

Douglas MacKay (dmackay@email.unc.edu)
Associate Professor, Department of Public Policy, UNC-Chapel Hill

Government Policy Experiments and the Ethics of Randomization
Philosophy and Public Affairs 48 (2020): 319-352

Abstract

Governments are increasingly using randomized controlled trials (RCTs) to evaluate policy interventions. RCTs are often understood to provide the highest quality evidence regarding the causal efficacy of an intervention. While randomization plays an essential epistemic role in the context of policy RCTs however, it also plays an important distributive role. By randomly assigning participants to either the intervention or control arm of an RCT, people are subject to different policies and so, often, to different types and levels of benefits. In this paper, I identify one necessary condition as well as a set of sufficient conditions for the permissible use of random assignment by government agencies. I argue first that random assignment is permissible only if it is consistent with governments' duty to realize morally important outcomes. I argue second that random assignment is permissible in cases where investigators are in a state of genuine equipoise regarding all arms of the experiment and the policy to which people have a claim of justice. Finally, I defend a set of conditions under which random assignment is permissible in cases where one or more arms of a policy RCT are reasonably expected to be either superior or inferior to this policy.

Keywords: Research Ethics; Randomized Controlled Trials; Distributive Justice; Randomization; Public Policy Research

Text

Governments are increasingly using randomized controlled trials (RCTs) to evaluate policy interventions.¹ RCTs are often understood to provide the highest quality evidence regarding the causal efficacy of an intervention.² By randomly assigning participants to

Previous versions of this paper were presented at the 2nd Annual Philosophical Issues in Research Ethics Workshop, the 2019 Annual Conference of the American Society for Bioethics and Humanities, the Bioethics at UNC Workshop, the Instituto de Ciencia Política at the Pontificia Universidad Católica de Chile, and the 2nd Annual Public Issues and Public Reason Conference at Carleton University. I thank audiences at these events for their questions. For helpful comments and conversations, I am grateful to Angela Ballantyne, Ariella Binik, Kirstin Borgerson, Jurgen De Wispelaere, Sudhanshu Handa, Steven Hemelt, Jonathan Herington, Rebecca Kreitzer, Emily Largent, Benjamin Levinstein, Alex John London, Joseph Millum, Vida Panitch, Cristian Pérez Muñoz, Bridget Pratt, Katherine Saylor, Rieke van der Graaf, Charles Weijer, and Danielle Wenner. Finally, I thank three Associate Editors of *Philosophy and Public Affairs* for their challenging and insightful feedback, and their deep commitment to helping me improve the paper.

¹ For good overviews of the increasing use of RCTs, see Steven D. Levitt and John A. List, “Field Experiments in Economics: The Past, the Present, and the Future,” *European Economics Review* 53 (2009): 1-18; Judith M. Gueron and Howard Rolston, *Fighting for Reliable Evidence* (New York: Russell Sage Foundation, 2013); Delia Baldassarri and Maria Abascal, “Field Experiments Across the Social Sciences,” *Annual Review of Sociology* 43 (2017): 41-73; Andrew Leigh, *Randomistas: How Radical Researchers are Changing Our World* (New Haven: Yale University Press, 2018); and Luciana de Souza Leão and Gil Eyal, “The Rise of Randomized Controlled Trials (RCTs) in International Development in Historical Perspective,” *Theory and Society* 48 (2019): 383-418.

² William R. Shadish, Thomas D. Cook, and Donald T. Campbell, *Experimental and Quasi-Experimental Designs for Generalized Causal Inference* (Boston: Houghton Mifflin, 2002), 226-278. For a critical evaluation of RCTs, see Angus Deaton and Nancy Cartwright, “Understanding and Misunderstanding Randomized Controlled Trials,” *Social Science & Medicine* 210 (2018): 2-21.

intervention and control groups, for example, investigators can minimize selection bias – i.e. systematic differences between those subject and not subject to the intervention that are potentially correlated with the outcome of interest. With observational studies, investigators can never be certain whether outcomes are the result of the intervention, or of some distinctive characteristic(s) for which they have not controlled.

Randomization thus plays an essential *epistemic* role in the context of policy experiments. However, it also plays an important *distributive* role as participants are randomly assigned to different policies and so, often, to different types and levels of benefits. This is troubling when some are randomly assigned to a promising intervention, while others are assigned to the seemingly inferior status quo policy or receive nothing. For example, some may receive cash transfers while others do not;³ some may receive housing vouchers while others do not;⁴ some students may be assigned to smaller class sizes while others remain in larger classes;⁵ and some children may receive a deworming treatment while others do not.⁶

³ For example, see T. Paul Schultz, “School Subsidies for the Poor: Evaluating the Mexican Progresa Poverty Program,” *Journal of Development Economics*, 74 (2004): 199-250; and Sudhanshu Handa et al., “The Social and Productive Impacts of Zambia’s Child Grant,” *Journal of Policy Analysis and Management* 35 (2016): 357-387.

⁴ Raj Chetty, Nathaniel Hendren, and Lawrence F. Katz, “The Effects of Exposure to Better Neighborhoods on Children: New Evidence from the Moving to Opportunity Experiment,” *American Economic Review* 106 (2016): 855-902.

⁵ Word, E., Johnston, J., Bain, H., Fulton, B.D., Zaharias, J.B., Achilles, C.M., Lintz, M.N., Folger, J., and Breda, C. (1990). *The State of Tennessee’s student/teacher achievement ratio (STAR) Project*. Tennessee: Tennessee Board of Education.

⁶ Edward Miguel and Michael Kremer, “Worms: Identifying Impacts on Education and Health in the Presence of Treatment Externalities,” *Econometrica* 72 (2004): 159-217.

It would be convenient if random assignment were not only epistemically ideal but also ideal from the standpoint of distributive justice. However, commentators rightly question whether it is always fair to allocate access to a promising intervention by lottery.⁷ Since research designs must be ethical in their treatment of participants as well as scientifically robust, it is imperative to determine when – if ever – it is fair to employ random assignment. There are moral constraints on the means we may use to produce knowledge, even when this knowledge promises to greatly facilitate governments’ ability to promote people’s wellbeing or secure a more just social order.

In this paper, I identify one necessary condition as well as a set of sufficient conditions for the permissible use of random assignment in policy RCTs. While not necessarily exhaustive,

⁷ See Edward Diener and Rick Crandall, *Ethics in Social and Behavioral Research* (Chicago: University of Chicago Press, 1978), 135-140; Rebecca Kukla, “Resituating the Principle of Equipose: Justice and Access to Care in Non-ideal Conditions,” *Kennedy Institute of Ethics Journal* 17 (2007): 171-202; Stéphane J. Baele, “The Ethics of New Development Economics: Is the Experimental Approach to Development Economics Morally Wrong?” *The Journal of Philosophical Economics* VII (2013): 2-42; Dawn Langan Teele, “Reflections on the Ethics of Field Experiments,” in *Field Experiments and Their Critics: Essays on the Uses and Abuses of Experimentation in the Social Sciences*, Edited by Dawn Langan Teele (New Haven: Yale University Press, 2014): 115-140; Rachel Glennerster and Shawn Powers, “Balancing Risk and Benefit: Ethical Tradeoffs in Running Randomized Evaluations,” in *The Oxford Handbook of Professional Economic Ethics*, Edited by George DeMartino and Deirdre McCloskey (New York: Oxford University Press, 2016): 367-401; Stephen T. Ziliak and Edward R. Teather-Posadas, “The Unprincipled Randomization Principle in Economics and Medicine,” in *The Oxford Handbook of Professional Economic Ethics*, Edited by George DeMartino and Deirdre McCloskey (New York: Oxford University Press, 2016): 423-448; Douglas MacKay, “The Ethics of Public Policy RCTs: The Principle of Policy Equipose,” *Bioethics* 32 (2018): 59-67; and Margarita S. Rayzberg, “Fairness in the Field: The Ethics of Resource Allocation in Randomized Controlled Experiments,” *Science, Technology, & Human Values*, 44 (2019): 371-398.

these conditions provide guidance for the principal scenarios in which government agencies consider and employ RCTs. Part 1 provides a brief account of the nature of government policy RCTs and the obligations that are directly relevant to random assignment. I argue that random assignment is permissible only if it is consistent with governments' duty to realize morally important outcomes. Parts 2 and 3 identify a set of sufficient conditions for permissible randomization. Part 2 re-interprets the familiar norm of clinical equipoise for policy RCTs, arguing that random assignment is permissible when the social science community occupies a state of uncertainty regarding the interventions under study. Part 3 considers whether it is ever permissible to use random assignment when equipoise is not satisfied – i.e. when there is strong evidence that one intervention is superior to the alternatives. I argue that it is and defend a set of conditions under which doing so is permissible. One interesting conclusion of my analysis is thus that equipoise is not a necessary condition of ethical randomization as is sometimes proposed.⁸ Part 4 refines the necessary condition introduced in part 1 and gives some examples of RCTs that violate it.

