to believe it to justify one's acting on this belief without further confirmation. Of course, how certain one must be of something in order to be rational in acting on it in practice is entirely contextual and depends upon what is at stake in being right or wrong.

3. Classic discussions include Fried (1974), Freedman (1987), Hellman and Hellman (1991), and Passamani (1991). For recent discussions see Note 1. As far as I know, only Alex John London (2006b, p. 2880) has explicitly called into question the grounding of the principle of equipoise in the dual loyalties of the physician-researcher, seeking instead to ground it "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."
4. There have been two major kinds of attacks on the validity of PE in the last few years. Franklin Miller and Howard Brody (2002) are among those who have defended a "difference position," claiming that the impetus for PE has come from a conflation of the ethics of clinical medicine and the ethics of research. On their view, physician-investigators occupy two ethically distinct roles that ought to be sharply distinguished in practice, and in the context of research they have no special therapeutic obligations to their subjects. Thus, they conclude, there is no reason to accept any version of PE as an ethical constraint on research. Meanwhile, Winston Chiong (2006) recently has argued that although, contra Miller and Brody, research *does* need to be governed by the ethics of clinical medicine, physicians never had the kind of "uncompromisingly patient-centered" therapeutic duty that has been presumed in most discussions of equipoise. In fact, Chiong argues, physicians can and do make ethical tradeoffs that compromise their individual patients' care. Hence, he concludes, there is again no reason to accept PE as a constraint on research. But both of these arguments undercut PE only if one presupposes the grounding of the principle in the unlimited therapeutic duties of physician-investigators.

5. The 2000 "note of clarification" on this paragraph softens the rule, allowing some exceptions, but provides no theoretical justification for these exceptions. See *http://www.wma.net/e/policy/b3.htm* (accessed 15 October 2007) for the Declaration, including the note of clarification.
6. On the other hand, Veatch's point that one should take patient values and preferences seriously in establishing equipoise is an important corrective to most of the literature, and I shall return to it later in this discussion.
7. Whether the authors of the *Belmont Report* managed to capture and operationalize these three principles properly is a different issue and beyond the scope of this paper.
8. This is one way of filling out what Miller and Brody (2002, p. 6), among others, tried to capture when they suggest replacing PE with the "obligation not to exploit participants for the sake of scientific investigation."
9. As an anonymous referee emphasized, it is never possible to extricate political from material conditions completely. This is part of the murkiness of everyday life, and I do not see that it should slow down the attempt to determine local *de jure* standards and then to put them into practice. One may use various (imperfect) conceptual devices as aids in doing so, including, for instance, the quasi-Rawlsian formula that John Arras and Robert Crouch (1998) call the "liberal consensus view," according to which a just distribution of health care is one that "informed, rational, and prudent individuals would choose for themselves against a background entitlement to a fair share of their society's resources."
10. In his most recent articles, London (2001; 2006a & b) has pointed out that whether an intervention is *beneficial* to an individual overall depends not only upon its biological effects, but upon contextual features such as how convenient it is to go through with the intervention, how much support is available to patients during the treatment, whether ongoing access to follow-up care will be needed, and so forth. Later in this discussion, I endorse this broadened sense of "benefit." But notice that even if an intervention is beneficial *to individuals* in this broader, richer sense, this does not prove that it ought to be part of the local standard of care—for instance, such care might be spectacularly expensive, and making it generally accessible might take away from other services that are yet more beneficial.
11. For just one example among many, see London (2001).
12. Arras and Crouch (1998) defend a standard very close to the one I defend here, by insisting that participants are entitled to the "best" treatment, where the "best" does not mean the most therapeutically efficacious, as it presumably does in the Declaration of Helsinki, but rather the most appropriate and

just treatment, given the context. My MSP is slightly weaker, since it insists only on the care to which participants are *entitled*, given the local resources and context, and not necessarily to the care that is the *best* in that local situation.

