

Uncertainty, Bias, and Equipoise: A New (Old) Approach to the Ethics of
Clinical Research

Michael Goldsby and William P. Kabasenche

michael.goldsby@wsu.edu wkabasenche@wsu.edu

School of Politics, Philosophy, and Public Affairs

Washington State University

Abstract

The concept of equipoise is considered by many to be part of the ethical justification for using human subjects in clinical research. In general, equipoise indicates some uncertainty about the relative merits of the experimental intervention compared to existing treatments. Relieving this uncertainty gives scientific value to an experiment, thereby making the risks to human subjects in the trial acceptable, other considerations notwithstanding. But how to characterize equipoise remains controversial since Freedman's ground-breaking publication on the subject. We offer a new account of equipoise, which draws on and extends an option Freedman discarded. After establishing the importance of some account of equipoise as part of ethically justified clinical trials, we revisit Freedman's distinction between clinical and theoretical equipoise. We raise concerns about Freedman's preferred clinical equipoise and then rehabilitate theoretical equipoise. In particular, we use a variety of arguments from epistemology to show how Freedman was too hasty in rejecting theoretical equipoise. In addition, we argue that theoretical equipoise is a subspecies of epistemic equipoise, which we characterize as a form of uncertainty that is the result of the possibility of error. This type of uncertainty can best be alleviated by research, which produces beliefs supported by strong statistical evidence, which is a key aim of clinical trials. Further, this type of uncertainty can explain why even clinicians with more first hand experience with an intervention than their peers, and who may not be in theoretical equipoise, could still justifiably support a trial designed to ameliorate the possibility of error due to cognitive bias.

Introduction

Although there seems to be a broad consensus that some kind of equipoise is necessary for the ethical conduct of biomedical research using human subjects, commentators continue to debate how best to characterize the requisite equipoise. Since Freedman's original, path-breaking publication (1987), which itself attempted to summarize a number of previous conceptions of equipoise before arguing for clinical equipoise as filling the needed role, there have been continuing discussions of how to understand this critical concept. Some (e.g., Miller and Brody, 2003) have attempted to move the debate away from clinical equipoise in favor of other ways of supporting ethical biomedical research. Others (e.g., Ubel and Silbergleit, 2011) have attempted to refine Freedman's original concept.

In this article we argue for epistemic equipoise, of which theoretical equipoise is a special case, as filling the needed role, an option Freedman considered and discarded. We begin by articulating the connections between some form of equipoise and the ethical conduct of clinical research using human subjects. In the second section we challenge the place of *clinical* equipoise to provide the ethical support for such research and argue that *theoretical* equipoise is more robust than Freedman thought. In Section 3, we use the expanded notion of epistemic equipoise to account for the ethical justification of clinical trials when theoretical equipoise has been displaced. We argue that epistemic equipoise is characterized by a healthy respect for cognitive bias and show that this is sufficient to do what clinical equipoise cannot, namely, characterize the uncertainty that makes a particular research trial scientifically and clinically valuable.

1. Equipoise and the Ethics of Clinical Trials

When biomedical researchers ask individuals to participate in a research trial, this raises the possibility that participants may receive a less effective treatment or even be harmed by the experimental intervention. Researchers ask subjects to take on this risk in order to know how better to treat future patients of a similar type. Thus, it uses some people for the benefit of others, and therefore biomedical research involving human subjects requires ethical justification. The notion of equipoise has been considered for some time to be an important part of the ethical justification of biomedical research (Freedman 1987). It requires, in general, some kind of

uncertainty on the part of researchers about the relative effectiveness (or safety) of a particular intervention. This uncertainty means that a clinician-researcher is not in a position to endorse, in a sense we hope to further characterize in this paper, the claim that intervention A is better than either a placebo or intervention B. Given the uncertainty, the clinician-researcher can be justified in enrolling subjects in a trial with some confidence that in doing so he or she is not depriving them of the best treatment.

We will sometimes use the term 'clinician-researcher.' But Miller and Brody (2003) might object to our attempts to further clarify equipoise. They have argued that the worlds of clinical care and research are different and have different ethical norms because they have different practical aims. Clinical therapy aims to heal or more generally to provide effective therapies, and research aims to discover information of scientific and clinical value. Thus, they argue, the basic ethical responsibility of clinicians is to provide therapies that are effective and appropriate for a given patient, and the basic ethical responsibility of researchers is to ensure that they do not expose subjects to risks in "valueless research" (2003, p. 26). Thus, Miller and Brody seem to us to replace the function of clinical equipoise with the notion of scientifically valuable research. But they say little about the concept of scientific value, and we hope to clarify that concept by way of a renewed effort to articulate the kind of equipoise that serves as a part of the ethical justification of biomedical research. Scientifically valuable research ought to be research that reduces the uncertainty clinicians and patients face when considering new treatment options. This uncertainty is central to our notion of equipoise.

We suspect that Miller and Brody's account either fails to appreciate the connections between biomedical research and clinical care (which is unlikely given their vast experience in those areas) or it implicitly assumes connections between the two arenas without fully acknowledging the interdependence between them. What a physician in the clinic is ethically permitted or obligated to offer appropriate patients depends in large measure on the current state of understanding and evidence in biomedical research. And what researchers are permitted or obligated to investigate in a study depends, to some degree, on the current knowledge gaps in clinical health care. Thus, though the two arenas are different and though there may be few true clinician-researchers, as opposed to pure clinicians or pure researchers, the epistemic state of the research community bears on the

ethical obligations of the clinician and vice versa. The term 'clinician-researcher' is meant to capture this. Despite Miller and Brody's objections to the concept of equipoise, we believe this project has value. The clinical and research realms are linked in important ways; indeed, the scientific value of research is linked to its clinical relevance.