My analysis is limited in two ways. First, I focus on experiments that are conducted or authorized by governments rather than non-governmental organizations. Governments have distinct duties to their citizens that are directly relevant to the conduct of research and so the ethics of randomization in government policy experiments is deserving of separate treatment.⁹

⁸ Indeed, I defend this position in “The Ethics of Public Policy RCTs.”

⁹ For discussion of the importance of this distinction in the context of clinical research ethics, see Douglas MacKay, “Standard of Care, Institutional Obligations, and Distributive Justice,” *Bioethics* 29 (2015): 262-273; and Leah Pierson and Joseph Millum, “Health Research Priority Setting: The Duties of Individual Funders,” *American Journal of Bioethics* 18 (2018): 6-17. For an overview of governments' involvement in policy RCTs, see de Souza

Second, I focus on the permissibility of *random assignment*. Policy RCTs must no doubt meet a number of additional requirements to be permissible – e.g. informed consent, social value, and fair participant selection – but a full specification of these requirements is beyond the scope of this paper.¹⁰

1 Government Policy Experiments and the Requirements of Justice

Governments are increasingly making use of experiments rather than observational research to evaluate public policies. Unlike observational research, policy experiments involve the introduction of a policy intervention, or the alteration of an existing policy intervention, as part of a systematic investigation aimed at generating knowledge about the intervention’s impact. They may also involve the collection of data that would not otherwise be collected. Policy RCTs are policy experiments that employ random assignment.¹¹

Government policy experiments are policy experiments that are conducted or authorized by a government agency. For example, a government agency may introduce and evaluate a

Leão and Eyal, “The Rise of Randomized Controlled Trials (RCTs) in International Development in Historical Perspective.”

¹⁰ For a discussion of when informed consent is necessary for government policy experiments, see Douglas MacKay and Averi Chakrabarti, “Government Policy Experiments and Informed Consent,” *Public Health Ethics* 12 (2019): 188-201.

¹¹ My definition of a policy RCT is very similar to David Greenberg and Mark Shroder’s definition of a *social experiment*. According to them, a social experiment involves: (1) random assignment, (2) a policy intervention, (3) follow-up data collection, and (4) evaluation of the intervention’s effectiveness. See David Greenberg and Mark Shroder, *The Digest of Social Experiments*, Third Edition (Washington, DC: The Urban Institute Press, 2004), 4.

particular intervention itself or grant permission to a non-governmental organization or private firm to do so. Governments also authorize policy experiments when they fund them.

One might wonder why government agencies' use of RCTs raises any ethical questions at all. With clinical RCTs, ethical questions concerning randomization are sometimes thought to arise because investigators are physicians and so have a fiduciary obligation to promote participants' health-related interests.¹² Social scientists, by contrast, have no such obligations and legitimate governments possess the authority to make policy decisions. However, governments possess duties of justice and these obligations are directly relevant to the conduct of policy research, placing constraints on the types of policies to which they may randomly assign people.

Consider that on prominent liberal theories of justice, governments have duties to protect residents' liberty, promote their health, combat poverty, and ensure access to high-quality educational opportunities.¹³ I shall use the concept of *target outcomes* to refer to 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. Target outcomes may include specific levels of health insurance coverage, specific degrees of personal security and liberty, certain literacy outcomes, or a minimal standard of living. Target outcomes may vary from state to state,

¹² For example, see Samuel Hellman and Deborah S. Hellman, "Of Mice but Not Men – Problems of the Randomized Clinical Trial," *The New England Journal of Medicine* 324 (1991): 1585-1589.

¹³ For a range of different theories that advocate duties of this kind, see John Rawls, *A Theory of Justice*, Revised Edition (Cambridge: Harvard University Press, 1999); Ronald Dworkin, *Sovereign Virtue: The Theory and Practice of Equality* (Cambridge, Harvard University Press, 2000); Norman Daniels, *Just Health: Meeting Needs Fairly* (New York: Cambridge University Press, 2008); and John Tomasi, *Free Market Fairness* (Princeton: Princeton University Press, 2012).

depending on cultural variation and resources, and governments may have different duties towards citizens and non-citizens.

To realize any particular target outcome, governments must implement laws and policies that satisfy a number of criteria. First, they must implement policies that are evidence-based. K-12 students, for example, have a claim of justice to education policies that can be reasonably expected to realize literacy and numeracy outcomes. Governments must also, second, implement policies that are consistent with people's rights. For example, they must not implement policies that, while effective in realizing a particular target outcome, unjustly restrict people's liberty. Third, governments must implement policies that are consistent with the realization of other target outcomes. They must not put so many resources into achieving educational outcomes that they are unable to adequately promote people's health.

For any particular target outcome therefore, governments should implement policies that are: (1) evidence-based; (2) consistent with people's rights; and (3) consistent with the realization of other target outcomes. Governments should therefore implement policies that are best proven (to satisfy 1), but also, to borrow a concept from Alex John London, morally and practically attainable and sustainable (to satisfy 2 and 3).¹⁴ A policy is morally and practically attainable and sustainable, I suggest, if and only if:

1. It is consistent with people's rights; and

¹⁴ Alex John London, "Equipose and International Human-Subjects Research," *Bioethics* 15 (2001), 327.

2. It can be implemented for an appropriate period of time given a just system of resource procurement and allocation.¹⁵

The policy that satisfies these conditions is thus the most effective, rights-respecting policy a government could implement to realize a particular target outcome if it acted justly and effectively in the procurement and allocation of resources. Call this the *best proven morally and practically attainable and sustainable policy*, or, *BPA* (best proven attainable) policy for short.¹⁶

BPA policies should not be understood as Platonic ideals that apply to all societies. Societies differ reasonably in terms of their norms, political cultures, political institutions, and established policy regimes, among other factors. Societies also differ in terms of bureaucratic expertise, fiscal capacities, and social capital. A set of policies that are effective at realizing certain health outcomes in one country should not therefore be assumed to be BPA policies for another. Such policies may not work in the same way in the latter country, and the latter country may lack the fiscal and administrative capacities necessary to properly implement them. BPA policies are thus likely to be highly context specific, and in many cases may simply be the status quo policies. For the status quo policy not to be the BPA policy, after all, there must be some policy that is known to be more effective for the particular context in question, and that can be

¹⁵ MacKay, “The Ethics of Public Policy RCTs,” 62. I use the phrase “appropriate period of time” rather than “long-term” to recognize that there may be cases when it is reasonable to implement a policy for only a short period of time rather than, as is often the case, for the foreseeable future.

¹⁶ Ibid. As will become clear below, this concept of the BPA policy plays an analogous role in discussions of the ethics of policy RCTs as the concept of “standard of care” plays in discussions of the ethics of clinical RCTs. For a helpful overview of this latter concept, see Alex John London, “The Ambiguity and the Exigency: Clarifying ‘Standard of Care’ Arguments in International Research,” *Journal of Medicine and Philosophy* 25 (2000): 379-397.

implemented sustainably given the government's expected fiscal and bureaucratic capacities – a high standard to meet.

Governments' obligations to realize target outcomes and enact BPA policies are relevant for the conduct and design of policy RCTs. First, governments may have duties of justice to conduct policy research. Even in reasonably just societies, there is likely to be a gap between the outcomes governments have a duty to realize (and which are in principle realizable), and the outcomes realized by existing policies. For example, it is arguable that governments have a duty to secure fair equality of opportunity for their citizens, that is, to ensure that people's life prospects are not unduly influenced by their parent's socio-economic status.¹⁷ By comparison with other countries, Canada does a reasonably good job of satisfying this principle.¹⁸ However, there is still a sizeable gap between the outcomes Canada's existing policies realize, and the outcomes demanded by fair equality of opportunity. The federal and provincial governments of Canada should close this gap to the extent that it is feasible to do so, and one strategy is to conduct research to identify policies that would be more effective in improving the mobility rates of low-income children.

As a concrete example of such research in the U.S., the Seattle and King County Housing authorities have recently partnered with researchers to conduct an RCT to identify the reasons why low-income families do not move to low-poverty neighborhoods which offer expanded

¹⁷ Rawls, *A Theory of Justice*, 73.

¹⁸ OECD, *A Broken Social Elevator? How to Promote Social Mobility* (OECD Publishing, Paris, 2018) Available at: <https://doi.org/10.1787/9789264301085-en>.

opportunities for upward income mobility.¹⁹ Housing voucher recipients in Seattle and King County were randomized to either receive or not receive a suite of services designed to reduce barriers to moving to high-upward-mobility neighborhoods, including help in the search for homes and short-term financial assistance. Researchers found that while only 14% of families in the control group moved to a high-upward-mobility neighborhood, 54% of families in the intervention group did so, a finding useful for the design of affordable housing policies for low-income residents.²⁰ It is arguable that the Seattle and King County Housing Authorities have an obligation to conduct this type of RCT as a way to improve the economic opportunity of low-income Americans.