13. One cannot say that the research should be geared toward improving the health and welfare of the test subjects themselves, of course. Not all research is therapeutic, and much of it is directed at helping future generations, sick family members of healthy subjects, and so forth.
14. What counts as accessibility is also context-dependent. A purely voluntary cosmetic surgery procedure probably counts as accessible as long as it is safely available to interested buyers who can afford it. At the other end of the spectrum, something like contraception might count as accessible only once a society offers it for free, provides convenient and confidential access to it, and does community outreach to encourage people take advantage of it.
15. Remember that advocates of PE invariably also believe that no trial arm should receive what is known to be less than the best available care.
16. If we are not sure whether an intervention is beneficial *at all*, then of course cutting it does not necessarily *lower* the standard of care, but might instead raise it. Now that we know that episiotomies do more harm than good, for instance, eliminating them from routine vaginal delivery practice should be counted as an improvement rather than a reduction in the standard of care.
17. Toroyan and colleagues do not cite the study itself; it appears that they ran the study, and that their background is in education.
18. Indeed, there appears to have been only one such study, conducted in 1985, and it focused only the effects of WIC on the children of women who smoked during pregnancy (Metcoff et al. 1985). Reports of observational studies on WIC's effects on health include Devaney, Bilheimer, and Schore (1992); Hamilton, Fox, and Biing (2004); and Oliveira (2002).
19. Some conservative commentators have suggested just that, insisting that until WIC's benefits have been proven in a randomized controlled trial, we have no reason *at all* to believe in them. Having examined the data that is available, and knowing how little of our actual working "knowledge" in fact derives from randomized controlled trials, I find this view sufficiently implausible to warrant no further consideration.
20. The U.S. federal code governing human subjects research, Title 45 Part 46 of the "Common Rule," exempts research on benefits programs from the normal informed consent requirements. If my argument is correct, then regardless of

this legal loophole, informed consent is *ethically* required for participation in such a trial.

21. See, for instance, London (2006a), Chiong (2006), and the many replies to Chiong in the same issue of the *American Journal of Bioethics*.
22. Although, as an anonymous referee reminded me, it is not clear that even the short course regiment realistically could have been made widely available in most developing nations.
23. Significantly, despite his promissory note in the 2001 paper, in this later paper London does not appeal to the notion of equipoise at all, except in referring back to the first paper (London 2005, note 46).
24. Thanks to Richard Manning for suggesting this form of example.
25. This point was made to me by Holly Taylor in private conversation.

REFERENCES

Angell, Marcia. 1997. The Ethics of Clinical Research in the Third World. *New England Journal of Medicine* 337: 847–49.

Arras, John, and Crouch, Robert. 1998. AZT Trials and Tribulations. *Hastings Center Report* 28 (6): 26–34.

Besharov, Douglas J., and Germanis, Peter. 2001. *Rethinking WIC: An Evaluation of the Women, Infants, and Children Program*. Washington, DC: AEI Press.

Chiong, Winston. 2006. The Real Problem with Equipoise. *American Journal of Bioethics* 6 (4): 37–47.

Devaney, B.; Bilheimer, L.; and Schore, J. 1992. Medicaid Costs and Birth Outcomes: The Effects of Prenatal WIC Participation and the Use of Prenatal Care. *Journal of Policy Analysis and Management* 11 (4): 573–92.

Edwards, S. J., and Kirchin, S. 2002. Rationing, Randomizing, and Researching in Health Care Provision. *Journal of Medical Ethics* 28 (1): 20–23.

Freedman, Benjamin. 1987. Equipoise and the Ethics of Clinical Research. *New England Journal of Medicine* 317: 141–45.

Fried, Charles. 1974. *Medical Experimentation: Personal Integrity and Social Policy*. New York: Elsevier.

GAO. General Accounting Office, United States. 1992. *Early Intervention: Federal Investments Like WIC Can Produce Savings*. Washington, DC: General Accounting Office.

Gifford, Fred. 2007. Pulling the Plug on Clinical Equipoise: A Critique of Miller and Weijer. *Kennedy Institute of Ethics Journal* 17: 203–26.

Halpern, Scott D. 2006. Evidence-Based Equipoise and Research Responsiveness. *American Journal of Bioethics* 6 (4): 1–4.

Hamilton, W.; Fox, M. K.; and Biing, H.L. 2004. *Effects of Food Assistance and Nutrition Programs on Nutrition and Health, Vol. 3: Literature Review*. Washington, DC: Abt Associates.

Hansson, Sven Ove. 2006. Uncertainty and the Ethics of Clinical Trials. *Theoretical Medicine and Bioethics* 27: 149–57.

Hellman, Samuel, and Hellman, Deborah S. 1991. Of Mice but Not Men: Problems of the Randomized Clinical Trial. *New England Journal of Medicine* 324: 1585–89.

Jansen, Lynn A. 2005. A Closer Look at the Bad Deal Trial: Beyond Clinical Equipoise. *Hastings Center Report* 35 (5): 29–36.

Karlawish, Jason H., and Lantos, John. 1997. Community Equipoise and the Architecture of Clinical Research. *Cambridge Quarterly of Healthcare Ethics* 6: 385–96.

Levine, Robert J. 1986. *Ethics and Regulation of Clinical Research*. 2d ed. Baltimore, MD: Urban & Schwarzenberg.