We do agree with Miller and Brody that *clinical* equipoise is problematic. Their reasons sometimes seem to be centered on the suggestion that the word "clinical" implies that research should have therapeutic intent. As noted above, research is different from therapy and has different, though interrelated, aims. We agree with this concern, and we will also offer an additional critique of clinical equipoise (in Section 2 below). However, whereas Miller and Brody are content to point to scientific value as part of the ethical basis for doing research that uses human subjects, we hope to further clarify the kind of uncertainty that makes a study scientifically, and ultimately clinically, valuable.¹

The ethics of equipoise affects a number of different kinds of practical decisions, further requiring us to keep a connection between these different arenas in mind. For a researcher, the value of research she or he might do depends on what is currently uncertain. As Miller and Brody point out, it would be unethical to do research that exposes subjects to risk in order to find out some piece of information that is of little to no value. Equipoise is meant to capture one kind of value in research. Thus, despite Miller and Brody's objection, we believe the concept of equipoise remains useful, and our paper aims to show its continuing value for the ethical justification of biomedical research. In the next section, we aim to recast the requisite uncertainty in terms of theoretical or epistemic equipoise, as opposed to the clinical variety endorsed by Freedman.

For the clinician the question of whether to suggest or recommend that a patient enroll in a study as a subject should depend, at least in part, on the state of uncertainty that currently exists about potential treatments. It would be at least ethically suspect for a clinician to recommend that a patient enter a study unless there was sufficient uncertainty that the current treatment is more effective than the experimental one. Similarly, for patients, the decision to enter a study would depend on the current state of uncertainty regarding their own treatment. Proper self-regard would most likely prevent them from entering a study if doing so might deprive them of a needed therapy already known to be effective.

The classic concern about the claims made above is that it is not clear that everyone shares the same level of uncertainty. Some clinicians, perhaps in virtue of their clinical experience or their understanding of the basic science or by some other source of insight, will have knowledge such that for them there is no equipoise (Hellman and Hellman 1991). This experienced clinician-researcher, who admittedly has more evidence at his or her disposal, has an obligation to recommend treatments he or she knows to be effective. That there is uncertainty in the larger medical community has no bearing on the experienced clinician-researcher because this researcher is in a position of having the best, or at least better, available evidence. Given this researcher's position, what reason might he or she have to support or to call for a trial on a given intervention? This researcher is not in theoretical equipoise, on Freedman's account of it. So that cannot be used to justify the experienced clinician-researcher's support of the trial. Nonetheless, the clinician-researcher might be in clinical equipoise, assuming there is disagreement in the medical community. However, it is unclear to us how mere disagreement between community members affects this clinician-researcher's position with respect to his or her patient's care or participation in the proposed trial. There are, as we point out in Section 2, different approaches to disagreement. These differences bear on the question of what the experienced clinician-researcher is in a position to recommend to his or her patients about entering trials. We address the question of what grounds the experienced clinician-researcher might still have to support a trial in Section 3 of this paper. However, before doing that, we critique Freedman's account of clinical equipoise and argue for the value of theoretical and epistemic equipoise for justifying biomedical research using human subjects.

2. Freedman got it Backwards

Freedman (1987) distinguishes between two types of equipoise - theoretical and clinical. He further argues that many of the problems that have been raised against equipoise arise precisely because theoretical equipoise is a problematic concept that cannot ethically justify randomized clinical trials. Consequently, he claims that clinical equipoise, which he thinks is relatively unproblematic, can succeed where theoretical equipoise fails. The contention that we will explore in this section is that Freedman was right to suggest that, between clinical and theoretical equipoise, one of them was

conceptually troubled and incapable of playing a part in randomized clinical trials. However, we further contend that he picked the wrong horse. We endeavor to show that his criticisms of theoretical equipoise miss the mark, due to a conceptual misunderstanding on his part, but first we will show that it is actually clinical equipoise that is conceptually troubled. Additionally, we will look to recent work in the epistemology of disagreement to show that even if we get an untroubled version of clinical equipoise, it is at least controversial as to whether clinical equipoise is ethically relevant to the individual clinician-researcher's decision. Finally, we intend to argue that theoretical equipoise is both conceptually stable and more ethically relevant.

2.1 Worries about Clinical Equipoise

Clinical equipoise, according to Freedman, "is a situation in which there exists (or is pending) an honest disagreement in the expert clinical community regarding the comparative merits of two or more forms of treatment for a given condition" (1992, p. 231). Freedman maintains that it is just this sort of disagreement that licenses clinical trials even when the clinician is not in theoretical equipoise. With regard to the resulting disagreement, it is his opinion that the results of something like Phase III clinical trials would be sufficient to settle any such disputes.²

One immediate worry concerns what we might call an "honest disagreement." Is one dissenting clinician working on the fringe of legitimate medical practice sufficient enough for an honest disagreement? We don't think so, and we also think that to push such a line of criticism against Freedman is tantamount to straw-manning his position. He does not intend by disagreement the failure of the *entire* medical community to act unanimously. Instead it seems that he requires dissent from at least some significant minority of the medical community before it would count as an honest disagreement. Nonetheless, it is a bit worrisome that he gives no guidance as to how many dissenters it would take to have an honest disagreement. This suggests some mild conceptual trouble in identifying exactly when we have an "honest disagreement" within the medical community. This is problematic since it is the disagreement that is supposed to license randomized clinical trials. However, if we are unable to identify when such a disagreement obtains, then it is difficult to see how we can know if we are permitted to conduct said trials, particularly when we are in a situation where it is unclear that there is significant disagreement.

While we hold that it is difficult at best to determine when an honest disagreement obtains, we are more concerned with how Freedman thinks such a disagreement can function as part of the justification for clinical trials. The aim of clinical trials, according to Freedman, is to resolve the disagreement (1987, 144). So, it is the resolution of the disagreement that is supposed to play some part in the ethical justification of clinical trials. There are two ways that we might understand this claim, and it seems that Freedman straddles both. Recall that the problem that clinical equipoise is supposed to address is the case of the experienced clinician-researcher, who is not in theoretical equipoise. This clinician-researcher, as the story goes, has enough evidence to determine that one treatment arm is better than another, and furthermore, this evidence is not available to the larger medical community. One way that clinical trials might resolve the dispute is to make the evidence that the clinician-researcher has more public. We admit that doing so is good, but it seems to ignore the impact of the problem. If the experienced clinician-researcher is subjecting his or her patients to randomized clinical trials where some will receive what he or she believes to be an inferior and possibly harmful treatment (given the evidence available to the researcher), then the experienced clinician-researcher is exposing those patients knowingly to harm for the benefit of others. We take this worry seriously, as we have stated above, and we are unsure how this is supposed to ease our worries.

