RCTs in Development Economics: An Ethical Audit

By

Marcos R. Picchio

A dissertation submitted in partial fulfillment of
the requirements for the degree of

Doctor of Philosophy
(Philosophy)

at the
UNIVERSITY OF WISCONSIN-MADISON
2023

Date of final oral examination: 03/24/2023

The dissertation is approved by the following members of the Final Oral Committee:
Harry Brighouse, Professor, Philosophy
Daniel Hausman, Emeritus Professor, Philosophy
J. Paul Kelleher, Associate Professor, Philosophy / Medical History and Bioethics
James Goodrich, Assistant Professor, Philosophy
Priya Mukherjee, Assistant Professor, Agricultural and Applied Economics
Norman Fost, Emeritus Professor, Medical History and Bioethics

Table of Contents

Acknowledgments

I would like to first and foremost thank Priya Mukherjee and Paul Kelleher. I owe my interest in the ethics of RCTs in development economics to Priya and Paul. In the Fall of 2020, Priya was looking for a research assistant to help with a project on the ethics of development RCTs. Knowing that I had a general interest in philosophical issues in economic methodology, Paul kindly forwarded my contact information to Priya, who subsequently contacted me and offered me the position. Priya provided me with literature to review, and it was while doing this initial literature review that I realized that the already sparce literature on the ethics of development RCTs lacked contributions from philosophers. I was quickly captivated by RCTs in development economics and decided to abandon my original dissertation topic to focus on an opportunity to use my philosophical training to comment on policy-relevant and cutting-edge work in economics. But since there virtually no philosophical literature on the ethics of development RCTs, I did not know where to start to my dissertation research. Paul was incredibly helpful in getting me oriented with the literature on the ethics of medical research, which I draw on considerably in this dissertation and now have independent research interests in.

I would also like to thank my co-advisors: Harry Brighouse and Dan Hausman. Had it not been for their encouragement, I would not have felt comfortable writing on a topic which so little had been said about. I have also learned a considerable amount from them during my time at UW-Madison. My thinking on the ethics of development RCTs is influenced by both their philosophical interests and approach to philosophy more generally.

Doug MacKay and Norm Fost also deserve special mention. Doug's pioneering work on the ethics of policy experiments was a source of philosophical inspiration. My numerous conversations with Norm also helped shape my approach to research ethics. Both provided me with detailed written comments on early drafts of multiple chapters.

I would lastly like to thank my family and friends. My parents, Matilde Picchio and Gaston Picchio, and my two sisters, Camila Picchio and Delfina Picchio, have always encouraged me to pursue my academic interests. Completing a PhD would not have been possible without their support. I would also like to acknowledge Megan Fritts and Frank Cabrera, who, as fellow graduates of UW-Madison's Philosophy PhD program, helped guide me through the program in countless ways.

Abstract

The use of randomized controlled trials (RCTs) in development economics has gained considerable prominence over the past 20 years. Despite this, very little has been said about the ethical permissibility of conducting these kinds of experiments. This is surprising, given that the use of RCTs in medical research is subject to intense ethical scrutiny by physicians and bioethicists alike. In my dissertation, I identify and analyze ethical issues with the use of RCTs in development economics. To do so, I draw on concepts from medical research ethics and connect them to contemporary debates in political philosophy. A major theme that runs throughout my dissertation is that various aspects of the experimental approach to development economics are ethically justifiable, despite initial appearances to the contrary.

In chapter 1 ("**RCTs in Development Economics**"), I provide background on development RCTs and the history of policy experiments, more generally. I identify and discuss three reasons why development economists became interested in the experimental method: (1) the credibility revolution in applied economics; (2) disillusionment with macroeconomics; (3) ethical concerns about foreign aid effectiveness. I close the chapter with discussion of some influential development RCTs.

In chapter 2 ("**What Is the Standard of Care in Experimental Development Economics?**"), I analyze the routine practice of using the status quo as a control. In medical research, participants in all trial arms are entitled to a *standard of care*, i.e., a pre-specified level of medical resources and attention. Placing research participants below this standard is wrong, which raises a serious worry when one considers that the status quo in developing countries is unjust. In part one I address the question: What level of socioeconomic resources and opportunities are owed to research participants by development economists? At its core, this is a

question about distributive justice, and in particular, distributive justice in non-ideal circumstances. Drawing on lessons from controversies with the use of placebo controls in international medical research, I argue that the "standard of care" in experimental development economics is the status quo. However, to address the concern that my account is overly permissive, I also argue that development economists conducting experiments are bound by the natural duty of rescue. This, in turn, provides the basis for a *minimum* standard of care.

In chapter 3 ("**What Role Should Equipoise Play in Experimental Development Economics?**"), I analyze the role that the concept of *equipoise* should play in experimental development economics. In its most general guise, equipoise is a state of uncertainty about the relative merits of a set of interventions. The most famous version of equipoise is *clinical equipoise*, which obtains when the expert medical community is uncertain about the relative therapeutic benefits of a set of interventions. Many physicians and bioethicists believe it is imperative to establish clinical equipoise before a trial goes forward, because without it, researchers would violate the fiduciary obligations they have as physicians. Yet development economists routinely fail to establish clinical equipoise, since the interventions they evaluate are known to provide welfare benefits to individual recipients. I argue that equipoise has only a minor role to play in experimental development economics. I allow that it has some importance, because it provides evidence that an experiment is socially valuable and worth carrying out. I proceed to develop the concept of *social equipoise*, which requires that (1) there not only be uncertainty about the research questions an experiment answers, but also that (2) the experiment be designed in a way that reduces the degree of uncertainty in question. I use the latter condition to address some of the methodological disputes with the use of RCTs in development economics.

In chapter 4 (**"Who Has the Right to Cluster Randomize?")**, I analyze the use of an experimental technique known as "cluster randomization." Instead of randomizing individuals into treatment and control groups, development economists instead randomize social entities, e.g., villages or neighborhoods. While cluster randomization has methodological advantages, it often makes obtaining individual consent impractical. Cluster randomization raises a serious challenge since individual consent is widely regarded as a necessary condition of permissible experimentation. I link the issue of research without consent to questions about political legitimacy and maintain that governments can legitimately exercise their authority over citizens even if individuals do not consent to specific government actions. I leverage this argument to show that *governments* have the right to conduct research without consent. However, development RCTs are often sponsored and conducted not by governments but by private actors, such as philanthropic foundations and non-profit research firms. This leads me to propose what I call the *authorization condition*, which requires that private actors obtain government authorization to conduct research without consent.