London, Alex John. 2000. The Ambiguity and the Exigency: Clarifying "Standard of Care" Arguments in International Research. *Journal of Medicine and Philosophy* 25: 379–97.

———. 2001. Equipoise and International Human-Subjects Research. *Bioethics* 15: 312–32.

———. 2005. Justice and the Human Development Approach to International Research. *Hastings Center Report* 35 (1): 24–37.

———. 2006a. The Moral Foundations of Equipoise and Its Role in International Research. *American Journal of Bioethics* 6 (4): 48–51.

———. 2006b. Reasonable Risks in Clinical Research: A Critique and a Proposal for the Integrative Approach. *Statistics in Medicine* 25: 2869–85.

Menikoff, Jerry. 2003. Equipoise: Beyond Rehabilitation? *Kennedy Institute of Ethics Journal* 13: 347–51.

Metcoff, J.; Costiloe, P.; Crosby, W. M.; et al. 1985. Effect of Food Supplementation (WIC) During Pregnancy on Birth Weight. *American Journal of Clinical Nutrition* 41: 933–47.

Miller, Franklin G., and Brody, Howard. 2002. What Makes Placebo-Controlled Trials Unethical? *American Journal of Bioethics* 2 (2): 3–9.

Miller, Paul B., and Weijer, Charles. 2003. Rehabilitating Equipoise. *Kennedy Institute of Ethics Journal* 13: 93–118.

———. 2007. Revisiting Equipoise: A Response to Gifford. *Kennedy Institute of Ethics Journal* 17: 227–46.

Oliveira, Victor J. 2002. *The WIC Program: Background, Trends, and Issues*. Washington, DC: U.S. Department of Agriculture, Economic Research Service.

Passamani, Eugene. 1991. Clinical Trials—Are They Ethical? *New England Journal of Medicine* 324: 1584–92.

Rossi, Peter H. 1998. *Feeding the Poor: Assessing Federal Food Aid*. Washington, DC: AEI Press.

———, and Hamilton, William. 2002. *Effects of Food Assistance and Nutrition Programs on Nutrition and Health*. Washington, DC: U.S. Department of Agriculture, Economic Research Service, Food Assistance & Nutrition Research Program.

Steinberg, David. 2002. Clinical Research Should Not Be Permitted to Escape the Ethical Orbit of Clinical Care. *American Journal of Bioethics* 2 (2): 27–28.

Toroyan, Tami; Roberts, Ian; and Oakley, Anne. 2000. Randomisation and Resource Allocation: A Missed Opportunity for Evaluating Health Care and Social Interventions. *Journal of Medical Ethics* 26 (5): 319–22.

Veatch, Robert M. 2002. Indifference of Subjects: An Alternative to Equipoise in Randomized Clinical Trials. *Social Philosophy and Policy* 19: 295–323.

Weijer, Charles. 2002. When Argument Fails. *American Journal of Bioethics* 2 (2): 10–11.

Contributors

Robert B. Baker, Ph.D., is Director and Professor of Bioethics, The Union Graduate College-Mount Sinai School of Medicine Bioethics Program, and William D. Williams Professor of Philosophy, Union College, Schenctady, NY.

Marion Danis, M.D., is Head, Section on Ethics and Health Policy, Department of Bioethics, and Chief of the Clinical Center's Bioethics Consultation Service, National Institutes of Health, Bethesda, MD.

Fred Gifford, Ph.D., is Professor of Philosophy, Michigan State University, East Lansing.

Samia A. Hurst, M.D., is Maître Assistante, Institute for Biomedical Ethics, University of Geneva Medical School, Switzerland.

Rebecca Kukla, Ph.D., is Professor of Philosophy and Internal Medicine, University of South Florida, Tampa.

Laurence B. McCullough, Ph.D., is Professor of Medicine and Medical Ethics, Center for Medical Ethics and Health Policy, and a Faculty Associate, Huffington Center on Aging, Baylor College of Medicine, Houston, TX.

Paul B. Miller, J.D., Ph.D.(cand.) is Law Clerk to the Honorable Justice Ian Binnie of the Supreme Court of Canada.

Peter A. Ubel, M.D., is Professor, Division of General Internal Medicine; Director, The Center for Behavioral and Decision Sciences in Medicine; and Adjunct Professor, Department of Psychology, University of Michigan, Ann Arbor; and Investigator, HSR&D Ann Arbor Center of Excellence, Department of Veterans Affairs, Ann Arbor, MI.

Charles Weijer, M.D., Ph.D., is Associate Professor and Canada Research Chair in Bioethics, Department of Philosophy, University of Western Ontario, London, Ontario, Canada.