

Douglas MacKay

Policy Equipoise and Interventional Superiority

Forthcoming at *Journal of Development Effectiveness*

### Abstract

According to the norm of policy equipoise, it is permissible to randomly assign participants to two or more interventions in a public policy randomized controlled trial (RCT) when there is meaningful uncertainty among the relevant expert community regarding which intervention is superior. While this norm is gaining traction in the research ethics literature, the idea of interventional superiority remains unclear. Is one intervention superior to another if it is reasonably expected to realize one outcome of interest more effectively, even though there is uncertainty regarding other outcomes of interest? Or must an intervention be reasonably expected to realize *all* outcomes of interest more effectively? I address this question in this paper. My aim is to develop and defend an account of interventional superiority for policy RCTs that are authorized, funded, or conducted by government institutions. I defend the *greatest value view*, according to which one intervention is superior to another if and only if it is reasonably expected to more effectively realize a set of outcomes with greater value.

### Text

Suppose you are a policymaker in a low-income country considering whether to supplement the existing social safety net with a small, monthly cash transfer targeted at the poorest households. Suppose further that there is high quality evidence from randomized controlled trials (RCTs) in similar jurisdictions that the transfer will improve people's consumption and food security. There is

uncertainty, however, regarding additional outcomes of great concern, including the prevalence of low birth weights, child nutrition, and the use of healthcare services. Would it be permissible to conduct a large-scale cluster RCT, randomizing poor households in some districts to the transfer program and poor households in other districts to the status quo? Or should you implement the cash transfer on the grounds that it will improve people's consumption and food security, even though its broader effects are uncertain?

To decide this question, you might appeal to the norm of 'policy equipoise,' according to which it is permissible to randomly assign participants to two or more interventions when there is meaningful uncertainty among the relevant expert community regarding which intervention is superior (Asiedu et al. 2021; MacKay 2020; MacKay 2018). But is there meaningful uncertainty in this case? On the one hand, there is high quality evidence that the cash transfer will more effectively improve households' consumption and food security compared to the status quo. But does the uncertainty regarding the additional outcomes imply that the cash transfer is not superior? More generally, for one intervention to be superior to another, must it be reasonably expected to realize *all* outcomes of interest more effectively? Or is this standard too demanding?

I address this question in this paper. My aim is to develop and defend an account of interventional superiority for policy RCTs that are authorized, funded, or conducted by government institutions. I defend the *greatest value view*, according to which one intervention is superior to another if and only if it is reasonably expected to realize more effectively a set of outcomes with greater value.

I begin in part 1 with a brief overview of the problem of interventional superiority. In part 2, I consider existing approaches to this problem, arguing that policy research ethics scholars largely ignore it and that the accounts developed by clinical research ethics scholars are not useful for policy research. In part 3, I develop the *advantage view*, which addresses a problem with the dominance view,

but faces problems of its own. In part 4, I develop and defend the *greatest value view* which resolves the problems faced by both the dominance and advantage views. In part 5, I offer guidance regarding its application to RCTs.

My paper contributes to the growing literature on the ethics of public policy experiments. While some scholars have recognized the problem of interventional superiority (Asiedu et al 2021; MacKay 2020, 328; Petticrew et al. 2013), no one has yet developed a systematic account that is tailored for public policy research. My paper aims to fill this gap, showing how the norm of policy equipoise can be applied to an important set of cases.

The focus of my paper is limited in two ways. First, I provide an account of interventional superiority for policy RCTs authorized, funded, or conducted by government institutions and agencies. Government and non-government institutions have different obligations, and developing an account for all research sponsors is too large a task for this paper.<sup>1</sup> Second, my paper does not provide a comprehensive account of the permissibility of random assignment in policy experiments, but rather focuses on a particular justification for random assignment, namely, policy equipoise. I recognize that the scarcity of an intervention may also sometimes justify random assignment (see MacKay and Cohn 2023).

## 1 The Problem of interventional Superiority

Consider a simple case. There are two policy interventions, A and B, and there is one outcome of interest, X. The expert community occupies a state of equipoise regarding these interventions if and only if they are collectively uncertain regarding whether A or B will most

---

<sup>1</sup> For discussion of the obligations of different types of research funders, see Pierson and Millum 2018 and MacKay 2015.

effectively realize X (London 2024, 216; MacKay 2020, 328). This uncertainty may be due to a lack of evidence regarding the effectiveness of A and B, or it may be that the evidence is not conclusive, providing some experts good reasons to favor A, and others good reasons to favor B.

Importantly, although the term ‘equipoise’ suggests that the expert community must be ‘equally poised’ between A and B, it is widely recognized that for equipoise to obtain it must only be the case that there is reasonable uncertainty among the expert community regarding A and B (London 2024, 227). The expert community is thus in equipoise regarding the comparative effectiveness of A and B even if 70% of experts think A is more effective, or if all experts are 60% confident that A is superior to B. Were equipoise to require the expert community to be equally poised between A and B, very few RCTs could be permissibly started, and only a small number of these could be permissibly completed, since accumulating evidence would change experts’ minds (London 2024, 219). Governments would be prohibited from conducting research necessary to resolve reasonable disagreements and determine which intervention will most effectively promote people’s interests.

For equipoise to be disturbed, the evidence regarding the comparative effectiveness of A and B in realizing X must meet an “epistemic threshold” such that the reasonable uncertainty among experts is eliminated or sufficiently reduced (London, 2024, 218, 230). Where this epistemic threshold has not been met, I shall say that the expert community is *uncertain* regarding which intervention will more effectively realize a particular outcome. Where this threshold has been met, I shall say that one intervention is *reasonably expected* to realize a particular outcome more effectively than another.

Consider a more complex case. There are two policy interventions, A and B, and there are three outcomes of interest, X, Y, and Z. As illustrated in Table 1, there is uncertainty regarding

whether A or B is more effective at realizing Y and Z, but A is reasonably expected to be more effective than B with respect to X.