Second, and more pertinently, governments' duty to realize target outcomes places constraints on the use of random assignment: random assignment is permissible only if it is consistent with this duty. This condition would not be terribly constraining if it were permissible for governments to realize target outcomes to different degrees for different people. However, for many target outcomes, governments will likely have a duty to realize them to the same extent for all residents. For example, governments arguably have obligations to ensure that all residents have access to the same minimal level of health care, enjoy the same minimal level of personal security, and have access to adequate nutrition. Even in cases where there are target outcomes that are specific to a subset of the population, any policy RCT is likely to focus on members of that subset. For example, governments may have special obligations to children to realize

¹⁹ Peter Bergman, Raj Chetty, Stefanie DeLuca, Nathaniel Hendren, Lawrence Katz, and Christopher Palmer, "Creating Moves to Opportunity: Experimental Evidence on Barriers to Neighborhood Choice," NBER Working Paper No. 26164 (August 2019): 1-46.

²⁰ *Ibid.*, 2.

educational outcomes, but any RCT evaluating educational interventions is only going to enroll children, not adults.

Government agencies may only randomly assign participants to different interventions therefore if doing so is consistent with their duty to realize target outcomes for participants and non-participants. Call this necessary condition of permissible random assignment *NC*.

NC: Government agencies may authorize or conduct policy experiments in which participants are randomly assigned to different interventions only if such assignment is consistent with their duty to realize target outcomes for participants and non-participants. Since, for most policy RCTs, agencies will have a duty to realize the same target outcomes for participants and non-participants, and since RCTs involve assigning people to different interventions which may be expected to realize such outcomes to different extents, NC would seem to place strict constraints on the permissible use of random assignment. Indeed, one might wonder whether random assignment is ever permissible.

In the next section of the paper, I begin to outline a set of sufficient conditions for permissible randomization. Doing so will also make it possible to refine NC and outline examples of RCTs that violate it, a task I return to in part 4.

2 Permissible Randomization 1: Policy Equipose

To determine when random assignment is permissible, it is helpful to consider when it is permissible for governments to use lotteries to distribute goods. Peter Stone has made important headway on the morally appropriate role of lotteries in decision-making and the allocation of

scarce goods.²¹ Stone focuses on fair lotteries – i.e. lotteries for which equal probabilities can be assigned to outcomes – since they are more unpredictable than weighted lotteries and unpredictability is the central contribution of lotteries to decision-making.²² My analysis also focuses on fair lotteries as the method of random assignment in policy RCTs, though I briefly discuss the use of weighted lotteries in the conclusion.

Stone argues that the purpose of fair lotteries is to sanitize the decision-making process from the influence of reasons. The distinguishing feature of lotteries as a decision-making process, Stone claims, is that the outcomes of the process are unpredictable.²³ When agents make use of reasons to make decisions, by contrast, the outcome of the process is predictable – it is the outcome that is responsive to the reasons in question. This role of lotteries in decision-making is valuable, Stone argues, in cases of indeterminacy, that is, cases where the process of reasoning does not pick out one option as preferable.²⁴ Where decision-makers are in a state of indeterminacy, fair lotteries are a permissible way to make decisions; where there is a possibility of appealing to bad reasons, fair lotteries are obligatory.²⁵

Stone's account of the role of lotteries is useful for identifying when random assignment is permissible. Since it is permissible to use fair lotteries in cases of indeterminacy, it should be permissible for government agencies to employ random assignment when they are in a state of

²¹ See Peter Stone, *The Luck of the Draw: The Role of Lotteries in Decision Making* (New York: Oxford University Press, 2011).

²² *Ibid*, 24.

²³ *Ibid*, 35-36.

²⁴ *Ibid*, 30-31.

²⁵ *Ibid*, 39-40.

indeterminacy. One such state is when government agencies lack sufficiently strong reasons to prefer that participants be subject to either the intervention or control arm. Another is when they have sufficiently strong reasons to prefer the intervention arm but are in a state of indeterminacy regarding who should have access to it. I explore the former condition in this section before turning to the latter in the next.

From the standpoint of their duty to realize target outcomes, government agencies lack sufficiently strong reasons to prefer either the intervention arm or the control arm of an RCT when there is reasonable disagreement within the social science community regarding which arm is superior for the realization of these outcomes. The simplest example of this type of case is where there is one target outcome and there is reasonable disagreement within the social science community regarding which arm is more effective at realizing it. There may also be more complex cases, however, where the interventions under study are expected to realize multiple target outcomes, A and B, each intervention is expected to realize A and B to different degrees, but there is reasonable disagreement regarding which balance of A and B is preferable from the standpoint of justice.

Where government agencies are in this state of indeterminacy between the interventions under study, they are in a state of *equipoise*. The norm of *clinical* equipoise is a principal norm governing the design of clinical RCTs.²⁶ For some scholars, it promises to reconcile physician-researchers' fiduciary obligation to promote the medically-related interests of patient-subjects,

²⁶ See Benjamin Freedman, "Equipoise and the Ethics of Clinical Research," *The New England Journal of Medicine* 317 (1987): 141-145.

with the need to randomly assign them to an intervention.²⁷ *Policy* equipoise functions in a similar way, offering government agencies a way to reconcile random assignment with their duty to realize target outcomes for people.²⁸ In cases of reasonable disagreement regarding the relative efficacy of the interventions under study, government agencies lack sufficiently strong reasons to conclude that one intervention will better realize the outcomes in question than the other. This need not mean that the reasons in support of the interventions in question are equally balanced, but rather that the social science community is collectively uncertain regarding which intervention is superior.²⁹ Some experts have good reasons to think that intervention C is superior to D while others have good reasons to think that D is superior to C.

The mere fact that government agencies are in a state of equipoise regarding the interventions under study does not entail that participants are treated justly. People have claims of justice to be subject to BPA policies, and so government agencies treat participants justly when there is reasonable disagreement regarding all arms of the study and the BPA policy. The following is therefore a sufficient condition of permissible random assignment:

Sufficient Condition (SC) 1: Government agencies may randomly assign participants to different policy interventions if they are in a state of genuine equipoise regarding all arms of the study and the BPA policy.

²⁷ P.B. Miller and C. Weijer, "Trust Based Obligations of the State and Physician-Researchers to Patient-Subjects," *Journal of Medical Ethics* 32 (2006): 542-547.

²⁸ MacKay, "The Ethics of Public Policy RCTs," 63.

²⁹ For further discussion on this point, see Alex John London, "Equipoise: Integrating Social Value and Equal Respect in Research with Humans," in *The Oxford Handbook of Research Ethics*, edited by Douglas MacKay and Ana Iltis (New York: Oxford University Press, forthcoming).

Where this principle is satisfied, no participant is subject to a policy for which there are sufficiently strong reasons to judge that it will realize the relevant target outcomes to a greater or lesser degree than the BPA policy.

Further work is necessary to identify when RCTs satisfy this condition and when they do not. In particular, it is necessary to provide an account of the “epistemic threshold” that an intervention must meet if it is to be judged superior to alternative interventions with respect to the realization of target outcomes.³⁰ Such an account would provide policymakers with a systematic way to evaluate and weigh the different types of evidence that are relevant to the interventions in question, and specify a threshold such evidence must meet for equipoise to be disturbed, and thus for government agencies to have sufficiently strong reasons to favor one intervention over another.

To my knowledge, there has been no systematic study regarding how often something like SC1 is satisfied by policy RCTs, or even whether equipoise is considered a binding norm. Indeed, the concept of “equipoise” is rarely mentioned in the rather scant literature on the ethics of policy RCTs, despite its importance to the ethics of clinical RCTs.³¹ There is evidence however that some people are uncomfortable with RCTs that satisfy SC1. A recent study found that people judge a randomized experiment comparing two interventions – A and B – in a state

³⁰ See Ibid.

³¹ For exceptions, see Kukla, “Resituating the Principle of Equipoise;” Baele, “The Ethics of New Development Economics;” Ziliak and Teather-Posadas, “The Unprincipled Randomization Principle in Economics and Medicine;” and MacKay, The Ethics of Public Policy RCTs.”

of equipoise as less appropriate than the universal imposition of either A or B.³² While the study authors suggest this preference is unreasonable, one might argue that it reflects a commitment to the principle of equal treatment. For RCTs that satisfy SC1 may lead to one group being subject to an intervention that is ex post inferior or superior, whereas universal imposition avoids this possibility. Universal imposition may also better *express* the value of equality than an RCT.³³

We should reject this interpretation of equal treatment, however, since it has counter-intuitive implications. Suppose a proposed RCT evaluating A and B satisfies SC1 and, to simplify, that the mean expected target outcomes for individuals, based on existing evidence, is 6 for both policies. Suppose that in actual fact A yields an outcome of 8 whereas B yields an outcome of 6. According to the above interpretation of equal treatment, this RCT is objectionable because some participants would be subject to an ex post inferior policy – B – while others would be subject to an ex post superior one – A. Instead, either universal imposition of A or universal imposition of B is preferable to an RCT since both options avoid this inequality. But, if universal imposition of B is superior to the RCT, then we are committed to preferring an option where no participants are better off and some are worse off (those who would have been subject to A in the RCT).³⁴ Consider too that we need not abandon the principle of equal treatment to prefer the RCT to universal imposition of either A or B. Since the RCT satisfies SC1, all

³² See Michelle N. Meyer, Patrick R. Heck, Geoffrey S. Holtzman, Stephen M. Anderson, William Cai, Duncan J. Watts, and Christopher F. Chabris, “Objecting to Experiments that Compare Two Unobjectionable Policies or Treatments,” *PNAS* 116 (2020): 10723-10728.