The other way we might understand clinical trials as resolving the disagreement is that the disagreement, itself, constitutes some evidence that the experienced clinician-researcher is mistaken. As such the experienced clinician-researcher can be ethically justified in conducting a clinical trial because he or she is not in the epistemic position to recommend clinical trials given the disagreement and its evidential bearing. We are not so sanguine, however, about the ability of clinical equipoise to justify clinical trials, when theoretical equipoise has been displaced for an investigator. We base our findings on the relatively recent work done on the epistemology of disagreement. Most of this literature is new, and Freedman ought not be blamed for failing to consider this work. However, we think that this literature brings significant worries to light, especially as it relates to the ethics of clinical trials.

There are two prevailing theories regarding the rational response to disagreement: the right reasons view and the equal weight view. We will argue that each view fails to deliver a favorable verdict in support of clinical

trials. These two theories generally look at cases where both parties to the disagreement have access to the same body of evidence. As such, there is an obvious disanalogy between the cases considered in the disagreement literature and the case of the experienced clinician-researcher as the latter case involves a situation where one party is privy to evidence to which the other party does not have access. We, nonetheless, think it is instructive to look at how disagreement *might* be used to justify clinical trials.

The first theory that we will consider is *right reasons view*, which is defended by Thomas Kelly (2005) and Michael Titelbaum (forthcoming). According to the right reasons view, when two epistemic peers³ are considering the same evidence and come to different conclusions with regard to a particular claim, one of the two of them has made a mistake in reasoning. As such, the one who made the mistake (if that person wants to remain rational) must adjust his or her credence⁴ to be more in conformity with the credence of the other party. In other words, the one who got it right must "stick to his or her guns" and the one who got it wrong should adjust his or her credence (with regard to the claim under dispute) to be more in line with the correct reasoner. For example, suppose both Mike and Bill are working on proving a particular theorem in first order predicate logic and that the proof is a particularly tedious one requiring over 300 lines to complete. Further suppose that both Mike and Bill are competent, yet sometimes fallible, logicians, who are equally reliable (say 99%) at completing such proofs, and that they share roughly the same epistemic virtues. Mike, who has dropped a tilde in his proof, comes to the conclusion that the theorem cannot be proven. Bill, on the other hand, has made no such error, and completes the proof with (relative) ease. Thus, Bill comes to the opposite conclusion. Here we have a case of disagreement, but since Mike is in error he should adjust his credence that the theorem can be proven to be more in line with Bill's credence. Bill, on the other hand, is required to stick to his guns according to the right reasons view.

If the right reasons view is the correct way to rationally respond to disagreement about whether a particular treatment arm is better than another, then it is hard to see how one could cite said disagreement in the ethical justification for conducting a clinical trial. With the standard case of the experienced clinician-researcher, we are supposing that the experienced clinician-researcher is in the best position to know whether one treatment arm is better than another, since the researcher in question has access to more evidence than the rest of the medical community. Furthermore, what makes

the case interesting is the assumption that the experienced clinician-researcher has reasoned correctly about that evidence. If that is the case, then the experienced clinician-researcher has no reason to adjust his or her credence regarding the relative merits of the treatments under consideration, according to the right reasons view. Furthermore, the rest of the medical community is required by the norms of rationality to adjust their credences to be more in line with the experienced clinician-researcher. As such, they might be as morally required to adopt the treatment strategies endorsed by the experienced clinician-researcher without conducting any trial.⁵ If, however, we were to relax the assumption that the experienced clinician-researcher is the one who has reasoned correctly, the disagreement could still not be cited as part of the ethical justification for conducting a clinical trial. In such a case, the experienced clinician-researcher would be the one rationally required to adjust her credence regarding the medical treatments. Once again, no clinical trial is necessary.⁶ Thus, we maintain that if the right reasons view is the correct way to rationally respond to disagreement, then clinical equipoise cannot be used in an ethical justification of clinical trials.

There is some disagreement, however, concerning the rational response to disagreement. The main competitor to the right reasons view is the equal weight view endorsed by Richard Feldman (2007) and Adam Elga (2007). The equal weight view states roughly that when two epistemic peers who have access to *exactly* the same body of evidence disagree based on that body of evidence about whether a particular claim is true, then the constraints of rationality require that both parties should adjust their respective credences such that they meet somewhere in the middle.⁷ By way of example, consider two bright philosophy students, Carri and Sonia. Carri and Sonia are epistemic peers who have decided to take a class on the philosophy of religion. Over the course of the semester, both students carefully consider all the evidence and arguments (presented in class) for and against the existence of God. Additionally, they take the time to discuss any previous evidence and arguments that they may have had regarding God's existence. As a result, Carri and Sonia end up with the exact same body of evidence related to the existence of a supreme being. However, upon careful reflection Carri judges that the existence of God is highly unlikely, whereas Sonia comes to the opposite conclusion namely that existence of God is highly likely. According to the equal weight view, both students are rationally required to

adjust their credences such that they meet in the middle, where perhaps they hold that each possibility is equally likely.

At first glance, it seems that the equal weight view might support the claim that disagreement within the medical community can play some part in the ethical justification of clinical trials. One might argue that if there is disagreement about whether a particular treatment plan is the best option for treating a patient, then according to the equal weight view, we ought to suspend judgment (if we wish to remain rational) about whether said treatment plan is actually best, even if one of the parties to the disagreement is the experienced clinician-researcher. Additionally, if we are required to suspend judgment about which treatment plan is best, then rationality requires that we gather more evidence by conducting clinical trials. From there one might conclude that given the requirements of rationality and the equal weight view, disagreement can play some part in the ethical justification of clinical trials.

The problem with that line of reasoning is that it is false that the equal weight view requires that the experienced clinician-researcher adjust his or her credence regarding which treatment is best. That is because the equal weight view only applies to cases of disagreement between epistemic and evidential peers. The case that is supposed to create a problem for equipose assumes that the experienced clinician-researcher has significantly *more* evidence regarding the treatment options than the rest of the medical community. That is to say that the experienced clinician-researcher, has no evidential peers. Thus, while the equal weight view looked promising for the prospects of clinical equipose, it cannot be applied to the case of the experienced clinician-researcher. Furthermore, it means that our experienced clinician-researcher has no rational obligation to shift his or her credence, and as such it would be disingenuous of this individual to use clinical equipose as a justification for giving the patient what is thought to be an inferior treatment.