Table 1

<b>Outcome Intervention</b>	<b>X</b>	<b>Y</b>	<b>Z</b>
<b>A</b>	✓	?	?
<b>B</b>	X	?	?

In this case, the expert community is in equipoise regarding the comparative effectiveness of A and B in realizing Y and Z. Regarding Y and Z, we might say, the expert community occupies a state of *outcome equipoise*. Outcome equipoise obtains when the relevant expert community is uncertain regarding the comparative effectiveness of two or more interventions in realizing a specific outcome. But what we really care about is whether it is permissible to randomly assign people to A and B. We thus care about *interventional equipoise*: is there uncertainty among the relevant expert community regarding the superiority of A or B, with the implication that it is permissible to randomly assign participants to these interventions? Or is A superior to B because it is more effective in realizing X, with the implication that random assignment would be wrong? More generally, how should we determine which intervention is superior in these types of cases?<sup>2</sup>

An account of interventional superiority is necessary to determine when intervention-level equipoise is disturbed. This account would specify conditions that must be satisfied for one intervention to be superior to another. When a study only has one outcome, there is no difference between outcome equipoise and interventional equipoise. But when a study has two or more

---

<sup>2</sup> In a paper on clinical equipoise, Fred Gifford (2000, 423) draws a similar distinction between “one-dimension-considered” equipoise and “all-dimensions-considered” equipoise.

outcomes, these concepts may diverge. Regarding the example illustrated in Table 1, the relevant expert community is in outcome equipoise regarding A and B with respect to Y and Z, but may not be in intervention equipoise if the superiority of A with respect to X is sufficiently important to imply that A is superior to B. As I show below, judgments regarding interventional equipoise require researchers to make value judgments, not merely judgments of effectiveness (see also Gifford 2000, 417).

A full answer to the problem of intervention superiority requires, first, an account of the type and level of evidence that is necessary and sufficient to reach the epistemic threshold and so disturb outcome equipoise, that is, for the expert community to reasonably conclude that one intervention will more effectively realize a specific outcome than another. It also requires, second, a set of necessary and sufficient conditions for determining when an intervention is superior in cases where one intervention more effectively realizes a specific outcome than the alternatives, but not *all* outcomes of interest. That is, it requires an account of when interventional equipoise is disturbed in cases where it cannot be reduced to outcome equipoise. Call these ‘complex cases.’ This paper focuses on the latter task and so only offers a partial solution to the problem of interventional superiority. I shall assume throughout the paper that we have an account of the epistemic threshold interventions must meet for uncertainty regarding their effectiveness in realizing a particular outcome to be resolved or sufficiently minimized.

## 2 Existing Approaches

Prior to constructing a novel account of interventional superiority, it is prudent to review the research ethics literature for promising existing approaches. Unfortunately, policy research ethics scholars largely fail to acknowledge this problem, and clinical research ethics scholars have not developed accounts that are both applicable to policy research and normatively defensible.

Consider the policy research ethics literature first. Some authors don't consider complex cases and so fail to offer an account of interventional superiority that has the resources to address them (Abramowicz and Szafarz 2020; MacKay 2018, 62; Ziliak and Teather-Posadas 2016, 425; Baele 2013, 17-21). For example, in my 2018 paper, "The Ethics of Public Policy RCTs: The Principle of Policy Equipoise," I propose that one intervention is superior to another when it is known to realize the outcome of interest more effectively. Accordingly, the case I discuss to illustrate this conception of policy equipoise features only one outcome (MacKay 2018, 65).

Other authors recognize the problem posed by complex cases, but don't develop accounts of interventional superiority to address them (Asiedu et al 2021; MacKay 2020, 328; Petticrew et al. 2013). For example, in a passage that nicely illustrates the challenge complex cases pose, Edward Asiedu et al (2020) acknowledge the problem of interventional superiority but suggest that researchers should largely just ignore it by limiting the scope of their concern to a specific set of outcomes:

We highlight one further component of the above definition of policy equipoise: "achieving improvements in the relevant outcomes of the study." An RCT could, for example, be focused on testing alternatives for reducing child malnutrition. A treatment arm could be cash transfers. If we were to "know" that cash transfers create benefits beyond reducing child malnutrition (a reasonable conjecture, albeit one with uncertainty regarding the types and magnitudes of impacts), should the nonchild malnutrition outcomes be considered when examining if any treatment arm is viably "better" than all of the others? Doing so would require a full-blown multidimensional welfare analysis, weighing all ethically relevant outcomes. We argue that it is typically sufficient to constrain the comparison of treatment arms (and counterfactual policy) to the problem at hand (e.g., paths to improving the primary outcomes of the study). Realistically, anything that dominated in all dimensions with

certainty would likely not be “practically attainable.” We do see merit in this discussion, though.

Similarly, I recognize the problem of complex cases in recent work but fail to offer an account of interventional superiority which may help researchers work through them (MacKay 2020, 328).

Turning to the clinical research ethics literature, scholars writing on clinical equipoise are aware of complex cases, for clinical trials often have multiple outcomes of interest, with some related to mortality and morbidity and others related to quality of life - e.g. side effects (London 2022, 264; Gifford 2000, 413-417; Lilford and Jackson 1995, 555-556; Friedman 1987, 143-144; Fried 1974, 153-154). Some also propose accounts of interventional superiority which can address these types of cases.

First, Fred Gifford (2000, 418) argues that the superiority of a clinical intervention can be determined by reference to the values of individual patients (see also Chard and Lilford 1998, 893-895; Lilford and Jackson 1995, 555-556). For example, suppose it is known that intervention A is better than intervention B in terms of survival, but that the precise extent is uncertain. Suppose further that while B is known to be better with respect to side effect X, there is uncertainty which is better with respect to side effects Y and Z. Investigators may wish to compare A and B by means of an RCT, and, on Gifford’s (2000, 418) account, it is permissible to randomly assign patients to A or B provided these patients are indifferent among them given their values. One intervention is superior to another, on Gifford’s account, if and only if a patient, given existing information regarding likely outcomes, prefers it to the alternative. A proposed clinical RCT satisfies the norm of interventional equipoise for patients who are indifferent between the interventions being compared, given what is known and not known about the relevant outcomes.