³³ Thanks to an anonymous Associate Editor for raising these objections.

³⁴ Derek Parfit, “Equality and Priority,” *Ratio* X (1997), 211.

participants are treated equally in the sense of being subject to policies which are ex ante equivalent (roughly) with respect to the realization of target outcomes.

Indeed, even proponents of the former interpretation of equal treatment are unlikely to favor universal imposition over RCTs that satisfy SC1. Such egalitarians often address the counter-intuitive implications of this interpretation by noting that while equal states of affairs are preferable to unequal states of affairs in one respect, the latter may nonetheless be on balance preferable when they offer a better balance of equality and other values – e.g. utility.³⁵ This matters since a policy of permitting RCTs promises much greater utility than a policy of universal imposition since the latter would greatly hinder policymakers’ ability to learn about the effectiveness of different policies. In the above example, if RCTs are prohibited, policymakers might implement the inferior policy B and never come to know that A is superior. An RCT, by contrast, permits the possibility of unequal outcomes for the duration of the study, but then provides policymakers with evidence regarding which policy should be implemented long-term for all.

There are also cases where it is permissible for government agencies to conduct RCTs in which the experimental intervention is expected to be inferior to the BPA policy, but is in a state of equipoise with the status quo policy. Consider a situation in which a government agency wishes to conduct a policy RCT, but where citizens are not subject to the BPA policy:

Healthy Incentives Pilot: In 2011, the U.S. Department of Agriculture’s (USDA) Food and Nutrition Service designed the Healthy Incentives Pilot (HIP) with the aim of determining whether Supplemental Nutrition Assistance Program (SNAP) recipients,

³⁵ Larry Temkin, “Equality, Priority, and the Levelling Down Objection,” in *The Ideal of Equality*, edited by M. Clayton and A. Williams (New York: St. Martin’s Press, 2000): 81-125.

would increase their consumption of fruits and vegetables in response to financial incentives.³⁶ Participants assigned to the HIP received an incentive of \$0.30 on their electronic benefit transfer (EBT) card for each dollar of SNAP benefits they spent on fruits and vegetables from participating retailers. In 2011, 7,500 SNAP recipients in Hampden County Massachusetts were randomly selected to participate in the experiment. These participants received specially marked EBT cards as well as multiple notifications regarding their enrollment.³⁷ Data were collected from a variety of sources, including through voluntary surveys with samples of both HIP and non-HIP households, from EBT transaction data, and from surveys and observations of participating food retailers.³⁸

Assume that the current SNAP policy is not the BPA policy since its benefits are too low to realize the relevant target outcome – secure access to food adequate to meet nutritional needs – and a more generous policy is affordable in the U.S.³⁹ Given this assumption, the HIP pilot compared two policies that were expected to be inferior to the BPA policy, the status quo SNAP policy and the experimental healthy incentives policy. Assume further that the USDA Food and Nutrition Service occupied a state of genuine equipoise regarding the status quo SNAP policy

³⁶ Bartlett, S., Klerman, J., Olsho, L., Logan, C., Blocklin, M., Beauregard, M., Enver, A., Wilde, P., and Owens, C., (2014). *Evaluation of the Healthy Incentives Pilot (HIP): Final Report*. Available at: <https://fns-prod.azureedge.net/sites/default/files/ops/HIP-Final.pdf>.

³⁷ Ibid, 25.

³⁸ Ibid, 26-38.

³⁹ Kranti Mulik and Lindsey Haynes-Maslow, “The Affordability of MyPlate: An Analysis of SNAP Benefits and the Actual Cost of Eating According to the Dietary Guidelines,” *Journal of Nutrition Education and Behavior* 49 (2017): 623-631.

and the experimental healthy incentives policy.⁴⁰ Was random assignment between these two interventions permissible?

The U.S. government was not in a state of indeterminacy regarding the current SNAP policy, the healthy incentives policy, and the BPA policy. It ought therefore to have implemented the BPA policy. But the USDA's Food and Nutrition Service did not have the authority to do so since Congress is responsible for setting SNAP benefit levels. Although the USDA has the authority to design and conduct experiments on proposed changes to SNAP, it does not have the option of implementing the BPA policy. Participants in the HIP, like all SNAP recipients, were thus treated unjustly by the U.S. government since they were not subject to the BPA policy. But the USDA Food and Nutrition Service did not commit an additional wrong against participants by randomly assigning them to either the current SNAP policy or the healthy incentives intervention. Since it occupied a genuine state of equipoise, participants were not subject to a policy reasonably expected to be less effective at realizing the target outcome in question than those outside of the experiment. Participants were either subject to the status quo policy or the experimental healthy incentives intervention, which was not expected to be inferior.⁴¹

This analysis suggests there is an additional condition under which random assignment is permissible:

⁴⁰ Though it is intuitive to think that the health incentives policy would be more effective, government agencies face significant challenges in successfully communicating program changes to recipients. In the context of the HIP, only 62 percent of households receiving the incentive intervention reported having heard of it 4-6 months after they began participating in the study. Bartlett et al. *Evaluation of the Healthy Incentives Pilot (HIP)*, 85.

⁴¹ MacKay, "The Ethics of Public Policy RCTs," 64-65.

SC2: A government agency may randomly assign participants to different policy interventions if:

1. It occupies a state of genuine equipoise regarding all arms of the study and the non-BPA status quo policy; and
2. It does not have the authority to implement the BPA policy.

Random assignment is therefore permissible even if government agencies are not in a state of equipoise regarding the interventions under study and the BPA policy. However, neither SC1 nor SC2 permits agencies to authorize or conduct such a policy RCT where they have the authority to simply implement the BPA policy. If they have this authority, they should exercise it, not conduct or authorize RCTs which evaluate inferior policies.

Suppose governments agencies are not in a state of indeterminacy between the interventions in an RCT, might random assignment still be permissible? I explore this question next.

3 Permissible Randomization 2: Just Allocation of a Scarce Good

Government agencies may have reasons to conduct RCTs even when they are not in a state of equipoise. I argue that random assignment is permissible in such cases when agencies are in a state of indeterminacy regarding *who* should have access to the superior intervention. I first discuss RCTs that evaluate the BPA policy against a superior non-BPA policy. I then turn to RCTs that evaluate the BPA policy against an inferior non-BPA policy.

3.1 Random Assignment to the BPA Policy and a Superior Non-BPA Policy

Government agencies may have reasons to evaluate the BPA policy against a policy that is reasonably expected to be superior in realizing target outcomes, but that is not morally and practically attainable and sustainable. An RCT may help policymakers more precisely determine how much better the policy is and along which dimensions and may also enable them to more accurately estimate the implementation costs. Such evidence may be useful for estimating the cost-effectiveness of the policy and may inform decisions about whether to work to implement it in the future. An RCT may also yield a deeper understanding of the intervention's underlying mechanisms – i.e. *how* it causes superior outcomes.

Consider the Canadian province of Ontario's recent guaranteed income pilot project:

Guaranteed Income: Beginning in 2018, the government of Ontario randomly assigned over 6000 low-income participants in three locations to either an experimental guaranteed income intervention or the status quo policy.⁴² The experimental policy offered recipients a maximal cash transfer of \$16,989/year for individuals and \$24,027/year for couples, less 50% of any earned income.⁴³ The status quo policy offered recipients a maximum cash transfer of \$8,472 for individuals and \$13,140 for couples if their income fell below a certain threshold; and to receive this transfer individuals had to be willing to actively look for employment. Policymakers sought to understand the experimental intervention's impact on participants' ability to meet basic needs, food security, stress and anxiety,

⁴² Government of Ontario, *Ontario Basic Income Pilot*, available at: <https://www.ontario.ca/page/ontario-basic-income-pilot>.

⁴³ Ibid.

mental health, health and healthcare usage, housing stability, education and training, and employment and labor market participation.⁴⁴

The guaranteed income intervention, while undoubtedly superior to the status quo policy with respect to the target outcomes of income security, food security, and housing stability (among others), was not the BPA policy since its affordability was in doubt. A central purpose of the RCT was thus to determine whether the intervention would lower people's labor market participation, and lower demand for healthcare services. Assuming that the status quo policy was the BPA policy, participants were either subject to the BPA policy or a superior non-BPA policy, and so were subject, at minimum, to a policy expected to realize the relevant target outcomes to the extent promised by the BPA policy. Since the guaranteed income intervention is superior to the BPA policy, the central ethical question concerns the allocation of a scarce good: access to the experimental policy.