In chapter 5 ("**Policy Experiments, Philanthropy, and Public Goods**"), I address a political worry with how development RCTs are funded. I explain how wealthy philanthropists influence public policy through funding policy research—including specifically policy experiments such as development RCTs. I then suggest that the way policy experiments are funded sits in tension with expectations of political equality found in democratic societies, and therefore calls for closer examination. I argue that governments may permissibly authorize philanthropic foundations to conduct or sponsor policy experiments because these private actors are in a unique position to supply a public good, namely, policy innovation. I end chapter 5 with discussion of avenues for future research.

Chapter 1: RCTs in Development Economics

1. Introduction

In 2019, Abhijit Banerjee, Esther Duflo, and Michael Kremer were awarded the Nobel Prize in Economic Sciences for their work on development economics—the branch of economics which studies the development process in low-income countries. In the scientific background for the award, the Nobel Prize committee credits the laureates for turning development economics into a "blossoming, experimental field" (Committee for the Prize in Economic Sciences, 2019, p. 2). At the heart of this new approach to development economics is the use of field experimentation involving random assignment of research participants to either treatment or control groups. This experimental method is known as a randomized controlled trial (RCT). Though the methodological benefits of RCTs have been well-known by medical researchers for some time, the 2019 laureates were among the first to apply this experimental method to the study of economic development. In the eyes of the Nobel Prize Committee, the laureates' use of this methodological innovation has "dramatically increased the practical quantitative knowledge necessary to isolate key mechanisms behind poverty and behavioral responses to various policy interventions" (Committee for the Prize in Economic Science, 2019, p. 5). To put it succinctly, Nobel Prize Committee credits the laureates with deepening our understanding of poverty in the developing world.

Reading the Nobel Prize Committee's scientific background, one would get the impression that the ends undeniably justify the means. This attitude is reflected in Banerjee and Duflo's (2011) best-selling book, *Poor Economics: A Radical Rethinking of the Way to Fight Global Poverty*, which chronicles the discoveries and subsequent policy interventions made

possible using RCTs.[1] Nowhere in the text do the authors stop and question the ethical permissibility of randomly distributing a scarce resource to a poverty-stricken group of people for the sake of potentially learning something valuable.[2] Unlike medical researchers, economists have said little about the ethical implications of using RCTs in their discipline.[3] Economists can perhaps be excused. Research ethics is not a topic that economists have had to concern themselves with in the past.[4] Economists also lack the specialized training to properly identify and analyze the ethical issues posed by their new experimental methodology. This is a task for philosophy, and this dissertation is an attempt to advance the current state of ethical discourse surrounding development RCTs.

Though there is little research on the ethics of development RCTs, it is a mistake to suggest that there is no philosophical literature bearing on this topic. On the contrary, the philosophical literature on bioethics—specifically the literature on medical research ethics— is sophisticated and rich. In this dissertation, I draw on concepts and principles from medical research ethics to investigate the ethics of development RCTs. However, it is also a mistake to think that concepts and principles from medical research ethics can be mechanically imported to development RCTs. Development RCTs raise questions that are not in the purview of traditional bioethics, e.g., questions about distributive justice, social value, and political legitimacy. As this

[1] See Karlan and Appel (2011) for a similar introduction to development RCTs. Also see Duflo et al. (2004) for an early survey.
[2] Banerjee and Duflo (2009, p. 101) do mention ethics in a programmatic article on the experimental approach to development economics. However, they do not do themselves any favors when they suggest that randomization is seen as fair because poor people are already used to arbitrariness in their daily lives.
[3] There are some exceptions. Contributions from economists include Abramowicz and Szafarz (2020), Alderman et al. (2016), Glennerster and Powers (2016), and Ziliak and Teather-Posadas (2016). Also see Asiedu et al.'s (2021) discussion of an ethics appendix for social science articles involving field experimentation.
[4] However, DeMartino (2011) argues otherwise and calls for an ethical code of conduct for all economists—not just those who conduct experiments. Also see DeMartino and McClosekey (2016) for a recent anthology on professional ethics in economics.

dissertation also reveals, contemporary debates in political philosophy are also important to the ethical analysis of development RCTs.[5]

By taking insights from bioethics and political philosophy, I conclude in this dissertation that—despite initial appearances to the contrary—various aspects of experimental development economics are ethically justifiable. Though I will be defending the experimental approach to development economics, this should not be taken to suggest that there have not been development RCTs that are unethical. My motivation for writing this dissertation is to defend development economists from general ethical objections that do not withstand scrutiny. I am ultimately interested in the ethical *foundations* of development RCTs and will have little to say about practical measures for ensuring that development RCTs are carried out in an ethical manner. While philosophers should certainly weigh in on the formulation of regulatory guidelines for ensuring development RCTs are ethical, their skillset is not tailored for this task. Collaboration is key, but so is respecting a division of labor.

The rest of this chapter serves as an introduction to RCTs in development economics. In section 2, I sketch the conceptual basis for using RCTs. In sections 2 and 3, I provide a quick overview of the history of experiments in economics and their relation to scientific experimentation, more generally. The main contribution of this chapter is section 4. In this section, I identify three factors that led to the ascent of experimental methodology in development economics. Section 5 closes with discussion of some noteworthy development RCTs.

[5] Contemporary political philosophy is also relevant to many controversies in bioethics. See Millum and Emmanuel (2012) for an anthology devoted to the topic.

2. Why Randomize?

In this section, I sketch the conceptual justification for randomization. This section largely draws on Deaton and Cartwright (2018) and Angrist and Pischke (2014, ch. 1).