Gifford’s solution is in principle workable for clinical RCTs since: (1) patients consent to enrollment in clinical trials and so could be asked whether they are indifferent between the arms of

the study; and (2) it's at least possible that patients' value judgments regarding trade-offs between survival and quality of life, among other factors, don't correlate with the outcomes of interest, thus ensuring that trial participants are representative of the broader patient population. Unfortunately, it's unlikely to be workable for policy RCTs. First, many policy RCTs proceed without securing participants' informed consent, meaning there would be no possibility of only enrolling people indifferent between the arms of the study (MacKay and Chakrabarti 2019). Second, while there is unlikely to be a correlation between patient's values and the biological outcomes of clinical interventions, it's highly likely that participants whose values lead them to be indifferent between arms of a policy RCT are different in behaviorally relevant ways from people who are not, thus raising serious concerns regarding the generalizability of policy RCTs that satisfy a conception of interventional equipoise consistent with Gifford's proposal.

A second solution is developed by Alex John London in *For the Common Good: Philosophical Foundations of Research Ethics*. Importantly, while London's (2022, 159) view of the philosophical foundations of research ethics covers clinical *and* social science research, his account of interventional superiority applies most directly to clinical research. London need not therefore be committed to the claim that his account is applicable to policy research, but it offers the basis for a plausible solution to the problem of interventional superiority in this area and so is worth discussing.

London's (2022, 264) account is presented in the context of a practical test for determining when it is permissible to assign a research participant to a randomized controlled trial (RCT):

Uncertainty Regarding Basic Interests: For each individual with a particular condition (e.g., a health problem), the care and protection afforded to that individual's basic interests satisfies the condition of equal concern just in case every intervention to which that person might be allocated in a research study is admissible. An intervention is admissible for an individual just in case there is either uncertainty among, or conflict between, experts about whether it is

dominated by any other intervention as a means of safeguarding the basic interests of that individual.

For London therefore, A is superior to B if and only if there is no uncertainty that A dominates B in the safeguarding of an individual's basic interests. Basic interests are those interests all people share in being able "to develop and exercise their basic intellectual, affective, social, and physical capacities in order to be able to formulate, pursue, and revise a meaningful life plan, including forming and maintaining relationships of significance with others (London 2022, 133)." It is therefore permissible to enroll someone in an RCT if and only if there is no intervention that is reasonably expected to dominate the others, that is, be superior with respect to all basic interest-related outcomes. In cases where this practical test is satisfied, participants in a clinical trial can expect to receive a treatment that is favored for their condition by at least a reasonable minority of clinicians (London 2022, 265).

London's view offers a promising solution to the problem of interventional superiority as it arises in clinical research. But, like Gifford's view, it requires the consideration of each individual's situation to determine if one intervention is superior to another. As such, it is not workable for policy research which typically involves the application of an intervention to a population. It is worth considering, however, whether an account inspired by London's view, namely, one which understands superiority in terms of dominance, offers a promising answer to the problem of interventional superiority in policy research. Call this the *dominance view*:

Dominance View: A is superior to B if and only if A dominates B in the safeguarding of people's basic interests - i.e. is reasonably expected to realize all basic-interest related outcomes more effectively.

To return to the example in Table 1, even if A is reasonably expected to more effectively realize basic interest-related outcome X compared to B, provided there is uncertainty regarding whether A

will more effectively realize either basic interest-related outcome Y or Z than B, A is not superior to B, and so it is permissible to randomly assign people to A and B.

Unfortunately, there are two problems with the dominance view. First, an intervention must meet an incredibly high bar to count as superior to another, being more effective in the realization of *all* outcomes of interest. As a result, it permits policy RCTs that involve the random assignment of participants to interventions that are *prima facie* inferior. Consider the following case. In 2010, the Zambian Ministry of Community Development, Mother and Child Health (MCDMCH) began rolling out the Child Grant Program which offered poor households with a child under five an unconditional cash grant of roughly \$12 per month (Handa et al. 2016, 360-361). Households were randomized to either an experimental group, which started receiving payments in January 2011, or a delayed entry control group. Researchers studied the intervention's effect on consumption, food security, school enrollment, children's material wellbeing, children's nutrition, productive investment, asset accumulation, and child health (Handa et al. 2016, 358). The researchers acknowledged that cash transfers had been proven to improve consumption, food security, and school enrollment but that their effects on the other outcomes were uncertain (Handa et al. 2016, 361). They therefore did not appeal to uncertainty to justify random assignment, but argued instead that randomization was "ethically feasible" on the grounds that the MCDMCH did not have sufficient resources or capacity to provide the intervention to all eligible households (Handa et al. 2016, 361).

The dominance view implies that the cash grant is not superior to the control - i.e. no intervention - for there is uncertainty regarding whether it dominates the control. Although the cash grant is reasonably expected to more effectively promote consumption, food security, and school enrollment, there is uncertainty regarding whether it will more effectively realize other basic-interest related outcomes, including children's material wellbeing, nutrition, and health. According to the

dominance view therefore, random assignment of people to the intervention and control in this case is permissible, and would be permissible *even if the cash grant was not scarce*.

In my view, the researchers are correct in their judgment that the cash transfer is superior to the control for there is strong evidence that it will make a meaningful impact on people's basic interests related to consumption, food security, and school enrollment. Although there is uncertainty regarding whether the cash transfer will also improve children's material wellbeing, nutrition, and health, people have a strong claim to interventions known to combat hunger and facilitate school enrollment. Adequate food and access to school are basic needs. It's also unlikely that the cash transfer will reduce children's material wellbeing, nutrition, and health outcomes to a degree that would put the cash transfer's superiority in question. That is, it is unlikely that the cash transfer would reduce these outcomes to such a degree that it would no longer be a net beneficial intervention and so superior to no intervention.