Following Stone's work on lotteries, random assignment is permissible in these types of cases when policymakers are in a state of indeterminacy regarding who should receive access to the superior non-BPA policy. A state of indeterminacy obtains, Stone argues, if there is a group of people with equally strong claims to the good, and not enough of it to satisfy all the claims.⁴⁵ The use of a fair lottery, Stone argues, realizes the value of impartiality, which demands that agents only allocate scarce goods on the basis of claim strength, not on the basis of irrelevant factors. Where parties have equally strong claims to a scarce good, agents must allocate it on the basis of no reasons at all, and since fair lotteries sanitize the decision-making process from the

⁴⁴ Ibid.

⁴⁵ Ibid, 52-53.

influence of reasons, they offer an impartial way to allocate it.⁴⁶ For this type of policy RCT therefore, random assignment is permissible if participants have equally strong claims to access the non-BPA policy.

One might argue that such RCTs will rarely be permissible since people's claims to scarce goods often vary in strength. For example, it is arguable that the worse off have stronger claims to scarce goods than the better off. However, where the scarce good is access to a non-BPA policy, this is less likely to be true since people do not have claims of justice to superior non-BPA policies. Superior non-BPA policies differ from BPA policies in that they are not morally and practically attainable and sustainable. As such, governments have no duty to implement them. Access to the non-BPA policy is a *discretionary good*, a good the government has no obligation to offer to citizens, and so to which no citizen has a claim right.⁴⁷

The fact that access to the non-BPA policy is a discretionary good does not imply that the government may allocate it in any way it pleases. Michael Blake puts the point nicely:

[A] state which gave new cars to all and only white men would be guilty of racial discrimination even if it had no duty to give cars to anyone. Those who did not receive such cars would quite rightly complain that no valid principle could ground such a discriminatory form of distribution. Their complaint would not be mollified by the response that the state, in giving cars, was providing a benefit it had no obligation to provide to anyone.⁴⁸

⁴⁶ Ibid, 82-83. See also John Broome, "Fairness," *Proceedings of the Aristotelian Society* 91 (1990-1991), 97-98.

⁴⁷ Michael Blake, "Immigration and Political Equality," *San Diego Law Review* 45 (2008), 966.

⁴⁸ Ibid, 970.

Even when allocating discretionary goods, governments still have a duty to treat their citizens as moral equals and so may not allocate the good on the basis of reasons that violate people's moral equality. However, because the good in question is discretionary, governments have more leeway in how they distribute it.

For example, suppose the U.S. government creates a program that offers ordinary people the opportunity to travel to space alongside professional astronauts.⁴⁹ It would be reasonable to hold a lottery among willing and able candidates. The U.S. government has no obligation to allocate the opportunity in a way that creates the most utility or gives priority to the worse off. By contrast, if it is making a decision about how to allocate scarce vaccines in a pandemic, holding a lottery is unlikely to be a just approach.⁵⁰ Where the scarce good is non-discretionary, people often have claims of varying strength and so allocation via lottery will be impermissible; in cases where the scarce good is discretionary, this is less likely to be so.

One might object that the use of random assignment in *Guaranteed Income* did not treat people as moral equals. First, only those who met the eligibility requirements concerning income and geographic location could participate. Second, the study involved the unequal treatment of low-income participants since some were granted access to the non-BPA policy and some were not. Indeed, this line of argument was found persuasive in the case of Finland's basic income experiment. Investigators initially proposed randomly assigning people to different benefit

⁴⁹ This is not a terribly fanciful example. In 1984, President Ronald Reagan introduced the Teacher in Space Project, which granted teachers the opportunity to travel to space as Payload Specialists.

⁵⁰ See Govind Persad, Alan Wertheimer, and Ezekiel J. Emanuel, "Principles for Allocation of Scarce Medical Interventions," *The Lancet* 373 (2009): 423-431.

levels, including levels far greater than that of the status quo policy.⁵¹ However, this proposal was rejected on the grounds that it would violate the Finnish constitution, which requires that all individuals are to be treated equally.⁵² Investigators were thus required to test only one basic income intervention at a level that was roughly equivalent to the existing policy.⁵³

I agree that governments have a duty to treat their citizens equally where this involves not treating them differently on the basis of some morally arbitrary ground. However, this duty does not always forbid governments from using random assignment in cases where the BPA policy is evaluated against a superior non-BPA policy.

I suggest that whether a ground of differential treatment is morally arbitrary depends on two factors: whether the ground is irrelevant to the treating agent's realization of one or more of its purposes; and whether the treating agent's purpose is more valuable than the purposes that would be frustrated by the differential treatment.⁵⁴ The moral arbitrariness of grounds of differential treatment is thus context dependent. Skill may be morally arbitrary when a government agency is deciding how to allocate scarce vaccines in an epidemic; it may not be when that agency is deciding whom to hire.

Although governments have a duty not to treat people differently on the basis of morally arbitrary grounds, differential treatment is thus consistent with this duty if:

⁵¹ Kela, *From Idea to Experiment: Report on Universal Basic Income Experiment in Finland*. 2016. Available from: <https://helda.helsinki.fi/bitstream/handle/10138/167728/WorkingPapers106.pdf>, 12-13.

⁵² Ibid, 60.

⁵³ Ibid.

⁵⁴ Douglas MacKay, "Immigrant Selection, Health Requirements, and Disability Discrimination," *Journal of Ethics and Social Philosophy* 14 (2018), 56.

- A. The differential treatment is expected to significantly advance the realization of one of the government's purposes;
- B. There is no non-differential treatment by which the government may significantly advance the realization of its purpose that would not result in undue burdens on it; and
- C. The purpose of the government is more valuable than the purposes frustrated by the differential treatment.⁵⁵

Conditions (A) to (C) together express the idea that governments may treat people differently on the basis of grounds that are not morally arbitrary. Conditions (A) and (B) specify that the ground of differential treatment must indeed be relevant – i.e. *non-arbitrary* – to the government's realization of one of its purposes. Condition (C) specifies that the ground of differential treatment is *morally* non-arbitrary by ensuring that the purpose in question is more valuable than the purposes that would be frustrated by the differential treatment.⁵⁶

The underlying idea of this account is that differential treatment is morally worrisome when and because it involves disadvantaging people in the pursuit of their purposes.⁵⁷ Hence there is a strong justification for laws and policies that prevent public and private agents from treating people differently – i.e. anti-discrimination laws. However, it also recognizes that prohibitions on differential treatment may prevent public and private agents from pursuing morally important purposes. Conditions (A) to (C) provide a framework for adjudicating conflicts between the purposes of different agents, suggesting that differential treatment is

⁵⁵ Ibid, 57.

⁵⁶ Ibid, 55-56.

⁵⁷ Ibid, 50.

permissible when it is necessary for the realization of some purpose, and this purpose is more valuable than the purposes that would be frustrated by it.⁵⁸

Might RCTs such as *Guaranteed Income* satisfy (A) to (C)? To answer this question, an account of how to weigh and compare the value of different purposes is needed. Developing such an account is beyond the scope of this paper; but the following sketch offers what I take to be a promising approach that is sufficient to motivate the claim that random assignment is sometimes permissible in cases like *Guaranteed Income*.

Consider first that within theories of justice, those purposes that are morally most important are the subject of certain deontic categories.⁵⁹ Agents have *duties of justice* to fulfill purposes whose realization is morally imperative, and people have *claims of justice* that these purposes be fulfilled. Similarly, agents are granted *rights* to protect their ability to pursue purposes that are particularly valuable. One important distinction among purposes therefore is that those that are the subject of duties or claims of justice are more valuable than those that are not. For liberal theories of justice broadly construed, states have duties of justice to ensure residents have access to those goods they require to set, revise, and pursue a conception of the good life; and residents have claims of justice to these goods, including basic rights and liberties, education, health care, opportunities, and income. Residents do not however have claims to those goods necessary to realize particular conceptions of the good life. I have a claim of justice on the state to provide me with healthcare, not to satisfy my preference to travel widely.

⁵⁸ Ibid, 57-58. For further discussion on the need to limit the scope of anti-discrimination law to secure the ability of agents to pursue valuable purposes, see Sophia Moreau, "What is Discrimination?" *Philosophy & Public Affairs* 38 (2010), 160-169.

⁵⁹ MacKay, "Immigrant Selection, Health Requirements, and Disability Discrimination," 55-61.