We begin with a research question. Suppose we are interested in the effects of health insurance (treatment) on some measure of health (outcome) in a population of interest. An RCT provides researchers with *one* way of measuring the effect insurance has on health in a target population. The first step is to create a trial sample. The trial sample depends on the unit of analysis. The unit might be individuals or groups of peoples (households, neighborhoods, etc.).[6] In our health insurance example, individuals are the unit of analysis. A trial sample will contain individuals (or whatever the unit of analysis is) from a population of interest who are then randomly assigned to treatment or control groups. Random assignment is sometimes confused with selecting individuals from a population of interest randomly (random sampling). To conduct an RCT, the trial sample can be comprised of individuals from a population of interest who are (non-randomly) willing to participate in an experiment, or simply individuals who are available to researchers conducting a trial. While random sampling has additional methodological benefits, it is not essential to an RCT, and I set discussion of sampling techniques aside for the purposes of this overview.

Given random assignment into treatment and control arms, there will be data from two frequency distributions, $F(Y_1)$ and $F(Y_0)$, associated with two outcomes, Y_1 and Y_0, which correspond to the randomly assigned treatment and control groups, respectively. In our health insurance example, Y_1 can be the outcome associated with insurance (treatment) and Y_0 can be the outcome that is associated with no insurance (control). The means from these two frequency

[6] In chapter 4, I analyze ethical difficulties with randomizing groups of people rather than individuals.

distributions can be subtracted from one another to obtain an average treatment effect (ATE). ATE estimates have practical applications across medicine and social science. Though RCTs can be used for other purposes, for many the main attraction of an RCT is that it can be used to reliably estimate the ATE of some intervention.

To correctly interpret the ATE, and more generally understand the conceptual justification behind randomization, we can represent the outcome of each unit in either the treatment or control group with the following model:

$$Y_i = \beta_i T_i + \sum_{j=1}^{J} y_j x_{ij} \qquad (1)$$

where

- Y_i is the outcome for unit *i.*
- β_i is the individual treatment effect on *i.*
- T_i is a dichotomous (1,0) treatment "dummy" indicating whether *i* is treated.
- x_{ij} are $j=1,\ldots,J$ observable and unobservable, unit-specific factors that may affect outcome Y_i (note: *J* can be very large).
- y_j is the effect of x_j on Y_i and may be negative or positive.[7]

In our health insurance example, *T* stands for insurance status. The factors other than *T* (the *x'*s) on the righthand side of (1) are *covariates*. Potential covariates include factors such as income, education, employment status, age, race, and sex. But there may also be unobservable and unknown covariates among the *x'*s. As we see below, the chief theoretical virtue of an RCT is that it allows a researcher to control for observable *and* unobservable covariates.

[7] Deaton and Cartwright (2018) omit the *i* subscript on y_j: "we do not need *i* subscripts on the γ's that control the effects of the other causes; if their effects differ across individuals, we include the interactions of individual characteristics with the original *x*'s as new *x'*s" (Deaton and Cartwright, 2018, p. 4).

Equation (1) can be given a causal interpretation using the Potential Outcomes Framework, otherwise known as the Rubin Causal Model (Rubin, 1974). Suppose *T* is dichotomous such as in our health insurance example; either an individual *i* has insurance or she does not. For each *i* there are two possible outcomes then: Y_{i1} and Y_{i0}. The former occurs if individual *i* has insurance and latter if they do not. Using (1), we can subtract the difference between the two outcomes ($Y_{i1} - Y_{i0}$) to define the individual treatment effect of insurance on health. The problem is that in the real world no individual can both be treated and not treated at the same time. Only one of the outcomes occurs—the other is an unobservable counterfactual. If we could somehow access the outcome not realized, we would be able to estimate the difference insurance status makes on an individual's health.

Though we never observe individual treatment effects, with a sufficiently large *N,* the law of large numbers allows access to "the path not taken" for a population but not for each individual that comprises it. Using the Potential Outcomes Framework, we can interpret the mean in the treatment group as a counterfactual for the mean in the control group. Consider an experimental set up with a treatment group (*T=1*) and control groups (*T=0*). If we subtract their means, we get:

$$\overline{Y_1} - \overline{Y_0} = \overline{\beta_1} + \sum_{j=1}^{J} y_j \left(\left(\overline{x}_{1ij} - \overline{x}_{0ij} \right) \right) = \overline{\beta_1} + \left(\left(\overline{S_1} - \overline{S_0} \right) \right) \quad (2)$$

The term of interest on the far-right hand side is β_1, or the ATE of the trial sample. The second term is the "error term": the sum of the net average balance of other causes (the covariates) across the two groups. Some commentators also refer to this term as "selection bias". In the best-case scenario for any research design, the difference between $\overline{S_1} - \overline{S_0}$ is zero. In such a case there is *perfect balance,* which allows for one to cleanly identify β_1 as the average casual effect of the treatment in the trial sample.

What randomization allows us to do is achieve perfect balance *in expectation*. This means that the term on the far right will average out to zero after a hypothetically infinite number trial repetitions. We do not need to know anything about the covariates—whether observable or unobservable—for this to hold. Repeated randomization will result in identical distributions of covariates and potential outcomes in treatment and control groups. So, in expectation, β_1 is the true ATE in the trial sample. But it is important to stress is that in a one-off RCT, the error term will not automatically be zero—there is no guarantee of perfect balance. Calculating the differences in means between treatment and control group means will get you the ATE *plus* the error-term, which is unlikely to be zero. Though it is not possible to run an infinite number of trials, there are ways of addressing this limitation. Most notably, larger sample sizes in one-off RCTs will reduce the size of the error term. It is also possible to identify important observable covariates and ensure the treatment and control groups appear balanced before administering a treatment. With an ATE estimate, one can then test the null hypothesis $\beta_1 = 0$ using the appropriate statistical test and determine the statistical significance of the finding.

An RCT is not the only way to achieve balance. With enough background knowledge of relevant covariates, a carefully controlled laboratory experiment can reliably estimate an ATE without randomly assigning a sample into treatment and control groups. Another option is *matching* (Angrist and Pischke, 2014, ch. 2). With matching, researchers find individuals that closely resemble other individuals in all causally relevant ways *except* treatment status. Matching, if done right, allows researchers to achieve balance and identify ATEs without randomization. In our health insurance example, researchers could try to match uninsured individuals with insured individuals who otherwise share similar known characteristics. But finding the relevant individuals could prove challenging. It is for this reason that virtues of

randomization should not be undersold, and this is especially true in the case of development economics. The kinds of interventions developing economists are interested in cannot be studied in a laboratory. And for many interventions, there are covariates that are unknown or unobservable and therefore difficult—if not impossible—to control for without doing an RCT.