The second problem with the dominance view is that random assignment can be made permissible simply by adding another outcome to the list of the outcomes under study.<sup>3</sup> Suppose A dominates B with respect to outcomes X and Y, meaning it would be impermissible for investigators to randomly assign participants to either A or B. If investigators simply add a third outcome - Z - for which the impacts of A and B are uncertain, it is now permissible to randomly assign participants to A or B. Perhaps there are cases where the addition of Z should mean that random assignment is permissible. But surely the mere addition of a further outcome of interest for which the causal impacts of the interventions are uncertain should not imply that random assignment is now permissible.

---

<sup>3</sup> Thanks to Joseph Millum for raising this point.

If the dominance view is deficient for these reasons, how might we improve upon it? What view explains the judgment that in the case of the Zambian RCT, the cash transfer is superior to no intervention?

### 3 The Advantage View

Let us take a step back. Researchers employ RCTs to determine the causal impact of interventions on outcomes of interest. For education researchers, these outcomes may be college enrollment or math scores; for researchers studying the social safety net, they may be income or housing security; and for public health researchers they may be the incidence or prevalence of certain diseases. But following London (2022, 3), policy research should be understood in broader terms as the activity whereby societies produce the knowledge that social institutions require to safeguard and promote the interests of community members. Governments have duties of justice to realize certain outcomes for their residents, including reasonable life expectancies, safe neighborhoods, and an adequate quality of life, and policy research is necessary to produce the knowledge governments require to effectively discharge these duties.

Elsewhere, I call these outcomes that are the object of governments' duties of justice "target outcomes (MacKay 2020, 323)." Target outcomes are "specifications of the type and amount of goods or services governments have a duty to provide, and the outcomes they have a duty to realize (MacKay 2020, 323)." They may include certain levels of food security, health insurance coverage, access to nutritious food, or specific educational outcomes.

Because governments have a duty to realize target outcomes, it is plausible to evaluate interventions in terms of their effectiveness in realizing them, all else being equal. Thus, if there is no uncertainty that A is better than B in realizing target outcome X, A is superior to B. Similarly, if there is no uncertainty that A is better than B in realizing target outcome X, and there is uncertainty

regarding whether A is better than B in realizing Y and Z, A is superior to B, for A promises to realize more target outcomes than B. Indeed, the researchers leading the *Zambian child grant RCT* may have had something like this in mind when deciding that scarcity, not uncertainty, offered a legitimate justification for randomization. Because the child grant intervention is better than the control with respect to consumption, food security, and school enrollment, and not known to be inferior with respect to the realization of any other target outcome, it is the superior intervention.

This analysis of interventional superiority supports the *advantage view*:

Advantage View: A is superior to B if and only if for target outcomes X, Y, Z, and so on, A is reasonably expected to realize a larger number of these target outcomes more effectively than B.

On this view, A is superior to B if and only if it has an advantage over B in the realization of target outcomes. Interventional equipoise is thus disturbed when one intervention is reasonably expected to realize a greater number of target outcomes than the alternatives. Outcome equipoise regarding at least one outcome is thus necessary for interventional equipoise, but not sufficient as is the case with the dominance view.

One strength of the advantage view is that it yields the correct answer in the *Zambian RCT* case: the cash grant is superior to the control because it is reasonably expected to realize three target outcomes more effectively - i.e. consumption, food security, and school enrollment - and there is uncertainty regarding whether it or the control will more effectively realize the other target outcomes. It can also handle more complicated cases where A is better on some target outcomes but where B is better on others. Consider the following:

Table 2

Target Outcome Intervention	X	Y	Z	XX	YY	Total

<b>A</b>	✓	X	✓	?	?	2
<b>B</b>	X	✓	X	?	?	1

A is superior to B, according to the advantage view, since it is reasonably expected to more effectively realize two target outcomes whereas B only more effectively realizes one.

There are two problems with this view, however. First, some target outcomes are more important than others from the standpoint of justice. There are many outcomes governments have a duty to realize, but some are more fundamental to people’s wellbeing than others, or are more closely tied to their ability to exercise their freedom. Ensuring people are safe from violent attack, are food secure, or have access to essential medical interventions which prevent or treat debilitating disease or disability are arguably more basic duties of justice than ensuring certain literacy outcomes for high school seniors or ensuring people from lower socioeconomic backgrounds have the same, real opportunities to attend institutions of higher education as people from higher socioeconomic backgrounds, as important as these latter outcomes are. To return to the Zambian child grant RCT, improving food security and consumption is arguably more important than improving families’ levels of productive investment. The advantage view does not recognize this for it treats target outcomes as equally valuable. A is superior to B, on this view, so long as A is reasonably expected to realize a larger number of target outcomes than B.

The second problem with the advantage view is that interventions might impact target outcomes to greater or lesser extents. There may be cases where A is reasonably expected to realize X more effectively than B, but the impacts on X are very small. If A’s impacts on X are very small, and B’s reasonably expected impacts on Y are very large, it may be reasonable to judge that B is superior to A. The advantage view ignores this consideration.

These problems illustrate the shortcomings of the advantage view. It does not consider two factors that are relevant for determining which intervention is superior: (1) the moral importance of the outcomes; and (2) the size of each intervention's impact on the outcomes. Depending on the target outcomes in question and the size of the interventions' impacts, it may be reasonable to judge one intervention to be superior than another, even though the latter is favored by the advantage view. For example, in the case outlined in Table 2, there may be circumstances where it is reasonable to think that B is superior to A, for example, because B's impacts on Y are large and Y is much more important than X or Z. Simply considering whether one intervention is better than another in the realization of target outcomes ignores relevant information.

#### 4 The Greatest Value View

Compared to the dominance view, the advantage view takes a step in the right direction. Is there a view that rejects the contention that outcome equipoise is sufficient for interventional equipoise, but also incorporates considerations of the moral importance of outcomes and the size of interventions' impacts?