A second consideration concerns the closeness of the relationship between an agent's purposes and its morally permitted central goals and values. By agents' own lights, purposes that are closer to such goals and values are more valuable than those that are incidental to them. For collective agents – e.g. government agencies, private for-profit corporations, and private civic associations – the value of a particular purpose depends on its closeness to the defining objective of those agents.⁶⁰ Where institutions or organizations are established to realize some morally permissible objective, for example, to practice a particular religious faith in the case of houses of worship, or to provide certain goods or services in the case of for-profit corporations, purposes that are closely tied to this objective should be given greater weight.⁶¹ The same is true of individuals: particular purposes are more valuable the closer they are tied to their conception of the good life and the values they hold dear.

With these considerations in hand, we can show how my account justifies cases of prima facie permissible differential treatment. First, it explains why it is sometimes permissible for American colleges receiving public dollars to favor racial minorities for admission in order to secure the educational benefits of having a diverse student body. Such a policy satisfies (A) since providing a rich educational environment for college students is a core purpose of colleges; and it satisfies (B) when alternative admissions schemes fail to realize this purpose. It also satisfies

⁶⁰ Ibid, 56.

⁶¹ This is reflected in discrimination law where, as Tarunabh Khaitan notes, for an organization's discriminatory treatment to be justifiable it must be necessary to advance some goal that is closely tied to its "*raison d'être*." Tarunabh Khaitan, *A Theory of Discrimination Law* (New York: Oxford University Press, 2015), 181. Similarly, Sophia Moreau notes that most jurisdictions permit agents to justify a discriminatory rule by showing that it is a "Bona Fide Requirement" and so necessary for the "fulfillment of the organization's goals." Moreau, "What is Discrimination?" 164.

(C) since, whereas this purpose is central to the mandate of colleges and is arguably one they have a duty to carry out, no applicant has a claim of justice to attend any particular college.⁶² Similarly, my account explains why it is permissible for religious organizations, including houses of worship and schools, to favor applicants for clerical or educational roles on the basis of religious faith. Such discriminatory treatment satisfies (A) and (B) since such organizations can only successfully practice and promote a particular faith if people performing these roles are committed to that faith. It also satisfies (C) since the practice and promotion of a particular faith is the defining purpose of religious organizations whereas people have no claim to perform a particular job for a particular employer.⁶³

Policy RCTs like *Guaranteed Income* will sometimes satisfy conditions (A) to (C). Recall that the differential treatment in *Guaranteed Income* involved excluding ineligible people from participation and assigning some participants to the superior non-BPA policy. Both forms of differential treatment satisfy condition (A) since they facilitate the carrying out of an

⁶² MacKay, “Immigrant Selection, Health Requirements, and Disability Discrimination,” 60. The U.S. Supreme Court reasons in much the same way, recently re-affirming the judgment that colleges may employ diversity affirmative action policies when they are necessary to fulfill their compelling interest in securing the educational benefits of a diverse student body and the use of race is narrowly tailored to realizing that interest. See *Fisher v. University of Texas at Austin et al.*, 136 S.Ct 2198 (2016).

⁶³ MacKay, “Immigrant Selection, Health Requirements, and Disability Discrimination,” 58. My account therefore provides justification for Title VII of the Civil Rights Act of 1964 which prohibits employers from discriminating on the basis of religion but exempts any “religious corporation, association, educational institution, or society with respect to the employment of individuals of a particular religion to perform work connected with the carrying on by such corporation, association, educational institution, or society of its activities.” Civil Rights Act of 1964, 42 U.S.C. § 2000(e).

affordable RCT aimed at producing socially valuable knowledge: the effect of a guaranteed income on the behaviors of low-income people. These forms of differential treatment are also likely to satisfy (B) since random assignment may be necessary to produce high-quality evidence and other study designs are also likely to involve differential treatment – i.e. rolling out the program in one location rather than another. In addition, geographical limitation is arguably necessary to minimize logistical and financial costs.

Regarding (C), Ontario was aiming to produce knowledge that is relevant to raising the living standards of low-income residents. This purpose is very valuable since Ontario arguably has a duty of justice to realize it. The purposes frustrated by the differential treatment, namely, the interests of non-participants in having a chance to access the guaranteed income intervention, and the interests of participants in the control arm in having access to it, are much less valuable. Because this intervention was not the BPA policy, Ontarians had no claim of justice to it, and so the purposes in question are not as valuable for that reason.

Although I have shown that *Guaranteed Income* satisfies (A) to (C), my account of permissible differential treatment must be developed further if it is to be helpful in the weighing of competing purposes more generally. For example, the considerations discussed above do not address cases where differential treatment is necessary to realize some purpose required by justice but also frustrates some purpose an individual has a claim of justice to carry out. Nonetheless, I expect many policy RCTs comparing a BPA policy against a superior non-BPA policy will satisfy the conditions I have identified. The differential treatment such RCTs involve is often necessary to generate high-quality evidence regarding the superior non-BPA policy, and provided the expected knowledge is valuable, the value of this purpose will usually outweigh the interests of those wanting access to the non-BPA policy since they have no claim of justice to it.

Finally, one might object that my account of permissible differential treatment is too limited, falling into the family of liberty-based theories of discrimination according to which discrimination is objectionable insofar it disadvantages people in the exercise of their freedom on the basis of morally arbitrary features of their identity.⁶⁴ But some suggest that discriminatory actions and policies are wrong because of the message they express. For example, Deborah Hellman argues that discrimination is wrong when and because it *demeans* people, where demeaning action involves “an expression of the unequal humanity of the other.”⁶⁵ One might argue therefore that even if an RCT satisfies the above three conditions of permissible differential treatment, it is still wrong since it demeans those who are ineligible to participate as well as those assigned to the control arm.⁶⁶

It’s difficult to see how such an RCT would do so though. For Hellman, whether an act is demeaning is not dependent on whether its target *feels* demeaned, but rather on whether the act *is* demeaning, where ascertaining this involves complex interpretive judgments regarding social practices.⁶⁷ In *Guaranteed Income*, for example, the decision to limit prospective participants to low-income residents of particular geographic areas was justified by reference to the need to generate knowledge about the intervention’s impact on a particular population, and by the need to limit financial and logistical costs. In no way does the study rest on the judgment that low-income people residing outside the study areas are less worthy of income support. In addition, the use of a fair lottery to decide assignments to the intervention and control arms does not

⁶⁴ See also Moreau, “What is Discrimination?” and Khaitan, *A Theory of Discrimination Law*.

⁶⁵ Deborah Hellman, *When is Discrimination Wrong?* (Cambridge, MA: Harvard University Press, 2008), 35-38.

⁶⁶ Thanks to an anonymous Associate Editor for raising this objection.

⁶⁷ *Ibid*, 75-79.

express the judgment that the humanity of those assigned to the control arm is unequal in some way. Fair lotteries do not treat people differently based on their membership in socially salient groups; and they implicitly recognize the equal standing of those subject to them by granting them an equal chance to receive the good in question.

This analysis implies a third sufficient condition for the permissible use of random assignment:

SC3: Government agencies may randomly assign participants to different policy interventions if:

1. No participant is subject to an intervention reasonably expected to be inferior to the BPA policy;
2. Participants have equally strong claims to be subject to any superior non-BPA policy under evaluation; and
3. Differential treatment of eligible and ineligible people, and participants and non-participants, is practically necessary to produce knowledge relevant for the realization of target outcomes.

Condition (3) summarizes conditions (A) to (C) of the above account of permissible differential treatment. Differential treatment is practically necessary to produce knowledge relevant for the realization of target outcomes when government agencies cannot otherwise do so without incurring undue burdens.

4.2 Random Assignment to the BPA Policy and an Inferior Non-BPA Policy

Government agencies may also have reasons to evaluate the BPA policy against an inferior policy. First, policymakers may begin to implement a new BPA policy but introduce an

element of random assignment in the policy's roll-out to facilitate the collection of high-quality evidence regarding its effectiveness. Such policy RCTs violate the requirement of equipoise since there are sufficiently strong reasons to think that the new intervention – the BPA policy – is superior to the non-BPA status quo policy. However, it is often not possible to roll out the new BPA policy to all eligible people at the same time due to insufficient resources or logistical difficulties. A second type of case arises when the BPA policy has been rolled out but does not cover all eligible people due to resource constraints. In this type of case, access to the BPA policy may already be determined by lottery, thus making random assignment *prima facie* permissible.

In these types of cases, the BPA policy is a “scarce policy,” that is, a policy for which there are insufficient resources – e.g. money, bureaucratic capacity, etc. – to cover all eligible citizens. A scarce BPA policy is thus one that is attainable and sustainable, but only for some portion of eligible people, at least for some period of time. A scarce BPA policy is different from an experimental policy since the former is one that governments have decided to *implement*. RCTs evaluating scarce BPA policies are thus often opportunistic since the policy in question has been or will be implemented outside of the context of an RCT.