3. The Experimental Method

Historically, economics has not been regarded as an experimental science. Vernon Smith first began testing economic theory with experiments in the 1950s (Smith, 1962), but it was not until 2002 that his methodological contributions were recognized with a Nobel Prize (Kahneman and Smith, 2002). Experimental economics is now a vast subfield of economics, and many economists learn experimental design as part of their professional training. Development RCTs only represent a subset of the experiments conducted by economists.

The typical demarcation in experimental economics is between laboratory experiments and field experiments. Up until twenty years ago, laboratory experiments were the most common kind of experiments in economics. Laboratory experiments take place in carefully controlled conditions (such as university computer lab) and make use of a convenience sample, e.g., university students. Smith's experiments were laboratory experiments, and so were Daniel Kahneman's—Smith's co-laureate for the 2002 Nobel Prize—influential experiments on human judgment and decision-making (e.g., Kahneman and Tversky, 1979). By contrast, field experiments take place in "real-world" settings rather than the carefully controlled environment that a university provides.[8] What distinguishes development RCTs from previous experimental work in economics is that development RCTs take place in the field rather than a laboratory.

[8] This is a rough characterization, and Harrison and List (2004) provide a sophisticated taxonomy of the different kinds of field experiments in economics.

Throughout this dissertation, I will have little else to say about the kind of experimental work that figures such as Smith and Kahneman inspired.

Economists did not pioneer field experimentation and field experiments also do not necessarily involve randomization. However, there is a close historical link between field experimentation and randomization. List and Levitt (2009) identify three waves of field experiments. The first wave is associated with agricultural research in the 1920s and 1930s. This research was associated with agricultural productivity and involved estimating the effects of treatments such as fertilizers on randomly assigned plots of land. The most notable figure from this era is Ronald Fisher, who is often credited as the father of the RCT. Fisher (1926) provided the conceptual justification for randomization (discussed in section 2), and Fisher's (1935) *The Design of Experiments* laid the foundation for modern experimental design.

Though many aspects of modern experimental design can be traced back to this first wave of field experimentation, it is important to note that the RCT did not originate in the context of field experimentation. Hacking (1988) discusses how randomized experiments were first developed by psychologists in the 19th century. Notably, it was C.S. Peirce and his student who designed the first randomized experiment to test the leading psychophysical theory of the day (Peirce and Jastrow, 1885). Subsequent randomized experiments in this period were used to debunk the claims of telepaths and psychics, and Hacking (1988) traces the scientific popularity of randomized experiments to these episodes.

The second wave of field experimentation was the era of large-scale government-sponsored experiments, which began in the 1960s. Within economics, these experiments are referred to as "social experiments". Unlike experiments on agricultural plots, social experiments involve real flesh and blood human beings. And unlike development RCTs, these experiments

originally took place in high-income countries—most notably the United States. Before turning to the current third wave of field experiments, it is worth examining this second wave in some detail. Many of the earliest social experiments from this era were conducted by economists or with their input. Yet this wave of field experimentation did not pique the interest of the economics profession and was even criticized by some leading economists of the time. A notable critic was eventual Nobel-prize winning econometrician James Heckman, who published an influential critique of the use RCTs for social policy evaluation (Heckman, 1992; also see Heckman and Smith, 1995). As I suggest further on, it would take a whole new generation of economists facing different disciplinary pressures to make field experimentation part and parcel of economic methodology.

4. Social Experiments

Greenberg and Shroder (2004) define a social experiment as having at least four of the following features: (i) random assignment; (ii) policy intervention, (iii) follow-up data collection, and (iv) evaluation. But it would be misleading to suggest that the only motivation behind the experiments from this era was to evaluate policy interventions and thereby speak directly with policymakers. The earliest experiments from this era were also motivated by an interest in testing economic theory—an interest that faded quickly.[9] Social experiments from this second wave were conducted by state or federal government agencies and carried out in collaboration with non-profit research firms. Policy interventions studied during this period include negative income taxes, employment programs, health insurance, electricity pricing, and housing allowances. A detailed overview and defense of this wave of field experimentation is Judith

[9] Gueron and Rolston (2013, p. 12) discuss how "social experiments in the welfare area moved out of the researcher-controlled environment (testing behavioral responses with statistical designs driven largely by economic theory) into the more complex context of mainstream public agencies and became the federal standard for the valuation of real-world programs."

Gueron and Howard Rolston' (2013) *Fighting for Better Evidence*. In lieu of a full overview, I discuss three notable experiments.

4.1. The New Jersey Income Maintenance Experiment

In the U.S., debates around the welfare system provided the impetus for social experimentation. In the early 1960s, Milton Friedman's (1962) *Capitalism and Freedom* introduced the idea of a negative income tax (NIT) to a wide audience. The basic idea behind a NIT is that it would pay out cash (negative taxes) to families at the low end of the income scale. Friedman promoted NIT as a way of (i) strengthening the market economy and empowering poor people; and (ii) simplifying welfare payments and reducing government expenditure on welfare programs. Economists, policymakers, and the public were skeptical such a reform could promote desirable consequences—wouldn't a NIT just disincentivize work and create dependency? Heather Ross, then an economics Ph.D. student at the Massachusetts Institute of Technology (MIT) and fellow at the Brookings Institution, took note and proposed a novel way of advancing the policy debate: a large-scale randomized experiment (List and Levitt, 2009). Ross's proposal caught the attention of the U.S. federal government. The resulting experiment was sponsored by the Office of Economic Opportunity (OEO) and conducted by the Institute for Research on Poverty at the University of Wisconsin-Madison (where Ross was employed) and Mathematica, Inc., a non-profit research firm based in Princeton, NJ.