Consider the *greatest value view*:

Greatest Value View: A is superior to B if and only if it is reasonably expected to more effectively realize a set of target outcomes with greater value. The value of a set of target outcomes is a function of (1) the value of the target outcomes that are reasonably expected to be realized by each intervention, and (2) the degree to which each intervention is reasonably expected to realize these outcomes.

On the greatest value view therefore, A is superior to B if and only if it enables governments to realize valuable target outcomes more effectively, that is, if and only if it enables governments to realize a greater amount of valuable target outcomes. To put it differently, A is superior to B if and only if it

enables governments to discharge their most pressing duties of justice more effectively; the society in question would be *more just* if the set of target outcomes promised by A were realized, than if the target outcomes promised by B were realized. Interventional equipoise is thus disturbed when A is reasonably expected to realize a more valuable set of target outcomes than B.

The greatest value view retains what the advantage view gets right, namely, that one intervention can be superior to another even if there is uncertainty regarding some of its impacts: outcome equipoise is not sufficient for interventional equipoise. However, by considering the value of different target outcomes and each intervention’s degree of impact, it avoids the problems faced by the advantage view. A remaining challenge with the greatest value view lies with its application to real-world cases, a problem I return to in part 5.

One might argue that the greatest value view leaves out a factor relevant to determining the superiority of an intervention: the relevant expert community’s *degrees of belief* regarding whether one intervention more effectively realizes target outcomes compared to another. For the greatest value view, degrees of belief regarding the comparative effectiveness of interventions in realizing target outcomes only matter if they are sufficiently strong to meet the epistemic threshold, thus disturbing outcome equipoise. But consider a case such as this:

Table 3

<b>Outcome Intervention</b>	<b>X</b>	<b>Y</b>	<b>Z</b>
<b>A</b>	✓	X	70%
<b>B</b>	X	✓	30%

Supposing that target outcomes X and Y are equally valuable and A realizes X to the same degree that B realizes Y, the expert community occupies a state of interventional equipoise regarding A and B. But shouldn't the expert community's high degree of confidence that A will more effectively realize Z than B count, perhaps tipping the scales?

This high degree of confidence is surely relevant for policymakers deciding to implement either A or B. All else equal, policymakers should implement A since it promises the greatest expected value. But this doesn't mean that it would be wrong to conduct a study randomly assigning people to A and B. Since the quality of evidence is not sufficiently high to disturb outcome equipoise regarding Z, there is a reasonable chance that B will more effectively realize Z than A, meaning it is the superior intervention. If policymakers permanently implement A, there is thus a reasonable chance they've implemented the inferior intervention. To maximize expected value, policymakers should conduct an RCT to determine which intervention more effectively realizes Z, and then permanently implement the policy which does so. This will ensure that the superior policy is implemented. This approach is also defensible to people randomly assigned to B since it is ex ante uncertain whether B is inferior to A. Conducting an RCT thus shows sufficient concern for the interests of people randomly assigned to A or B, and ensures that people will not be permanently subject to an inferior policy.<sup>4</sup>

---

<sup>4</sup> This analysis also applies in cases where researchers are confident in the *direction* of an intervention's impact on an outcome, but uncertain regarding whether one intervention is superior to another. To consider the case in table 3, suppose A is reasonably expected to have a positive impact on Z, but the extent of this impact is uncertain. This is surely relevant to the question of whether policymakers should implement A or B, but I think it is permissible to conduct an RCT comparing A and B since there is still uncertainty whether A or B will more effectively realize Z. It might turn out that A has a small positive impact on Z while B has a large positive impact. If policymakers permanently implement A, they may end up implementing an inferior intervention. Thanks to an anonymous reviewer for raising this possibility.

One might also argue that the greatest value view is subject to the same problem as the dominance view, namely, that random assignment can be made permissible simply by adding another high value outcome.<sup>5</sup> But notice that merely adding another outcome, provided there is uncertainty regarding the interventions' impact on this outcome, has no effect on the expected value of the interventions in question. Consider this example:

Table 4

<b>Outcome Intervention</b>	<b>X</b>	<b>Y</b>	<b>Z</b>
<b>A</b>	✓	✓	?
<b>B</b>	X	X	?

Looking only at outcomes X and Y, intervention A is clearly superior to B on the greatest value view and so policy equipoise would not permit random assignment to A and B. If we add outcome Z, this does not change the expected value of X or Y and so it would still be impermissible - per policy equipoise - to randomly assign people to A or B. On the dominance view, by contrast, the addition of Z creates uncertainty regarding A's domination of B, and so random assignment becomes permissible. For the greatest value view, the *mere addition of outcomes* does not affect judgments of international superiority.

One might argue finally that cost as well as effectiveness is relevant to interventional superiority. An intervention which is less effective but affordable is surely better than one which is

---

<sup>5</sup> Thanks to an anonymous reviewer for raising this objection.

more effective but unaffordable. Since the above account of interventional superiority focuses entirely on effectiveness, it is therefore incomplete.<sup>6</sup>

In response, I agree that considerations of cost are relevant to determining which policy should be implemented and to the question of whether a particular RCT is permissible. But we need to be careful about how we integrate such considerations. First, there is a distinct question of interventional superiority that excludes considerations of cost: which intervention most effectively realizes a set of target outcomes? An account of interventional superiority provides a way to determine when one intervention is superior to another understood in this sense. This doesn't mean, second, that cost is irrelevant to questions of the permissibility of RCTs. As I argue elsewhere, it is important to formulate a policy equipoise criterion that includes considerations of cost (MacKay 2020, 328-329). On this view, it is permissible to randomly assign people to A and B provided that neither is reasonably expected to be superior or inferior to the best proven policy that is morally and practically attainable and sustainable given the government's resources. Crucially though, this criterion also requires that we have a concept of interventional superiority that allows us to consider how interventions compare on the realization of target outcomes. The cost of different interventions is thus relevant to the permissibility of RCTs, but it needs to be integrated carefully, and we still need an account of interventional superiority that excludes such considerations.