We think that the challenges that face clinical equipose are legion. First of all, it is unclear what constitutes an honest disagreement within the medical community. However, what is more troubling, as we think we have shown, is that even in those cases where it is clear that there is significant disagreement regarding the treatment options, mere disagreement is not enough to support the ethical justification of clinical trials, when one of the parties to the disagreement is the experienced clinician-researcher. Since clinical equipose was proposed specifically to deal with

such cases, we think the prognosis for clinical equipoise is grim. For that reason, we think it might be a good idea to dust off the notion of theoretical equipoise in order to get a better handle on the issue.

2.2 A Defense of Theoretical Equipoise

The intuitive idea behind theoretical equipoise is that there is insufficient evidence for the clinician-researcher to determine with reasonable confidence whether one treatment arm is better than the other. If that is the case, then theoretical equipoise could help us to avoid the worries raised in Section 1, because the physician will have no reason to prefer treatment A to treatment B. Yet, the problems with theoretical equipoise, according to Freedman, are largely twofold: (1) it is too "fragile" to support a protracted clinical trial; and (2) it cannot be used with complex treatment hypotheses.

Looking at the second issue (2) first, Freedman's point is that when we consider treatment options, we do not pay attention only to effectiveness. Other factors might also play a role. When a clinician decides upon a particular treatment plan, he or she does so with an eye to how well the treatment will address the condition, but the clinician must also take into account what sorts of side effects may accompany the treatment. For example, a particular drug might treat a given condition, but it might also put some stress on the liver. If the patient has a history of liver problems, the physician might choose to use a less effective treatment option that is not as harmful to the liver. Freedman expresses doubt as to whether theoretical equipoise can handle messy hypotheses that deal with such complexities.

We do not deny that clinical decisions are often complicated and messy in this sense. The point of clinical trials is to gather data so that clinicians can make such tailored decisions regarding their patient's care. A clinician takes the results of those trials and forms a hypothesis on what the best treatment option for his or her patient is. What is mysterious, however, is why one might think clinical equipoise is better suited for dealing with such vagaries. According to Freedman, if we are in a state of clinical equipoise, then that means that there is disagreement about a particular treatment plan for a particular patient, because we are unsure about its relative merits and drawbacks with regard to that patient. Why think that it is any more difficult to be in theoretical equipoise with regard to a treatment plan for a particular patient? The whole point of equipoise is that it allows us to recommend that *this particular patient*

participate in a clinical trial rather than recommend either treatment. As such the hypotheses in question must be relativized to a patient in order for equipoise to do its work, whether it be clinical, theoretical, or otherwise. Thus, we do not see this particular problem as a serious challenge to theoretical equipoise, if it is likewise unproblematic for clinical equipoise.

With regard to the first issue (1), Freedman believes that theoretical equipoise is too fragile because he believes that it can only be achieved "when, overall, the evidence on behalf of the two alternative treatment regimens is exactly balanced" (1987, p. 143). He goes on to say that theoretical equipoise is "balanced on a knife's edge" as any result whatsoever will dislodge the clinician's credence from that balance (*Ibid.*). What's the reasoning behind this claim? It follows from a picture of what is going on in clinical trials that he borrows from Chalmers (1978).⁸ In order to illustrate this view, let's consider two rival hypotheses H_A and H_B :

(H_A) treatment arm A is better than treatment arm B (for patient X);
and

(H_B) treatment arm B is better than treatment arm A (for patient X).

Let E be the total amount of evidence available to the investigator at any point in the trial and let $\Pr(H_i|E)$ represent the investigator's credence in the hypothesis H_i given the total evidence. According to the Freedman/Chalmers picture of equipoise, we would be in a state of theoretical equipoise between treatments A and B just in case equation 2.1 is true:⁹

$$\Pr(H_A|E) = \Pr(H_B|E) = 0.5 \quad (2.1)$$

The problem, as both Freedman and Chalmers correctly point out, is that (2.1) is rarely if ever true. Even worse, suppose that against all odds (2.1) holds at the very beginning of a trial. It is extremely unlikely that it will hold in the face of any result that comes from the experimental group of the study as our credence in H_A will either marginally increase or marginally decrease in time with those results. Thus, even in the unlikely event that we start the trial in theoretical equipoise as defined by Freedman, it evaporates quickly as the first results start to roll in and as a consequence it no longer offers any justification to continue the trial.

This is bad news, indeed, but only if the Freedman/Chalmers picture of theoretical equipoise is accurate. Our project is to show that this picture is *not* accurate. There are at least two major problems that we see with this picture. First of all, such a picture seems to assume that the above-mentioned hypotheses are both mutually exclusive and exhaustive. It is true that when we have two mutually exclusive and exhaustive hypotheses, the only way in which both hypotheses might be equally supported by the evidence is when equation (2.1) holds. For example, suppose we have two hypotheses concerning the result of a fair die when it is tossed. The first hypothesis is that the result will be even and the second hypothesis is that it will be odd. In this case, if it is reported to us that the result is less than 5, then the evidence supports both hypotheses equally, and our respective credences in both hypotheses should be around 50 percent. If however, it is reported to us that the result is less than four, then the hypothesis that the result is odd enjoys twice the support ($2/3$) of the rival even hypothesis ($1/3$).

One might think that such a picture is an accurate picture of clinical trials. For example, we might be investigating whether patients with osteoarthritis are better treated with arthroscopic surgery and physical therapy than treating them with physical therapy alone.¹⁰ In such a case both treatment arms hold physical therapy as a constant, and they only differ over whether we also include arthroscopic surgery as part of the treatment plan (treatment arm A) or not (treatment arm B). Thus, in such an experimental setup, one might think that the two hypotheses are mutually exclusive and exhaustive. The problem is that they are not. Recall that H_A is that treatment arm A is better than treatment arm B; whereas H_B is the hypothesis that treatment arm B is better than treatment arm A. There is, however, a third logical possibility, namely that neither treatment arm is better than the other.¹¹

Our point is that when there are more than two logical possibilities then two hypotheses might receive equal evidential support at some level well below 50-50. For example, if our only evidence is that a card drawn from a standard poker deck is a jack, then our confidence that it is a jack-of-hearts (25%) should be equal to our confidence that it's a jack-of-spades (25%). Each hypothesis is equally supported by the evidence that the card is a jack, although at a level considerably lower than 50-50. So, when we return to hypotheses like H_A and H_B , even if we grant that they are equally supported

by the evidence, we need not assign each a credence of 0.5. That is because a third hypothesis is needed to exhaust the logical possibilities, namely H_C :

(H_C) treatment arm A is no better than treatment arm B (for patient X)
and treatment arm B is no better than treatment arm A (for patient X).