The New Jersey Income Maintenance Experiment was the first large-scale social experiment in the U.S. The experiment began in 1968 in five urban communities in New Jersey and Pennsylvania: Trenton, Paterson, Passaic, and Jersey City in NJ, and Scranton, PA. 1,300 low-income families enrolled in the experiment and were then randomized into treatment and control arms. The primary outcome of interest was the effect of a NIT on labor supply. The

treatment received cash payments corresponding to an experimental NIT, whereas the control received no NIT payments. The results were first analyzed in Ross's dissertation, but were also published by the head of Mathematica, Inc. in *Scientific American* (Kershaw, 1972). Early analyses of the results showed that the NIT did not reduce labor supply in the treatment group. But there were issues with the scientific validity of the experiment. Most notably, the experiment suffered from attrition bias—the tendency for research participants to drop out of a study. This is a problem that many social experiments from this era encountered.[10]

Kershaw (1972) was aware that the New Jersey Income Maintenance Experiment had serious limitations. As he suggests in the *Scientific American* article, that real contribution of the experiment was that it would (and did) inspire more social experimentation (see Gueron and Rolston, 2013). Additional NIT experiments were subsequently conducted in U.S. and Canada, but the results are difficult to disambiguate (see Widerquist, 2005). What is notable is that the economics profession was not particularly inspired by calls for more experimentation. Here was an opportunity to test economic theory or learn something about important economic parameters, e.g., labor supply elasticity. Yet economists were not rushing out to learn experimental design.

Worries about scientific validity is one possible explanation for why the economics profession was unenthusiastic about field experimentation. But as we see below, there was at least one social experiment conducted during this era that proved experiments could be rigorously designed and carried out. A more innocuous explanation was the prohibitive cost of conducting social experiments. By conservative estimates, the New Jersey Income Maintenance

[10] More careful analysis by Moffit (1981) was less optimistic and suggested that labor supply did in fact decrease among treated participants—thereby confirming the suspicions of the skeptics.

Experiment cost around $5 million—$40 million in today's dollars.[11] However, as I suggest further on, the more likely explanation is that economists did not face the disciplinary pressures that I argue gave rise to the third wave of field experimentation.

4.2.The RAND Health Insurance Experiment

When it comes to the second wave of field experimentation (the era of social experiments), the experiment that is most often hailed as an exemplar is the RAND Health Insurance Experiment. The results of this experiment were first published in the prestigious *New England Journal of Medicine* (Brook et al., 1983). This experiment was not without its problems (including attrition bias), but today the RAND Health Insurance Experiment is regarded as one of the best social scientific experiments ever conducted (Aron-Dine et al., 2013). The RAND Health Insurance Experiment was conducted by the RAND Corporation, a private non-profit research firm. The RAND Health Insurance Experiment was also sponsored by the U.S. Department of Health and Human Services. Much of the credit for the design and execution of the experiment goes to economist Joseph Newhouse, whose book *Free for All? Lessons from RAND Health Insurance Experiment* (Newhouse, 1993) provides a detailed methodological overview of the experiment and discussion of its policy implications.

The RAND experiment sought to answer a policy question which (to this day) is mired in controversy, and which the U.S. government has a considerable stake in: Does free health care improve health outcomes? The experiment took place from 1971 to 1982. To conduct the experiment, RAND researchers used a cluster randomized experiment design involving 3,958 people between the ages of 14 and 61 who belonged to 2,005 different families. The participants

[11] As List and Levitt (2009) note, this make's Ross's dissertation one of most expensive in the history of economics. To contrast, funding for this dissertation has totaled a mere $31,367.78 (which was taxed) obtained through three fellowships (two summer terms, one spring term).

lived in various regions of the U.S., including Washington, Ohio, Massachusetts, and South Carolina. The families were randomly assigned in clusters to different experimental insurance plans, which fell into four general categories: free-care plans, individual deductible plans, intermediate plans, and catastrophic plans.[12] The plans variated cost sharing rates and deductibles but none of the participants were charged a premium for the plan they were assigned to. As part of the experiment, RAND had to create a private insurance company to administer the plans.

The headline result from the RAND experiment was that free healthcare does not improve the health outcomes of average participants and leads to wasteful overuse of medical care. However, free care did improve the health outcomes of some low-income subgroups. At the time, RAND researchers concluded that "although free care did not improve health status across the entire range of measures or incomes groups examined, it did confer demonstrable benefits for patients with selected conditions that physicians are trained to manage" (Brook et al., 1983, p. 1433).

The findings from the RAND experiment are not just an intellectual curiosity to those interested in policy debates. There is a case to be made that the findings from this experiment affect U.S. citizens to this day. Though it is hard to establish any definite causal link between the RAND experiment's findings and continual use (and increase) in cost-sharing for medical care in the U.S.,[13] one notable effect of the experiment has been documented. Newhouse and Normand (2017, p. 2161) report that, to this day, the Congressional Budget Office uses the findings from

[12] See Brook et al. (1984, ch. 2) for specific details of the various plans offered.

[13] Newhouse (2004, p. 108) writes that it is not hard to draw a "substantive link" between the results of the RAND experiment and the subsequent rise in health insurance plans with substantial cost-sharing through high deductibles. Yet Newhouse is not suggesting a suggesting a causal link. Elsewhere he suggests such links between RAND experiment findings are speculative. For example, Newhouse (2004, p. 109) reports of an initial increase in cost sharing for inpatient services leading to decreased hospital admissions shortly after the findings were made public but is cautious in attributing causality. Perhaps paradoxically, the cost of healthcare in the U.S. has not decreased despite the increase in cost-sharing and the RAND experiment's findings.

the RAND experiment to estimate the value of cost-sharing when providing value estimates of proposed healthcare-related legislation. This shows the potential power of experiments to alter policy long after their conclusion—a political issue I explore in the final chapter of this dissertation.

Though few development RCTs have been conducted at such a large scale as the RAND experiment, evidence from development RCTs has altered the policies of governments, global institutions, and non-government organizations (NGOs) in ways that have been documented. In section 6, I review some of the most influential development RCTs.