Government agencies may have a number of reasons for carrying out RCTs evaluating scarce BPA policies. First, even when there is evidence supporting a particular intervention, an RCT may yield higher-quality evidence regarding its effectiveness, or evidence regarding some particular feature of it. Second, governments may have *political reasons* to conduct an RCT in the context of policy roll-outs. The collection of high-quality evidence regarding the policy's effectiveness may make it more difficult for a subsequent government to eliminate it.

Is it permissible to use random assignment to evaluate scarce BPA policies? I first consider the case of policy roll-outs, working through the example of Mexico's roll-out of *Progresa*, a conditional cash transfer program. I then turn to the case of BPA policies for which there are insufficient resources to cover all eligible people.

Progresa: Beginning in 1998, the Mexican federal government began implementing *Progresa* in 495 poor rural communities, a program offering conditional cash grants to mothers below a specific poverty level with at least one child enrolled in grades 3 through 9.⁶⁸ The grants were conditional on the child attending 85% of school days with the size of the grant varying on the gender and grade level of the child, and inflation.⁶⁹ Partnering with investigators, the federal government first identified the poorest rural communities and then randomly selected 314 of these 495 localities to receive the program from summer 1998 to summer 2000, with the remaining 181 localities serving as controls until fall 2000 when they too would receive the program.⁷⁰ Mothers receiving the cash grant were informed that the program subsidies were assured only for three years – the time until the next election.⁷¹

The challenge for RCTs like this one is that the government would seem to be denying people access to an intervention to which they are entitled. Note too that the promise to expand the program once the RCT was completed does not address this problem since there would still be a two-year period when the people in question would not have access to it. Similarly, if I am

⁶⁸ Schultz, "School Subsidies for the Poor," 202.

⁶⁹ *Ibid.*

⁷⁰ *Ibid.*

⁷¹ *Ibid.*, 206.

entitled to government-provided health insurance, it is unjust to strip me of this benefit for the purposes of conducting an RCT, and the fact that the RCT is time-limited does not eliminate this injustice.

Commentators have argued, however, that it is sometimes permissible to conduct this type of RCT. They note first that if the BPA policy in question is genuinely a scarce policy, it is inevitable that not all citizens will access it. They argue second that random assignment is a fair way to allocate opportunities to be subject to the BPA policy.⁷² In these types of cases, these commentators conclude, random assignment is ideal not only from an epistemic perspective, but also from the perspective of distributive justice.

This line of argument is promising, but it needs further development. First, if such an RCT is to be permissible, the resource constraints must be such that there is a sufficient window of time to conduct a well-designed RCT. That is, there must be some number of people X without access to the intervention for time Y where X and Y are large enough to carry out a well-designed RCT. If one or more people is denied access to an intervention to which they are entitled for some period of time longer than would be dictated by genuine resource constraints, an injustice has occurred.

⁷² For examples of this line of argument, see Diener and Crandall, *Ethics in Social and Behavioral Research*, 136; Abhijit Banerjee and Esther Duflo, "The Experimental Approach to Development Economics," in *Field Experiments and Their Critics: Essays on the Uses and Abuses of Experimentation in the Social Sciences*, Edited by Dawn Langan Teele (New Haven: Yale University Press, 2014), 101; J.M. Gueron, "The Politics and Practice of Social Experiments: Seeds of a Revolution," in *Handbook of Economic Field Experiments*, Volume 1, edited by Abhijit Banerjee and Esther Duflo, (Amsterdam: North-Holland, 2017), 35, 49; and Glennerster and Powers, "Balancing Risk and Benefit," 380-381.

Second, random assignment may not always be a fair way to allocate opportunities to access the BPA policy. For random assignment to be permissible in these types of cases, eligible people must have equally strong claims to it. In *Progresa*, if one community had a stronger claim to the conditional cash transfer program than another, then it was wrong to employ a lottery to determine which community received it. In contrast to *Guaranteed Income*, because the conditional cash transfer program was the BPA policy, low-income Mexicans had a claim to access it.

Third, the randomization procedure must recognize the claims of all people eligible for the BPA policy. If people below some income threshold are eligible for a scarce cash transfer program and have equally strong claims to it, it will not be permissible to allocate this program by means of a fair lottery that only considers some claims and not others. That is, if government agencies wish to use random assignment to evaluate a BPA policy, they may not carve out a portion of the eligible population and conduct a fair lottery among them since doing so fails to recognize the equally strong claims of eligible people outside of this population. The randomization procedure must therefore be consistent with the conducting of a fair lottery among all eligible people.

The above analysis suggests a fourth sufficient condition regarding the permissibility of random assignment:

SC4: If the BPA policy is a scarce policy, government agencies may randomly assign participants to either the BPA policy or some alternative intervention or non-intervention (which may be inferior) if:

1. No participant is subject to an intervention reasonably expected to be inferior to the non-BPA status quo policy;

2. Participants have equally strong claims to be subject to the BPA policy; and
3. No person's claim to the BPA policy is devalued for the purposes of conducting the RCT.

Regarding (3), a person's claim to the BPA policy is devalued either if they are given no chance to receive the BPA policy or if their chance to receive the BPA policy is lesser than is warranted by the strength of their claim.

These conditions are also helpful for determining the permissibility of RCTs where the BPA policy is not temporarily scarce but permanently so. Consider the Moving to Opportunity RCT:

Moving to Opportunity: In 1992, the U.S. Congress authorized the U.S. Department of Housing and Urban Development to carry out the Moving to Opportunity for Fair Housing (MTO) study. Eligible to participate in the study were low-income families with children living in public housing or Section 8 project-based housing located in census tracts with poverty rates of at least 40 percent.⁷³ 4,608 families, in five large U.S. cities, participated in the study. Participating families were randomly assigned to one of the following three groups:

- A. The MTO Low Poverty Voucher Group: Families assigned to this group received a Section 8 voucher usable only in low-poverty areas (10 percent of the population below the poverty line in 1990) as well as counseling and assistance in finding a private unit to rent;
- B. The Traditional Voucher Group: Families assigned to this group received regular Section 8 vouchers – i.e. vouchers with no geographic restriction –

⁷³ Chetty, Hendren, and Katz, "The Effects of Exposure to Better Neighborhoods on Children," 860.

along with ordinary briefings and assistance from local public housing agencies (PHAs);

- C. The Control Group: Families assigned to this group received no vouchers but remained eligible for public or project-based housing and other social programs.⁷⁴

The MTO study had two research goals: comparing the costs of interventions (A) and (B); and determining the impact of intervention (A) on the well-being of families and their children, including “their housing conditions, mental and physical health, employment and earnings, receipt of social program assistance and income, education, and delinquent or risky behavior of children.”⁷⁵ Families were willing to participate in the MTO study because demand for Section 8 housing assistance outstripped supply.

When *Moving to Opportunity* was carried out, the BPA policy was a scarce policy. HUD and local PHAs did not have sufficient funding to provide section 8 housing vouchers to all eligible residents. Given the number of families involved in *Moving to Opportunity* and the size of demand for these vouchers, it is arguable that condition (3) of SC4 is satisfied. Conducting *Moving to Opportunity* did not prevent HUD or local PHAs from implementing the BPA policy in an appropriate way – i.e. *Moving to Opportunity* did not lead to fewer people overall receiving Section 8 housing vouchers. *Moving to Opportunity* also satisfies condition (1); no arm of the RCT was expected to be less effective than the non-BPA status quo policy, in this case no intervention.

⁷⁴ “A Summary Overview of Moving to Opportunity: A Random Assignment Housing Mobility Study in Five U.S. Cities,” available at: <https://www.nber.org/mtopublic/MTO%20Overview%20Summary.pdf>.

⁷⁵ Ibid.

The real question is whether the study satisfies condition (2). Participating families were randomly assigned to one of the above three groups, but some families may have stronger claims than others to the traditional voucher – e.g. those which are worse off. Indeed, local PHAs currently have the authority to deviate from a first-come, first-served allocation system by giving priority to particular types of families, including families who are homeless or living in substandard housing, and families who are paying more than 50% of their income in rent.⁷⁶

My aim here is not to determine whether the *Moving to Opportunity* was ethical or not. Doing so would require a good deal of research and defending a position on distributive justice. Rather, my aim is to show that where the BPA policy is a scarce policy, random assignment may sometimes be impermissible. To satisfy SC4, participants must have equally strong claims to be subject to the BPA policy. Importantly, this need not mean that all who are eligible for the policy must have equally strong claims, only that study participants do so. It may therefore be permissible to conduct a study similar to *Moving to Opportunity* if the eligibility criteria are designed so as to only include people who are similarly situated – i.e. members of the worse off group. Of course, such a study must also satisfy condition (3) and so not involve devaluing the claims of non-participants. The conditions I outline above for permissible random assignment thus also offer lessons for the ethical design of RCTs, not only their evaluation.