Thus, it is possible that the evidence equally supports all *three* evidential hypotheses, even when our credence for each (given the evidence) falls considerably short of 50 percent.

To make this more concrete, consider a study that investigates the relative effectiveness of an SSRI (selective serotonin reuptake inhibitor) and some natural substance (hypericum) for treatment of depression.¹² Here, there are at least three possibilities to consider. One is that the SSRI (call this treatment A) is better than hypericum (call this treatment B). Another is that treatment B is better than treatment A. But a third possibility, represented by H_C , is that neither is more effective than the other. The evidence at hand might equally support all three hypotheses, yet if that is so, one cannot rationally hold a credence of 0.5 for all three hypotheses.

The second problem with the Freedman/Chalmers picture of theoretical equipoise has to do with the assumption that theoretical equipoise requires that the evidence equally support both treatment arms. However, once we acknowledge that a third hypothesis (like H_C) is required to exhaust logical possibility, the rules of the game have changed dramatically. For example, one might hold that most likely there is no difference between the two treatment arms, but if there is a difference it is more likely that treatment arm A is better than treatment arm B. Such an attitude is entirely coherent, and seems to be indicative of equipoise even if the one who holds the attitude thinks that *if* there is a difference, then the evidence favors A over B. In other words one's credence, given the evidence, in H_C might be .45 while one's credence in H_A based on the same evidence is .35 while also having a credence of .2 in H_B . We think that such a set of credences is coherent and even highly likely in a clinical research setting. We further think that the picture that we are developing of equipoise as it relates to clinical trials can accommodate such a set of credences, which is something that the picture that Freedman and Chalmers supply cannot do. As such, this presents what we believe to be a fatal problem for the Freedman/Chalmers picture.

With our picture of clinical trials in hand, we will now develop what we think is tenable conception of theoretical equipoise. As a starting point

we would like to introduce the notion of *epistemic equipoise*. Roughly speaking, epistemic equipoise is a state of *intolerable* uncertainty about whether one of the two competing treatments under consideration is better for a particular patient. With any scientific hypothesis there is some uncertainty, but exactly how much is tolerable is a question that has received considerable attention from both scientists and philosophers of science. Karl Popper (1959) argues that there is no logical answer to the question at least when dealing with probabilistic theories like those used in medicine (191). As a result, he suggests that we must adopt a methodological rule for what constitutes a tolerable level of uncertainty. Some have suggested that uncertainty can only be tolerable when the level of said uncertainty is less than or equal to 0.05 (Fisher 1956). Others have argued that even when the level of uncertainty is at 0.05 it is still intolerable. We do not have the space to consider this question in detail, but we do think the lower the uncertainty the better, especially when a patient's well-being is at stake. However, in what follows we will show that even if one adopts the most permissive tolerance of uncertainty, one can appeal to theoretical or epistemic equipoise in the ethical justification of a clinical trial.¹³

Our claim is that if there is *intolerable* uncertainty regarding which of the three competing hypotheses (H_A , H_B , or H_C) under consideration is best supported by the evidence,¹⁴ there exists a state of epistemic equipoise. It is epistemic equipoise that we think is capable of playing a part in the ethical justification of clinical trials. If there is intolerable uncertainty about which treatment (if any) is best supported by the evidence, then it is hard to say how one can hold that one treatment arm is better than the other. Thus, we think that a clinician in epistemic equipoise can justifiably recommend his or her patient for a clinical trial.

Theoretical equipoise is a special case of epistemic equipoise that results from insufficient evidence to distinguish between the merits of each hypothesis. The basic idea is that none of three hypotheses under consideration is decisively dominating the other two hypotheses in terms of evidential support. That's the intuitive notion, even if it is admittedly vague. Spelling out the details of what it means for one of the hypotheses to decisively dominate its competitors is a particularly tricky issue, as we mentioned above. However, we will say that a necessary condition of decisive domination is that the hypothesis in question must *minimally dominate* the other competing hypotheses. For any three mutually exclusive and exhaustive hypotheses, H_i , H_j , and H_k , and body of evidence E , H_i minimally dominates H_j

and H_k just when $\Pr(H_i|E) > \Pr(H_j|E) + \Pr(H_k|E)$. It follows that such a hypothesis will minimally dominate when its probability given the evidence is greater than 50 percent. Likewise, we can say that a state of theoretical equipoise holds between three mutually exclusive and exhaustive hypotheses if none of the hypotheses under consideration minimally dominates the others. This we think is the most permissive level of uncertainty that one can have regarding the treatment options.

For illustrative purposes, let's consider a toy example where a clinician-researcher, let's call her Dr. Sally, wishes to investigate whether St. John's wort (hypericum) is an effective alternative to standard SSRIs for treating patients with depression. She might have some anecdotal evidence from patients who claim to feel better when taking hypericum as an herbal supplement. However, she also has access to all of the clinical evidence in favor of using SSRIs. She now has two treatment arms under consideration: SSRIs will constitute treatment arm A and the herbal supplement hypericum will constitute treatment arm B. With respect to each patient suffering from depression she now has three hypotheses to test:

- H_A : SSRIs (treatment arm A) are better for treating patients with depression than hypericum (treatment arm B);
- H_B : Hypericum (treatment arm B) is better for treating patients with depression than SSRIs (treatment arm A); and
- H_C : Neither treatment arm is better than the other.

Let's suppose that Dr. Sally decides to test these hypotheses by running a clinical trial. Perhaps she has been relatively unimpressed with the past performance of SSRIs to relieve the suffering of her patients, but she nonetheless takes the previous clinical evidence in favor of the effectiveness of SSRIs seriously. As a result her credence, given the initial evidence, is 0.35 in H_A . Let's further suppose that she takes the reports from her patients who have taken hypericum as an herbal supplement seriously, but since there is no clinical evidence in favor of hypericum, she thinks that is less likely that hypericum is better than SSRIs. As a result her credence in H_B is a little less than her credence in H_A . Let's say her credence in H_B is 0.3, given the anecdotal evidence. That would leave her credence in H_C at 0.35, exhausting the logical possibilities.