## 5 Applying the Greatest Value View

Readers may be convinced that the greatest value view is more promising than either the dominance view or the advantage view. Practically-minded readers, however, including researchers and research sponsors, may wonder how it can be applied to real-world cases, given its abstract

---

<sup>6</sup> Thanks to Alexandra Avdeenko for raising this objection.

formulation. I will not solve this problem here. As will quickly become evident, the application of the greatest value view requires the navigation of challenging normative questions regarding which there is reasonable disagreement. My hope instead is to show that there is a path forward for applying the view and to identify questions for further research. In what follows, I first specify the challenges of application, offer some thoughts on how they may be addressed, and then consider a case.

5.1 Meeting the Challenges of Application

To apply the greatest value view to a proposed policy RCT, researchers must resolve four challenges. They must: (1) identify the target outcomes for the interventions in question; (2) determine the value of these outcomes; and (3) determine whether for either intervention, the epistemic threshold is met for any particular outcome (meaning that the intervention is reasonably expected to be superior to the alternatives in the realization of this outcome), and, if it is met, specify the size of the impact on this outcome. These questions may be helpfully illustrated on the following table:

Table 5

<b>Outcome Intervention</b>	<b>X Value?</b>	<b>Y Value?</b>	<b>Z Value?</b>
<b>A</b>	Epistemic Threshold Met?  Size and Direction of Impact?	Epistemic Threshold Met?  Size and Direction of Impact?	Epistemic Threshold Met?  Size and Direction of Impact?
<b>B</b>	Epistemic Threshold Met?  Size and Direction of Impact?	Epistemic Threshold Met?  Size and Direction of Impact?	Epistemic Threshold Met?  Size and Direction of Impact?

They must then (4) make an all things considered judgment regarding which intervention - if any - is reasonably expected to yield the set of target outcomes with the greatest value. Researchers occupy a state of intervention equipoise if there is uncertainty regarding which intervention is superior in this respect.

Researchers are well placed to complete task 3 for it involves investigating the existing evidence regarding the likely impacts of the interventions under study. Tasks 1, 2, and 4, however, require challenging normative judgments. How might researchers make them in a defensible way?

Several approaches are available. First, researchers may appeal to principles and normative frameworks developed by political philosophers and applied ethicists. Such principles and frameworks specify the responsibilities of governments and so offer resources for identifying target outcomes and ranking their importance from the standpoint of justice. For example, John Rawls's (1999; 2005) theory of justice provides a priority ranking of principles governing the meeting of basic needs, the protection and distribution of rights and liberties, access to occupations, and the distribution of income and wealth. Similarly, theories which employ social welfare functions, including utilitarianism and prioritarianism, offer principles for comparing different states of affairs from the moral point of view (Adler 2019). In general, normative theories are likely to value target outcomes in terms of the degree to which they impact people's wellbeing and/or ability to set and pursue a plan of life.

A central challenge for this approach is that normative theorists reasonably disagree about the demands of justice, and the academic world has no formal decision-making process by which theorists collectively work out which view is most likely to be correct. But this problem should not be overstated. While normative theories often conflict at the abstract level of principle, they also often converge, and, more importantly, despite divergence at the level of principles, often lend support to similar sets of policies. For example, many liberal egalitarian theories of justice conflict at

the level of principle, but nonetheless support a similar set of public policies, including universal health insurance, public education, and a social safety net generous enough to meet people's basic needs. This is not terribly surprising for the goal of normative theory is to develop a set of principles that support and justify widely held considered judgments of justice. Researchers may find that there is widespread convergence among normative theorists regarding which outcomes are target outcomes and what their relative importance is, including the promotion of personal liberty, safety, literacy, numeracy, and health, and the provision of goods such as adequate nutrition and universal health insurance.

A second approach appeals to the views of stakeholders, including policymakers and program participants. Surveys may be used to elicit stakeholders' preferences regarding the outcomes programs should be promoting, their relative value, as well as which intervention, given its reasonably expected impacts, is best on balance (Baker et al. 2021). Indeed, researchers may elicit these preferences as part of the community engagement process, which some argue is a central component of ethical research (MacKay 2024).

A central problem with this approach is that stakeholders may have ill-informed, inconsistent, or unsupported views. Preference elicitation methods typically treat responses as equally valuable, and one might reasonably object that researchers should give greater weight to preferences that are informed and supported by strong arguments (Baker et al. 2021).

Researchers may, third, appeal to the reasoned views of the public. As deliberative democrats rightly argue, the views of the public, when adequately informed and formed through a reflective and deliberative process, have normative authority (Gutmann and Thompson 2004; Habermas 1996). While real-world democratic procedures are by no means perfect, they nonetheless offer a way to resolve reasonable disagreements that gives people the opportunity to 'weigh in' on the dispute. Researchers may therefore appeal to the conclusions of deliberative processes, including

processes of public consultation or existing laws and policies in sufficiently democratic societies, to identify the target outcomes for specific programs. Regarding the ranking of such outcomes, the establishment of agencies with mandates to realize certain goals and the differential devotion of resources to specific programs may sometimes express societal judgments regarding the relative importance of different target outcomes. This approach promises to solve the problems with the second approach for deliberative practices involve participants exchanging information and giving reasons for the proposals they favor.

A central challenge for this approach is that deliberate procedures struggle to be representative of large, diverse populations (Baker et al 2021). Existing laws and policies, even if implemented in democratic societies, may reflect the views of elites rather than the broader population. Small-scale deliberative procedures, including citizen juries and deliberative forums, struggle to include a representative sample of the population, particularly if participation is voluntary.