4.3.PROGRESA

Before turning to the rise of development RCTs, it is helpful to review a social experiment that can be seen as a bridge between the second and third waves of field experimentation: The PROGRESA Experiment. This experiment serves as a salient place to end the discussion of social experiments and transition into development RCTs because it was a large-scale social experiment that took place in Mexico—a middle-income country with considerably more poverty than the U.S. In addition, PROGRESA had an influence on the kinds of interventions and research questions studied using development RCTs.

PROGRESA (later re-named “Oportunidades”; now known as “Prospera”) was a government social assistance program in Mexico started in 1997. The program was designed to address extreme poverty and reduce the cycle of intergenerational poverty transmission in Mexico. Today, PROGRESA is recognized as one of the most successful anti-poverty programs ever implemented by a government. What made PROGRESA unique was that cash transfers to recipients were conditional on poor recipients adopting behaviors meant to improve health, nutrition, and education. These behaviors include prenatal care, well-baby care and

immunization, nutrition monitoring and supplementation, preventive checkups, and participation in educational programs regarding health, hygiene, and nutrition. An additional cash transfer was also given to households with school-aged children who regularly attended school. The cash transfers were significant: roughly one-third of household income. Another notable feature of the program was that the cash transfers were given to mothers because it was theorized that this would direct resources to children more effectively—an aspect of conditional cash transfer programs that draws criticism from some feminists.

One of the main architects of PROGRESA was Santiago Levy, then Mexican Minister of Finance, and a US-trained economist (Levy, 2006). Due to budgetary constraints, the Mexican government was unable to rollout PROGRESA to all eligible recipients from the start. This meant recipients would instead need to be phased in over time. Due to logistical constraints, PROGRESA could also only be implemented at the village-level. Well-aware of these constraints, Levy saw an important opportunity to evaluate the effects of PROGREA as an RCT, which he believed would help the program survive long-term. At the time, 50,000 communities were eligible for enrollment, but only 505 experimental villages were selected. Out of that sample, 320 randomized into the treatment arm and 185 villages randomized into a control arm. Treatment villages received benefits from PROGRESA in the summer of 1998. Control villages did not receive any benefits and would not receive benefits for another two years. Control villages were also not informed that they would receive benefits from PROGRESA in the future—something that caused ire among some Mexicans when the details and results of the experiment were made public in 2000 (Parker and Teruel, 2005, p. 210).

Today PROGRESA is considered a success. Gertler's (2004) analysis showed that preventive care utilization increased by more than 50 percent, and both children and adults

experienced significant improvements in health. Shultz's (2004) analysis also showed that school enrollment in treated villages increased, especially secondary school enrollment. Secondary school enrollment increased from 67 percent to about 75 percent for girls, and from 73 percent to about 77 percent for boys. Though perhaps the short-term effects appear small, it is worth bearing in mind the long-term effects extra schooling has on future wages (see Schultz, 2004).

The PROGRESA Experiment led to the program's nationalization, and today subsequent conditional cash transfer programs enjoy widespread political support in Mexico. The success of PROGRESA and its positive evaluation has also led to the proliferation of conditional cash transfer programs all over the world, thereby once again showing the influence randomized experiments can have on policy. Despite this, economists noted one serious flaw with the program worth highlighting. At the end of Gertler's (2004) analysis of PROGRESA's experimental results, he writes:

> While these results suggest that PROGRESA has had a positive effect on child health, they do not indicate which aspects of this complex program really matter. PROGRESA combines large cash transfers with requirements that individuals engage in a number of preventive health and nutrition activities. One cannot tell if the same results could have been achieved with just a large cash transfer and no behavioral requirements. In is also hard to distinguish between the relative effects of compliance with the various requirements. Answers to these questions would facilitate a better package and therefore improve the cost-effectiveness of the intervention (Gertler, 2004, p. 340).

This kind of explicit concern with cost-effectiveness is a distinguishing feature of many of the early development RCTs. Despite the success of PROGRESA as a poverty-alleviation program, Banerjee (2007) would eventually criticize PROGRESA for its cost-*in*effectiveness ($6,000 per additional child) and would emphasize the importance of development economists finding ways of achieving similar outcomes as PROGRESA with a less hefty price tag. Banerjee's criticism was motivated by a larger concern (discussed further on) that contributed to the rise of development RCTs: the importance of making foreign aid effective.

5. The Rise of Development RCTs

In this section I identify three factors which prompted the rise of experimental methodology in development economics. The three factors are (1) the credibility revolution in applied economics; (2) disillusionment with macro-centered approaches to development; and (3) the debate about foreign aid effectiveness. All three factors are interrelated, and I do not mean to suggest that each factor is completely independent of each other. Factors (1) and (2) overlap because the credibility revolution had not extended to macroeconomics. And factors (2) and (3) overlap because investment in foreign aid is one of the macro policy interventions that some development economists became skeptical of.

4.1 The Credibility Revolution

In the past thirty years, research in economics has undergone a drastic change. Looking at the articles published in top economics journals today and comparing them to the articles published sixty years ago, one is likely to notice less abstruse mathematics and more data analysis. Today's economists are not only working with empirical data but are also generating and collecting empirical data. When economists used empirical data in the past, it was typically borrowed data from publicly available data sets, e.g., government data on economic growth.

Economists' newfound interest in data analysis has been turned on itself, and a study by Hamermesh (2013) provides evidence that economics has undergone an "empirical turn". In samples from 1963 to 1983, over half the articles published in the top three economic journals were theoretical. In a 2011 sample, that share drops to less than 20 percent. The biggest change over the past six decades is an increase in empirical work involving data collection (as opposed to empirical work using borrowed data, which has remained constant). In 2011, 34 percent of articles in top economics journals fell under this category. This is a remarkable increase when

one considers the share of empirical work involving data collection in the samples from 1963 to 1993. The highest share that this category received was 8.8 percent in 1993 and the lowest was 2.4 percent in 1983. Also worth noting is the rise of experiments. The samples from 1963 and 1973 contain no experimental articles. In the 1993 sample, there is significant jump to 3.7 percent from .8 percent in 1983. In the 2011 sample, 8.2 percent of published articles in the top three journals are based on experimental findings.