Government agencies may also wish to conduct RCTs such as *Progresa* or *Moving to Opportunity* under conditions where people are not subject to the BPA policy. For example, a government agency rolling out policy A, which it knows to be inferior to the BPA policy, may wish to randomly assign participants to receive either A or the non-BPA status quo policy for the

⁷⁶ U.S. Department of Housing and Urban Development, *Housing Choice Vouchers Fact Sheet*, available at: https://www.hud.gov/topics/housing_choice_voucher_program_section_8

purposes of collecting higher quality evidence regarding A's effectiveness and mechanisms.

Might such random assignment be permissible?

I think so. As with *Healthy Incentives Pilot*, as long as the government agency in question lacks the authority to implement the BPA policy and people are not made worse off than they would otherwise be, random assignment is permissible. To cover these types of cases, we can therefore add a fifth sufficient condition 5:

SC5: If citizens are subject to an inferior non-BPA scarce policy, governments agencies may randomly assign participants to either the scarce non-BPA policy or some alternative intervention or non-intervention (which may be inferior) if:

1. The government agency responsible for designing and conducting the RCT does not have the authority to reform the non-BPA status quo policy or otherwise subject participants to the BPA policy;
2. No participant is subject to an intervention reasonably expected to be inferior to the policy to which eligible people would be otherwise subject if they are not subject to the non-BPA scarce policy;
3. Participants have equally strong claims to be subject to the non-BPA scarce policy; and
4. No person's claim to the non-BPA scarce policy is devalued for the purposes of conducting the RCT.

4 Refining the Necessary Condition

Part 1 concluded with a statement of the following necessary condition for the permissible use of random assignment:

NC: Government agencies may authorize or conduct policy experiments in which participants are randomly assigned to different interventions only if such assignment is consistent with their duty to realize target outcomes for participants and non-participants.

Having worked through a number of scenarios in which government agencies conduct RCTs, we can now refine this condition and give some examples of RCTs that violate it. In particular this condition should be refined to more clearly account for cases where government agencies lack the authority to implement the BPA policy, and cases where the BPA policy is a scarce policy. I suggest the following formulation:

NC*: Government agency A may randomly assign participants to different interventions only if no participant or nonparticipant's claim to have their target outcomes realized to the extent promised by the policy A has a duty to implement is devalued.

A person's claim is devalued when the chance of this claim being fulfilled is less than what it ought to be given A's duties of justice.

This formulation accounts for scenarios where government agencies wish to evaluate a promising intervention that is reasonably expected to be inferior to the BPA policy but lacks the authority to implement it. Government agencies only have a duty to implement the most effective policy they have the authority and resources to implement. It also accounts for scenarios where the BPA policy is scarce by allowing that the value of people's claims may vary. Where the BPA policy is not scarce, people have a claim to access it or a policy that is not reasonably expected to be inferior. An RCT that granted them only a 50% chance of accessing such a policy would therefore devalue their claim. Where a BPA policy is scarce, people may only have a claim to have a chance of accessing it, where the strength of their claim may vary, for example, based on need.

What are some types of RCTs that would violate NC*? First, it will be impermissible for government agencies to randomly assign participants to an intervention that is reasonably expected to be inferior to the BPA policy when they have the authority to implement it. This will be true even if participants are ex ante no worse off than they would otherwise be because the government agency in question has failed to implement the BPA policy. For example, a number of U.S. states have chosen not to expand their Medicaid programs under the framework of the Affordable Care Act. Under this legislation, states have the option to expand Medicaid coverage to low-income adults earning less than 138% of the federal poverty level, with the federal government committed to paying 100% of the costs of expansion through 2016, with the matching rate phasing down to 90% in 2020.⁷⁷ Suppose the North Carolina legislature, which has decided not to expand Medicaid, authorizes the North Carolina Department of Health and Human Services to conduct an RCT in which uninsured North Carolinians are randomly assigned to either no coverage, or a bare-bones catastrophic health insurance plan, with the aim of determining whether the latter plan leads to superior health and financial outcomes for recipients. This study would violate NC* since North Carolina has the authority to expand Medicaid, there is good evidence to suggest that having Medicaid leads to better health and financial outcomes than not having Medicaid,⁷⁸ and an expanded Medicaid program is clearly

⁷⁷ Patient Protection and Affordable Care Act, 42 U.S.C. § 18001 (2010). Note that under the initial legislation, Medicaid expansion was required. In June 2012, the U.S. Supreme Court ruled that Medicaid expansion is lawful only as an option for states. *National Federation of Independent Business v. Sebelius*, 567 U.S. 519 (2012).

⁷⁸ Larisa Antonisse, Rachel Garfield, Robin Rudowitz, and Madeline Guth, “The Effects of Medicaid Expansion under the ACA: Updated Findings from a Literature Review,” Kaiser Family Foundation (August 2019), available

attainable and sustainable in North Carolina given its resources and the funding commitments of the federal government.

The same will be true for government agencies that lack the authority to implement the BPA policy, but randomly assign participants to an intervention that is ex ante inferior to the policy they have a duty to implement. For example, it would be impermissible for the USDA's Food and Nutrition Services to randomly assign SNAP recipients to a version of SNAP that is ex ante inferior – but perhaps cheaper – than the current SNAP policy since doing so would devalue recipients' claims to the current policy.

It will also be impermissible to use random assignment where the BPA policy is scarce and an RCT devalues the claims of some participants or non-participants. Although the Section 8 housing voucher program was a scarce policy in *Moving to Opportunity*, it would have been impermissible to randomly assign participants to the three interventions under study if some had a stronger claim than others to the traditional voucher. If households at risk of homelessness have stronger claims to a voucher than others, random assignment is not permissible for it devalues their claims, treating them as equal to those of others. Similarly, RCTs are also impermissible if they devalue the claims of nonparticipants. Had the Mexican government in *Progresa* withheld the conditional cash transfer from more households than necessary in order to conduct a well-designed RCT in a particular region of the country, the Mexican government would have violated NC* by devaluing the claims of nonparticipants who were given no chance to access the program.

at: <http://files.kff.org/attachment/Issue-brief-The-Effects-of-Medicaid-Expansion-under-the-ACA-Findings-from-a-Literature-Review>.

There are no doubt other ways in which policy RCTs may violate NC*. My hope is that this discussion is helpful in illustrating NC*'s implications for the permissible conduct of RCTs.

Conclusion

My aim in this paper has been to specify conditions for the permissible use of random assignment by government agencies. I have defended one necessary condition in addition to five sufficient conditions. I do not claim that this list is exhaustive. For example, one possibility worth exploring is whether, in cases where it would be unfair to allocate access to a scarce BPA policy by means of a fair lottery, it would be permissible to do so by means of a weighted lottery. Some scholars argue that there are circumstances where weighted lotteries are a just way to allocate scarce goods,⁷⁹ and the use of such lotteries has been explored in the context of clinical RCTs.⁸⁰

One might worry that if the conditions I have identified in this paper are exhaustive (or nearly exhaustive), my account will constrain the development of valuable knowledge. For example, as I note above, it is possible that *Moving to Opportunity* was unethical and so should

⁷⁹ See Frances Myran Kamm, "Equal Treatment and Equal Chances," *Philosophy and Public Affairs* 14 (1985): 177-194; and Ben Saunders, "A Defence of Weighted Lotteries in Life Saving Cases," *Ethical Theory and Moral Practice* 12 (2009): 279-290.

⁸⁰ See Andrew L. Avins, "Can Unequal Be More Fair? Ethics, Subject Allocation, and Randomised Clinical Trials," *Journal of Medical Ethics* 24 1998: 401-408; Spencer Phillips Hey and Jonathan Kimmelman, "The Questionable Use of Unequal Allocation in Confirmatory Trials," *Neurology* 82 (2014): 77-79; and Alexander A. Iyer, Saskia Hendriks, and Annette Rid, "Advantages of Using Lotteries to Select Participants for High-Demand COVID-19 Treatment Trials," *Ethics & Human Research*, Forthcoming.

not have been carried out as designed. Such a conclusion would be regrettable in one respect since this study yielded highly valuable findings regarding the impact of neighborhoods on children's life prospects.

My first response to this concern is "tough." One of the great moral advances of the 20th century was the recognition that participants in clinical research must be treated with respect and fairness, and that these constraints may not be set aside simply because a study seems promising. So too, participants in policy research may not be treated unjustly simply because doing so is necessary for the development of valuable knowledge. There may be some things that it would be valuable to know, for example, because such facts would greatly aid in the development of effective public policies, but that we may not come to know because of the moral constraints regarding the treatment of people.

That said, there may be cases where the unfairness is minor enough, and the potential benefits of a policy RCT great enough, that the value of these potential benefits could outweigh the pro tanto wrongness of the unfair treatment. After all, it is sometimes on balance permissible to infringe people's rights or treat people unfairly when doing so is necessary to promote a greater good or to satisfy people's weightier claims to justice. That such cases are possible does not mean they are common. Moreover, given that the vast majority of policy RCTs enroll people who are disadvantaged along a number of dimensions, including political voice, it is crucial that requirements of fair treatment not be set aside without the most rigorous justification.