Finally, researchers may also opt for some combination of the above approaches, taking normative theory, stakeholder preferences, and the reasoned views of the public as sources of evidence. They may find that certain target outcomes are well-supported by all three sources and use the strengths of one approach to address the weaknesses of the other. In the face of reasonable disagreement among normative theorists, approaches which appeal to the preferences or reasoned views of the public offer a way forward. In the face of uninformed and incoherent stakeholder preferences, researchers may turn to normative theory to identify views best supported by arguments.

One might argue that the above discussion is not terribly helpful for researchers, research sponsors, and policymakers. To apply the greatest value view, they must not only consider the existing empirical evidence regarding the effectiveness of interventions, but also wade into, and take

positions on, challenging questions in political philosophy. In response, it is a mistake to expect easy, technocratic solutions to questions of ethical study design. In deciding whether to randomly assign people to either intervention A or B, researchers are making *policy decisions*, if temporary, for a country’s citizens. Since, as John Rawls (2005, 3) puts it, “Justice is the first virtue of social institutions, as truth is of systems of thought,” it should be no surprise at all that researchers experimenting with social institutions should have to confront questions of justice when determining whether a study design is permissible or not.

I turn now to a case to illustrate how these approaches may be used to apply the greatest value view to a proposed policy RCT.

## 5.2 Case

Suppose that children of poor families in a particular country are too often born underweight and suffer from malnutrition and stunting. Policymakers are considering whether to implement an unconditional cash transfer program targeted at mothers with children under five years old. However, while there is high quality evidence from RCTs in similar jurisdictions showing that the cash transfer will improve families’ food security and consumption a moderate amount, it is uncertain whether it will improve children’s health and nutrition outcomes. Policymakers propose a cluster-randomized controlled trial, randomizing districts to one of three groups: (1) cash transfer; (2) cash transfer + intensive informational campaign; and (3) no intervention. Table 4 outlines the planned RCT as well as the current state of the evidence.

Table 6

<b>Outcome Intervention</b>	<b>Food Security Value: High</b>	<b>Consumption Value: High</b>	<b>Reduce Low Birth Weight Value: High</b>	<b>Reduce Malnutrition Value: High</b>	<b>Reduce Stunting Value: High</b>
-----------------------------	----------------------------------	--------------------------------	--	--	------------------------------------

<b>Cash Transfer</b>	Epistemic Threshold: Yes  Impact: Positive + Medium	Epistemic Threshold: Yes  Impact: Positive + Medium	Epistemic Threshold: No	Epistemic Threshold: No	Epistemic Threshold: No
<b>Cash Transfer + Information</b>	Epistemic Threshold: Yes  Impact: Positive + Medium	Epistemic Threshold: Yes  Impact: Positive + Medium	Epistemic Threshold: No	Epistemic Threshold: No	Epistemic Threshold: No
<b>Control</b>	Epistemic Threshold: Yes  Impact: Neutral	Epistemic Threshold: Yes  Impact: Neutral	Epistemic Threshold: No	Epistemic Threshold: No	Epistemic Threshold: No

Does this proposed RCT satisfy policy equipoise, understood here as interventional equipoise?

To apply the greatest value view, we must first identify the target outcomes for the interventions in question. The principal outcomes of the study are presented above, and these are all *target* outcomes, that is, outcomes which governments have a duty to realize. Whether we appeal to normative theory or the results of deliberative procedures, for example international human rights documents (UN General Assembly 1989), it is widely recognized that governments have a duty to ensure that children have access to adequate nutritious food, are not born underweight, are not malnourished, and are not stunted.

We must second determine the value of these outcomes. All five target outcomes are highly valuable since they play an important role in enabling children to develop fully and live healthy lives. It might be that one or more of these outcomes plays a larger instrumental role in realizing this goal and so would be more valuable for that reason, but we need not explore this issue further since both interventions perform equally well on all five outcomes and do as well or better on all five outcomes compared to the control.

We must third specify whether the epistemic threshold is met for any particular outcome, meaning that the intervention is reasonably expected to be superior to the alternatives, a task we have already completed. Finally, we must make an all things considered judgment regarding which intervention is reasonably expected to yield the set of target outcomes with the greatest value. Since both interventions are reasonably expected to be superior to the control on two outcomes and the expert community occupies a state of outcome equipoise regarding the remaining three outcomes, both interventions are superior to the control on the greatest value view. Because we don't know whether the intensive information campaign will reduce the prevalence of low birth weight, malnutrition, and stunting, the cash transfer + information is not superior to the cash transfer.

The principal conclusion of our analysis therefore is that the two interventions are superior to the control, and so the proposed RCT does not satisfy policy equipoise. This does not mean that the proposed RCT is unethical, for policy equipoise is not a necessary condition for permissible randomization (MacKay 2020, 329). But, if policymakers wish to conduct the RCT as planned, they would need to argue that the cash transfer is scarce and that the communities in question have equally strong claims to it, implying that allocation by lottery is permissible (MacKay and Cohn 2023, 18-21; MacKay, 2020, 333-348).

My hope is that this case illustrates how the greatest value view may be applied to real-world RCTs. No doubt, this type of case is unlikely to be the most challenging that researchers may face. If researchers plan to study two or more interventions which meet the epistemic threshold on different outcomes, they will need to make difficult value judgments regarding which target outcome - if any - is more important. Fortunately, these types of cases are likely to be rare. In addition, while these types of cases are challenging, a virtue of the greatest value view is that it recognizes the challenge and provides researchers with the tools to work through it. Alternatives, including the dominance

and advantage views simply fail to acknowledge that some target outcomes may be more important than others from the standpoint of justice.

## Conclusion

My aim in this paper has been to contribute to the development of an account of interventional superiority, that is, to specify the conditions under which interventional equipoise is disturbed and it is permissible to randomly assign people to two or more interventions. I have defended the greatest value view, according to which one intervention is superior to another if and only if it is reasonably expected to yield a more valuable set of target outcomes. The greatest value view does not offer a complete solution to the problem of interventional superiority, for it does not address the question of when outcome equipoise is disturbed. I leave the task of addressing this task for a future paper.