A limitation of Hamermesh's (2013) study is that it only focuses on the top economics journals, so it may be the case that empirical work was still conducted and published in other journals (as is likely the case). But one can still conclude from Hamermesh's (2013) study that the economics profession now recognizes the importance of empirical research considerably more than it did sixty years ago.

A key factor for understanding this shift is what Angrist and Pischke (2010) have dubbed the "credibility revolution" in empirical economics. As I argue below, it was the credibility revolution that led economists to take up field experimentation—not inspiration from other economists conducting laboratory experiments, much less nostalgia for the second wave of field experimentation.

Angrist and Pischke (2010) trace the origins of the credibility revolution to Leamer (1983), who memorably urged applied economists to "take the con out of econometrics". Leamer (1983) complained that econometric analysis was not taken seriously because results were not robust to changes in specification (the process of selecting an appropriate functional form in a regression model and which variables to include) and that specification decisions made by his contemporaries were often arbitrary. The feature of Leamer's (1983) methodological outlook most worth highlighting is the presentation of the RCT as an ideal research design. Leamer

argued that "randomized experiments differ only in degree from nonexperimental evaluations of causal effects, the difference being the extent to which we can be confident that the causal variable of interest is independent of confounding factors" (Angrist and Pischke, 2010, p. 6). Leamer's methodological outlook was quickly corroborated by LaLonde (1986). LaLonde's (1986) influential article compared the results from a randomized evaluation of a work training program against the results that an econometrician using non-experimental data would have produced. The econometric results differed significantly from the results using random assignment (cf. Heckman and Smith, 1995), thereby demonstrating the value of randomization to a new generation of economists.

The chief influence Leamer (1983) had on applied economics is that research design should approximate an RCT as much as possible. This outlook is now reflected in Angrist and Pischke's (2009, 2014) popular econometrics textbooks, which begin with discussions of RCTs. But as Angrist and Pischke (2010) are careful to note, Leamer (1983) did not actually advocate that economists do randomized experiments. As I suggest below, it was the "demand shock" for credible econometric analysis that created disciplinary pressures that culminated in economists taking up field experimentation. As economists became preoccupied with delivering credible results and cleanly identifying casual effects, so did the pressure to approximate the ideal research design—the RCT.

An important methodological predecessor to the third wave of field experimentation was the use of "natural" experiments in the 1990s. Natural experiments are an intermediary between an observational study and a randomized experiment. In a natural experiment, serendipity allows a researcher to observe the effects of a causal factor outside their control against a controlled comparison. Examples of serendipitous circumstances include lotteries and arbitrary cut-offs or

thresholds. When treatment and control groups are formed through some serendipitous process, researchers can make the case that the serendipitous process approximates randomization, thereby allowing data to be analyzed like an RCT.

Perhaps the most famous natural experiment in economics is a study by Card and Krueger (1994) on the effects of minimum wage laws. In 1992, New Jersey raised the minimum wage by 80 cents per hour while Pennsylvania saw no such increase. Card and Kreuger (1994) seized the opportunity to test one of the most controversial policy-relevant implications of economic theory: raising the minimum wage leads to higher unemployment. Card and Kreuger (1994) collected employment data by conducting interviews with fast food restaurant managers in both states. They argued that the decision to build a fast-food restaurant in New Jersey over Pennsylvania was arbitrary. Since the restaurants near the New Jersey and Pennsylvania border were similar enough, the New Jersey restaurants formed a randomly assigned treatment group (minimum wage increase) and the Pennsylvania restaurants formed a control group—thereby allowing the data to be analyzed like an RCT. Though the results are still a subject of debate among economists, Card and Kreuger (1994) found that a small increase in the minimum wage did not lead to higher unemployment. Interestingly, it was not until two years *after* the methodological importance of RCTs was recognized with the Nobel prize that natural experiments had their turn in the spotlight. In 2021, David Card, Joshua Angrist, and Guido Imbens were awarded the Nobel Prize in economics for demonstrating how natural experiments can help answer important policy questions (Committee for the Prize in Economic Sciences, 2021).

Natural experiments have an obvious limitation: researchers must find chance events that produce “natural” treatment and control groups to analyze. For many theoretically interesting

research questions in economics, there may simply be no natural experiment to discover—"nature" is not always a willing research collaborator. Insisting on natural experiments as the only method of delivering credible econometric analysis would be unduly constraining on the economics profession.

One solution to the lack of supply for theoretically relevant natural experiments is simply to produce more of them.[14] Of course, once a researcher has control over the causal factor of interest, they are no longer conducting an observational study. Rather, they are experimenting. And so, the disciplinary pressures to deliver credible results ushered in a new wave of field experimentation in the 21st century. This ongoing wave of field experimentation is not just limited to development economics—mainstream microeconomics has also been transformed using field experiments. Research questions studied using field experiments include the use of monetary incentives in low-income schools, the real reasons people discriminate in the labor market, and the persistence of the gender pay gap. List and Gneezy's (2013) popular book, *The Why Axis*, provides an accessible survey of important developments in microeconomics made possible through field experiments.

There is documentation that the credibility revolution played an important role in prompting the rise of development RCTs. Some documentation comes from *Experimental Conversations* (Ogden, 2017), an anthology which contains interviews about development RCTs with prominent economists. Some of the interviewees discuss how the credibility revolution in economics led to the rise of development RCTs. One of the most notable interviews is with 2019 Nobel laureate Michael Kremer. Kremer's response to the question of what got him interested in

[14] Another is the use of instrumental variables (IVs), which I do not discuss here. See Angrist and Kreuger (2001) for discussion of IVs and their relation to natural experiments. Additional econometric methods are discussed in Angrist and Pischke (2009, 2014).

RCTs is noteworthy because it suggests that previous experimental work in economics had little to do with the rise of development RCTs. Kremer recounts:

> In 1994 when I started the work in Kenya, I was very much influenced by the movement for better identification in labor economics and public finance, but not by lab experiments. I see these traditions as independent, although there is now some convergence of the lab experiment and field experiment traditions (Ogden, 2017, p. 3).