Before the study even begins, Dr. Sally has some professional opinion, based on tenuous evidence with regard to each of the hypotheses under consideration. Despite the evidence, all of which Dr. Sally takes seriously, not a single hypothesis minimally dominates the others. As such, Dr. Sally is

theoretical equipoise. Furthermore, she is not in position to recommend one treatment over the other.

Let's further suppose that Dr. Sally gets some preliminary results from a small subset of the patients participating in the study and that those results suggest that those who are taking hypericum are doing modestly better than those taking SSRIs. This might constitute additional evidence in favor of H_B , but since the sample is small and it is early in the test, such evidence may only be enough to shift her credences a little. Perhaps her new credences given the total evidence E (new and old) is such that $\Pr(H_A|E) = 0.3$, $\Pr(H_B|E) = 0.4$, and the $\Pr(H_C|E) = 0.3$. The modest improvement in treatment group B has increased her credence in H_B , while decreasing her credences in the other two hypotheses. Yet since the evidence is modest, no hypothesis minimally dominates the others, allowing her to continue the study, as she is still in theoretical equipoise. If, however, the gains were more extraordinary, such gains might constitute enough evidence that would require her to revise her credences such that one hypothesis would minimally dominate the rest. In such a case, Dr. Sally would no longer be in theoretical equipoise, and as such she could not appeal to the uncertainty that results from lack of evidence to help justify the trials.

The moral of this story is that theoretical equipoise can handle evidence from a variety of sources, yet even when one takes that evidence seriously, it can still be used as part of the justification for clinical trials. However, there are limits to theoretical equipoise, and it is not the whole story. For example, it is possible by the end of Dr. Sally's trial that she will no longer be in theoretical equipoise. However, an important part of any scientific endeavor is replication. One might wonder how one can ethically justify replicating Dr. Sally's study once one of the hypotheses minimally dominates the others, assuming that those who choose to replicate the study take Dr. Sally's work seriously (as they ought). We will return to this issue in section 3 when we discuss epistemic equipoise that is not the result of a lack of evidence. At the risk of divulging spoilers, we will suggest that while there might be sufficient evidence such that one hypothesis could minimally dominate the others, it is possible that the evidence might be misinterpreted. We think that many cognitive biases can lead to such misinterpretations, and as such we should take those steps (such as perhaps replication and more expansive clinical trials) to ensure that we minimize any chance of doing so.

3. Uncertainty and Cognitive Bias

We think that theoretical equipoise can justify randomized clinical trials in a way that clinical equipoise cannot, for reasons given above. However, what about the case of the experienced clinician-researcher who is not in theoretical equipoise *ex hypothesi*? If the experienced clinician-researcher cannot appeal to clinical equipoise, does he or she have any reason to support trials in such cases? We think that it is possible for a researcher to remain in epistemic equipoise even if that researcher is no longer in theoretical equipoise. Theoretical equipoise depends on the available evidence. However, the broader notion of epistemic equipoise, as we characterize it, results from uncertainty due to either a lack of evidence or the possibility of misinterpreting the available evidence due to cognitive bias. This is particularly relevant in scenarios where the possibility of systemic error due to cognitive bias is great.

Although one might characterize the experienced clinician-researcher as by definition holding beliefs that are true (i.e., what she believes about the relative safety and effectiveness of an intervention really is the case), a less idealized version would hold that she could be mistaken in her view, despite having sufficient evidence, but only if interpreted correctly. This is because any one clinician-researcher's experience is limited and the lessons one might draw from that experience might be limited in important ways as well. A now well-known literature stemming from work by Kahneman and Tversky should give us pause before too confidently dismissing the possibility of error (Kahneman and Tversky, 1974; Kahneman, 2011).¹⁵ This is, we believe, the condition that even the best among clinician-researchers is generally in.

Consider as one example the availability heuristic. When faced with complicated choices or assessments people pretty consistently draw on the most readily available information they have. Events widely reported in the media, for example, influence individuals' beliefs about the likelihood of various causes of death. More "newsworthy" deaths are more frequently reported, making these pieces of information more readily available when someone is asked to estimate the likelihood of a particular cause of death. But the most cognitively available causes of death are not proportionately representative of the statistical likelihood of different causes. So this heuristic causes us to make errors despite the evidence that may in fact be available. This heuristic distracts us, it seems, from an assessment of

statistical verities. Consider a clinical example. When physicians are regularly bombarded by pharmaceutical drug company representatives with information about a condition like social anxiety disorder (SAD), they are more likely to make a diagnosis of SAD. Admittedly, whether someone has a condition like SAD is a clinical judgment, but even clinicians are vulnerable to the availability heuristic, leading to overdiagnosis.

To take a clinical research example, consider the case of Dr. Sally, presented above. At the start of the clinical trial investigating hypericum she is in a state of theoretical equipoise. However, as evidence begins to roll in, theoretical equipoise may be dislodged. Suppose this were so, could the trial continue?¹⁶ We argue above that epistemic equipoise can help provide the justification to continue. It constitutes a form of uncertainty that the trial aims to resolve, whether that uncertainty is due to insufficient evidence or worries about whether the researchers have accurately interpreted the evidence. One form of cognitive bias that might be relevant here is wishful thinking. For example the data might be interpreted in one of two ways. One way the data might be interpreted is as modest evidence of some harmful side effect. The other way it might be interpreted is as a statistical anomaly. Suppose that due to wishful thinking, Dr. Sally chooses to interpret the data as a fluke, but in reality it is modest evidence of harm. As more subjects are enrolled in the trial, however, and more data come in, it becomes increasingly difficult for Dr. Sally to indulge in wishful thinking. The additional data overcomes her cognitive bias. Thus, the real possibility of cognitive bias (in this case wishful thinking) is a legitimate source of uncertainty that supports continuing the trial to seek out a more statistically robust understanding of the effects of hypericum.

One core theme of cognitive bias research is that humans are bad folk statisticians; this is the issue that got Kahneman and Tversky going in their collaborative work. We make decisions that reflect a poor appreciation for statistical concepts, like base-rate probabilities, or the possibility of reversion to the mean, etc. The aim of a randomized controlled trial is precisely to produce a statistically robust conclusion about the relative safety and efficacy of an intervention that might ameliorate the effects of such cognitive biases.