The greatest value view follows nicely from the premises with which we started above. Governments have duties to realize target outcomes for residents and they discharge these duties by implementing policies. The purpose of policy research is to provide social institutions with the knowledge they require to design effective policies. In the context of policy research and more broadly, one intervention is better than another if it better enables governments to discharge their duties of justice, that is, if it better enables governments to realize a more valuable set of target outcomes.

## Acknowledgements

This paper has benefitted greatly from the comments and questions of many people. First, I am grateful to three anonymous reviewers for their insightful comments and criticism, and to Alexandra

Avdeenko for her excellent editorial guidance. Second, I am grateful to participants at the 2022 New Scholarship in Bioethics Annual Symposium (Kirstin Borgerson, Michael Garnett, Joseph Millum, and Marika Warren) for their detailed, thoughtful feedback. Finally, I thank audiences at the 2022 Annual Meeting of the Philosophy, Politics, & Economics Society Meeting and the 2022 Annual Conference of the American Society of Bioethics and Humanities.

## References

Abramowicz, Michel and Ariane Szafarz. “Ethics of RCTs: Should Economists Care About Equipoise?” In *Randomized Control Trials in the Field of Development: A Critical Perspective*, edited by Florent Bédécarrats, Isabelle Guérin, and François Roubaud, 280-292. New York: Oxford University Press, 2020.

Adler, Matthew D. 2019. *Measuring Social Welfare: An Introduction*. New York: Oxford University Press.

Asiedu, Edward, Dean Karlan, Monica P. Lambon-Quayefio and Christopher R. Udry. 2021. “A Call for Structured Ethics Appendices in Social Science Papers.” *PNAS* 118 (29): e2024570118.

Baele, Stéphane J. 2013. “The Ethics of New Development Economics: Is the Experimental Approach to Development Economics Morally Wrong?” *The Journal of Philosophical Economics* 7, no. 1: 2-42.

Baker, Rachel, Helen Mason, Neil McHugh, and Cam Donaldson. 2021. "Public Values and Plurality in Health Priority Setting: What To Do When People Disagree and Why We Should Care About Reasons as Well as Choices." *Social Science & Medicine* 277 (May): 113892.

Chard, J.A., and R.J. Lilford. 1998. "The Use of Equipoise in Clinical Trials." *Social Science and Medicine* 47, no. 7: 891-898.

Freedman, Benjamin. 1987. "Equipoise and the Ethics of Clinical Research." *New England Journal of Medicine* 317, no. 3: 141-145.

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

Gifford, Fred. 2000. "Freedman's 'Clinical Equipoise' and 'Sliding-Scale All-Dimensions-Considered Equipoise.'" *Journal of Medicine and Philosophy* 25, no. 4: 399-426.

Gutmann, Amy, and Dennis Thompson. 2004. *Why Deliberative Democracy?* Princeton: Princeton University Press.

Habermas, Jürgen. 1996. *Between Facts and Norms: Contributions to a Discourse Theory of Law and Democracy*. Translated by William Rehg. Cambridge, MA: MIT Press.

Handa, Sudhanshu, David Seidenfeld, Benjamin Davis, Gelson Tembo, and Zambia Cash Transfer Evaluation Team. 2016. "The Social and Productive Impacts of Zambia's Child Grant." *Journal of Policy Analysis and Management* 35 (2): 357-387.

Lilford, Richard J., and Jennifer Jackson. 1995. "Equipose and the Ethics of Randomization." *Journal of the Royal Society of Medicine* 88: 552-559.

London, Alex John. 2024. "Equipose: Integrating Social Value and Equal Respect in Research with Humans." In *The Oxford Handbook of Research Ethics*. Edited by Ana Iltis and Douglas MacKay, 216-235. New York: Oxford University Press.

London, Alex John. 2022. *For the Common Good: Philosophical Foundations of Research Ethics*. New York: Oxford University Press.

MacKay, Douglas. 2024. "The Ethics of Public Policy Research: Lessons from Clinical Research Ethics." In *The Oxford Handbook of Research Ethics*, edited by Ana S. Iltis and Douglas MacKay, 445-472. New York: Oxford University Press.

MacKay, Douglas. 2020. "Government Policy Experiments and the Ethics of Randomization." *Philosophy and Public Affairs* 48 (1): 319-352.

MacKay, Douglas. 2018. "The Ethics of Public Policy RCTs: The Principle of Policy Equipose." *Bioethics* 32 (1): 59-67.

MacKay, Douglas. 2015. "Standard of Care, Institutional Obligations, and Distributive Justice." *Bioethics* 29 (4): 262-263.

MacKay, Douglas and Averi Chakrabarti. 2019. "Government Policy Experiments and Informed Consent." *Public Health Ethics* 12, no. 2: 188-201.

MacKay, Douglas, and Emma Cohn. 2023. "Public Policy Experiments without Equipoise: When is Randomization Fair?" *Ethics and Human Research* 45 (1): 15-28.

Petticrew, M., M. McKee, K. Lock, J. Green and G. Phillips. "In Search of Social Equipoise." *BMJ* 346 (2013).

Pierson, Leah, and Joseph Millum. 2018. "Health Research Priority Setting: The Duties of Individual Funders." *The American Journal of Bioethics* 18 (11): 6-17.

Rawls, John. 2005. *Political Liberalism*. Expanded Edition. New York. Columbia University Press.

Rawls, John. 1999. *A Theory of Justice*. Revised Edition. Cambridge, MA: Harvard University Press.

UN General Assembly, *Convention on the Rights of the Child*, 20 November 1989, United Nations, Treaty Series, vol. 1577, p. 3.

Ziliak, Stephen T., and Edward R. Teather-Posadas. 2016. "The Unprincipled Randomization Principle in Economics and Medicine." In *The Oxford Handbook of Professional Economic Ethics*, edited by George F. DeMartino and Deirdre McCloskey, 423-454. New York: Oxford University Press.