Thus, even the experienced clinician-researcher might be in a state of epistemic equipoise, if this is characterized as a healthy respect for the possibility that cognitive biases might lead to a misinterpretation of the evidence. Awareness of such cognitive bias makes it reasonable for clinician-

researchers to remain uncertain about the accuracy of their interpretations of the evidence. This kind of uncertainty provides part of the ethical justification for imposing risk on trial participants in order to ameliorate the possibility of cognitive bias. Until they have conducted a trial with sufficient statistical validity, clinician-researchers are obligated to consider the possibility that what they believe about an intervention is due to cognitive error or bias.

Freedman worried, as we noted above, that theoretical equipoise is too fragile to give ethical support to a study. We have shown that it is not. However, we have taken the problem of the experienced clinician-researcher seriously, and we hope we have shown that we should be wary of cognitive bias and therefore retain some measure of uncertainty in the form of epistemic equipoise until there is good statistical evidence to correct for our biases. The experienced clinician-researcher poses a threat to theoretical equipoise because he or she is in a position to know more than other members of the medical community. Theoretical equipoise, as a species of epistemic equipoise, can (partially) justify many trials. However, for the experienced clinician-researcher who is no longer in theoretical equipoise, epistemic equipoise may persist and still provide the kind of uncertainty that makes a trial justifiable.

4. Final Remarks

Before we close, we would like to consider a couple of potential objections to our view. We think that we have shown that theoretical equipoise, insofar as it is a species of epistemic equipoise, can justify randomized clinical trials in many instances. Furthermore, we recognize that it is possible that an experienced clinician-researcher might leave the state of theoretical equipoise long before the completion of those trials. Nonetheless, we have argued that an awareness of the existence of certain cognitive biases can result in some uncertainty as to which treatment (if either) is better for a patient. This residual uncertainty can be intolerable thus licensing further trials, as the experienced clinician-researcher is in a state of epistemic equipoise. One might object by stating that it is possible that the experienced clinician-researcher has taken great care to avoid cognitive biases. Lets call this researcher the virtuous researcher. Perhaps the virtuous researcher has read the relevant work on biases and taken great care to avoid using heuristics that might lead to error. Would such a researcher be in epistemic equipoise? Perhaps he or she is not. However, while we think

it is possible and even likely in some cases that a researcher could be out of theoretical equipoise, we are not so confident that one could avoid all of the pitfalls of cognitive biases. Nonetheless, it is, in principle, possible. If a virtuous researcher exists, then that researcher would not be in epistemic equipoise and could not appeal to equipoise in the ethical justification of clinical trials.

However, we speculate that since the research on cognitive biases is still quite new, there is some reason to think that the biases identified by Kahneman and Tversky are just the tip of the iceberg. Furthermore, it cannot be underestimated just how much we are in the thrall of such biases. As such, we are skeptical whether one could, by oneself, take the requisite care to eliminate all such biases without help from institutions like peer review, replication, and yes even clinical trials. It is our opinion that such institutions perform a valuable function in helping us overcome the heuristic biases that might plague without such help.

A similar objection might arise from the idea of an *overconfident* clinician-researcher. This researcher has taken no care to eliminate biases, because he or she is blissfully unaware of the potential for bias. As a result, the overconfident researcher is not as uncertain in a personal sense about the results and perhaps not in a state of epistemic equipoise. Perhaps the overconfident researcher even has an amazing track record in cases like these, which further adds to his or her sense of certainty. We think that such a researcher has failed to meet his or her obligation to watch for and eliminate such cognitive biases. We think that such overconfidence is an epistemic vice. Given that the stakes of medical research are high and the consequences are dire should we get it wrong, we think that recognizing the potential for bias and taking steps to eliminate such biases is a moral obligation. This means that adopting some stance of epistemic humility is morally required for medical research. That stance we believe is realized by epistemic equipoise.

Summing up, we believe that epistemic equipoise is a state of intolerable uncertainty that may result from either a lack of evidence or the possibility of error due to cognitive bias. Furthermore, we think that institutions like clinical trials are one tool that can (1) gather evidence and (2) help reduce error due to cognitive bias. As such, we think that epistemic equipoise plays an important part in the ethical justification of clinical trials, in ways that clinical equipoise cannot. One might, of course, question whether current protocols regarding randomized clinical

trials is the best way of fulfilling the two functions just mentioned. It may be the case that after phase III trials, as they are currently conducted, intolerable uncertainty will remain. As such, we might ask whether we need to change our protocols to better gather evidence and eliminate bias. We think this is an important question, and one that should be the subject of further statistical and philosophical research.

We further think that such research might inform other debates in the literature such as when we are required to continue or discontinue trials. For example, given the importance of getting an intervention into active clinical use as early as possible, we might ask whether going as far a Phase III trial is necessary. We think timely access to new medical technologies is a very important good, and if safety can be assured, then the timelier the better. However, it does appear that whatever protocol we do adopt, some tradeoff between eliminating uncertainty and the timely propagation of new medical techniques must be managed. That means that when we decide how much uncertainty is tolerable, we must keep this tradeoff in mind. Achieving the optimum between the desiderata of eliminating uncertainty and making new medical technologies available in a timely manner should be the focus of future inquiry.

Works Cited

- Chalmers, Thomas. 1978. The Ethics of Randomization as a Decision-Making Technique and the Problem of Informed Consent. In T. L. Beauchamp and L. Walters (eds) *Contemporary Issues in Bioethics*. Dickenson: Encino, CA, pp. 426-29.
- . 1981. The Clinical Trial. *The Milbank Memorial Fund Quarterly - Health and Society* 59(3):324-339.
- Elga, Adam. 2007. Reflection and Disagreement. *Noûs* 41(3): 478-502.
- Emanuel, Ezekiel, David Wendler, and Christine Grady. 2000. What Makes Clinical Research Ethical? *Journal of the American Medical Association* 283(20): 2701-2711.
- Feldman, Richard. 2007. Reasonable Religious Disagreement. In L. Antony (ed) *Philosophers Without Gods*. Oxford University Press: Oxford, 197-214.
- Fisher, R. A. 1956. *Statistical Methods and Scientific Inference*. Oliver and Boyd: London.
- Freedman, Benjamin. 1987. Equipoise and the Ethics of Clinical Research. *New England Journal of Medicine* 317(3): 141-145.
- . 1992. A Response to a Purported Ethical Difficulty with Randomized Clinical Trials Involving Cancer Patients. *The Journal of Clinical Ethics* 3(3): 231-34.
- Hellman, Samuel and Deborah S. Hellman. 1991. Of Mice but not Men: Problems of the Randomized Clinical Trial. *The New England Journal of Medicine* 324: 1585-1589.
- Hypericum Depression Trial Study Group. 2002. Effects of *Hypericum Perforatum* (St. John's Wort) in Major Depressive Disorder: A Randomized Controlled Trial. *Journal of the American Medical Association* 287: 1807-1814.
- Kahneman, Daniel. 2011. *Thinking, Fast and Slow*. Farrar, Straus, and Giroux: New York.

- Kahneman, Daniel and Amos Tversky. 1974. Judgment Under Uncertainty: Heuristics and Biases. *Science* 185: 1124-1131.
- Kelly, Thomas. 2005. The Epistemic Significance of Disagreement. In J. Hawthorne and T. Gendler (eds) *Oxford Studies in Epistemology*, volume 1. Oxford University Press: Oxford, 167-96.
- Miller, Franklin G. and Howard Brody. 2003. A Critique of Clinical Equipoise: Therapeutic Misconception in the Ethics of Clinical Trials. *Hastings Center Report* 33(3): 19-28.
- Moseley, J. Bruce, Kimberly O'Malley, Nancy J. Petersen, Terri J. Menke, Baruch A. Brody, David H. Kuykendall, John C. Hollingsworth, Carol M. Ashton, and Nelda P. Wray. 2002. A Controlled Trial of Arthroscopic Surgery for Osteoarthritis of the Knee. *New England Journal of Medicine* 347: 81-88.
- Popper, Karl. 1959. *Logic of Scientific Discovery*. Hutchinson: London.
- Titelbaum, Michael. Forthcoming. Rationality's Fixed Point (Or: In the Defense of Right Reason). In *Oxford Studies In Epistemology*.
- Ubel, Peter A. and Robert Silbergleit. 2011. Behavioral Equipoise: A Way to Resolve Ethical Stalemates in Clinical Research. *American Journal of Bioethics*. 11(2): 1-8.

¹ Miller and Brody agree with Emanuel, Wendler, and Grady (2000) that scientific value is only part of the ethical justification of research. Emanuel, Wendler, and Grady argue for seven ethical requirements. While it is not the focus of our paper to offer a complete account of the ethical justification of biomedical research, we also agree that the feature of scientific uncertainty or equipoise does not, by itself, suffice to justify this research.

² Although see Ubel and Silbergleit (2011). They describe a case where a group of clinicians continue to dissent even after the conclusion of particular drug trial.

³ Epistemic peers are usually understood to be possessing of the same level of intelligence, as well as being equally responsive to evidence, etc.

⁴ In the formal epistemology literature, 'credence' is a technical term indicating the degree of belief or confidence that an agent has that a particular claim is true. These degrees of belief are generally represented as probabilities that conform to Kolmogorav's axioms. For example, the credence that someone might have in the proposition p is represented as a number between zero and one and is generally symbolized by the expression, $\Pr(p)$. It is further held that these probabilities are subject to revision in the face of a body of evidence E via updating by conditionalization. The updated credence, which is often referred to as the posterior probability, is expressed by the expression, $\Pr(p|E)$.

⁵ There are, of course, some legal and societal reasons why the rest of the medical community might not do so, however.

⁶ One might object that a trial is necessary in order to show whether the experienced clinician-researcher is the one who has correctly reasoned in this context. We are somewhat sympathetic to such a view as will be seen in Section 3. However, our point is that it is not mere disagreement that plays a part in the ethical justification of clinical trials, but rather the uncertainty that we might have that the experienced clinician-researcher has reasoned correctly.

⁷ Feldman (2007) deals with the rough-grained doxastic attitudes: belief, disbelief, and suspension of judgment. His view is that if two epistemic peers who consider exactly the same body of evidence come to a disagreement about the truth of a proposition p such that one party believes p and the other party disbelieves p , then both parties are rationally required to suspend judgment with regard to the truth of p . Elga (2007), on the other hand, endorses an equal weight view for more fine-grained doxastic attitudes (aka credences). On his view, if epistemic and evidential peers assign different credences to a claim p then they should each adjust their credences to meet in the middle.

⁸ See also Chalmers (1981).

⁹ Freedman writes, "In Chalmer's view, [theoretical] equipoise is disturbed when the odds that A will be more successful than B are anything other than 50 percent" (1987, 143).

¹⁰ This question was the subject of a study conducted by Moseley et al. (2002).

¹¹ One might argue that less invasive procedures are always better than more invasive procedures, and thus that if it is shown that effectiveness of the two treatment arms are equal, then the less invasive procedure can be said to be better. Perhaps that is right. However, phenomena like the placebo effect can often confound such results. In fact, Moseley et al. (2002) were quite aware of this possibility, which is why they tested one group of patients with arthroscopic surgery (treatment A) and the other group with a sham surgery (treatment B). Their results were that there was no significant difference between the two groups, constituting confirmation of the third hypothesis.

¹² This example also parallels an actual study (Hypericum Depression Trial Study Group 2002). We owe this (adapted) example to Miller and Brody (2003).

¹³ By adopting the more permissive stance regarding tolerable uncertainty, we hope to head off any worries that we might be requiring a level of certainty so high as to make equipoise trivial. We are grateful to Chris Herrera for pointing out this potential objection.

¹⁴ Our understanding of evidence is quite permissive, and we do not intend to limit it to merely data from clinical trials. It might be that evidence can come from a variety of sources - e.g., professional hunches, allied evidence from biochemical analogs, etc.

¹⁵ In a series of publications, summarized in Kahneman's recent book *Thinking, Fast and Slow* (2011), Kahneman and Tversky showed how a variety of heuristics bias decision makers in a variety of ways.

¹⁶ We recognize that we are now verging on questions regarding the ethics of stopping rules in clinical trials. We do think these issues are related because once theoretical equipoise has been displaced the question of whether to continue the trial is similar in important ways to the question of whether to conduct a trial. However, we cannot give a full consideration to the ethics of stopping rules here.