Responding to the question of when the RCT movement began, Jonathan Morduch gives the following answer:

> The obvious context is economists' obsession with causality. We're obsessed. And we should be. Situations in development are rife with possible biases, and they can greatly distort understanding. As a graduate student and assistant professor at Harvard, you watched every empirical paper presented be picked apart based on causal claims (Ogden, 2017, p. 51).

Finally, like Morduch, Lant Pritchett also recalls the disciplinary pressures at places like Harvard (where Kremer and Morduch obtained PhDs) and MIT (where he obtained his PhD):

> So there was this trend that I can sit in a seminar in Cambridge and whatever instrument you propose, I can concoct a story in which your instrument is wrong. The end game of that is that grad students increasingly despaired of getting their dissertations approved because their professors could think of some crazy way in which it might be the case that their instrument didn't produce completely clean [causal] identification. The logical consequence of that trend is people resorting to RCTs so they could take the debate about identification off the table (Ogden, 2017, p. 141).

The "trend' which Pritchett is referring to is what Kremer calls the "movement for better identification"—or as Morduch more bluntly puts it, the "obsession with causality". Though none of these commentators use the term "credibility revolution", it is not a leap to conclude that economists' "obsession with causality"—whether healthy or not—is a consequence of the credibility revolution in applied economics that Angrist and Pischke (2010) document and celebrate in their article.

The credibility revolution only provides part of the story of what prompted the rise of development RCTs. To fully understand why development economists turned to experimental

methods, one needs to go all the way back to Adam Smith. The full title of Smith's *magnus opus* is crucial: *An Inquiry into the Nature and Causes of the Wealth of Nations* (Smith, 1993). Development economists in the post WWII period once again took up this inquiry. But as I suggest below, disappointment with repeated failures to make progress on Smith's inquiry, as well concerns about foreign aid effectiveness, provide a more complete picture of how development RCTs rose to prominence.

4.2 Disillusionment with Macroeconomics

In late May 2008, the Brookings Institution hosted a conference entitled *What Works in Development? Thinking Big and Thinking Small.* The aim of the conference was to bring leading development economists together to discuss two competing approaches to fighting global poverty: the prevailing macro-oriented school of thought, which had dominated development economics for half a century; and the new micro-oriented experimental approach to development, which was still in its nascent stage. The contributions from the conference are collected in a similarly titled anthology (Cohen and Easterly, 2009a). In the introduction, the conference organizers and editors of the anthology write:

> The starting point for the contributions to this volume, and the conference for which they were prepared, is that there is no consensus on "what works" for growth and development. The ultimate goal of development research—a plausible demonstration of what has worked in the past and what might work in the future—remains elusive (Cohen and Easterly, 2009b, p. 1).

As I suggest below, it was frustration with macroeconomic approaches to development that inspired some development economists to start "thinking small".

The primary factor that led to skepticism about "thinking big" is the failure of past macroeconomic policy reforms to spur growth in developing countries. Cohen and Easterly (2009b, p. 2) identify three major episodes:

1. The early big push in foreign aid (especially in the most aid-intensive continent, Africa.)
2. Structural Adjustment (also known as the Washington Consensus) in the 1980s and 1990s.
3. "Shock Therapy" in the former Communist countries.

The failure of the early big push in foreign aid is especially salient since, to this day, there is still considerable academic and policy debate surrounding the effectiveness of foreign aid. As I suggest further on, this policy issue is especially important in understanding the rise of development RCTs. But first, I focus on disillusionment with macroeconomics more generally.

One does not need to be an expert in economics to conclude that the macroeconomic interventions listed above have failed to deliver on their ambitious promises. Some of the experts also took note, and Cohen and Easterly (2009b) suggest that these macro policy failures intensified some development economist's skepticism about prevailing macro methodology. Economics' status as "the Queen of the social sciences" was clearly in question, and the obvious culprit was prevailing methodology.

It is worth noting that Angrist and Pischke's (2010) praise of economics' empirical turn focuses on microeconomics, and a key section of their article suggests that macroeconomists are yet to fully join the credibility revolution. Cohen and Easterly single out the cross-country growth regression literature as one of the main culprits of the credibility crisis in macro-oriented development economics. Cohen and Easterly (2009b) echo Leamer (1983) and Angrist and Pischke (2010) when they elaborate on what they see as the failure of the growth regression literature:

> Having a long list of possible controls to play with, researchers, found it easy enough to arrive at significant results, and using the abundant heuristic biases that make it possible see patterns in randomness, convinced themselves that the significant results were from the "right" specification, and that the others (usually unreported) were from the "wrong" ones (Cohen and Easterly, 2009b, p. 3).

It is worth highlighting the intellectual honesty contained in this statement. While Jessica Cohen is one of the early practitioners of experimental development economics, William Easterly is a macro-oriented development economist. As I discuss in the next section, Easterly's macro-oriented research on foreign-aid effectiveness also played an important role in prompting the rise of development RCTs.

The repeat failure to identify the causal determinants of economic development led some key development economists to entirely give up on Adam Smith's inquiry. Abhijit Banerjee's perspective is especially revealing. After also criticizing the growth regression literature, Banerjee (2009, p. 219) ends his contribution to Cohen and Easterly's anthology on a pessimistic note:

> Perhaps making growth happen is ultimately beyond our control. Maybe all that happens is that something goes right for once (privatized agriculture raises incomes in rural China) and then that sparks growth somewhere else in the economy, and so on. Perhaps we will never learn where it will start or what will make it continue (Banerjee, 2009, p. 219)

But Banerjee's attitude is not entirely defeatist; he also tries to strike an optimistic tone:

> The best we can do in that world is to hold the fort until that initial spark arrives: make sure that there is not too much human misery, maintain the social equilibrium, and try to make sure that there is enough human capital around to take advantage of the spark when it arrives. (Banerjee, 2009, pp. 219-220).

While Banerjee never claims that "thinking big" is worthless, his perspective is important for understanding the shift towards "thinking small" in development economics. Though it is hard to gather evidence for such a claim, it is likely that Banerjee's frustration was shared by many of his contemporaries. If this is right, it is no surprise that many development economists jumped on the opportunity to try a different approach to studying development. Unlike with "thinking big", focusing on smaller questions allows for discernable short-term progress. And if we believe the narrative of Nobel Prize Committee and many other development economists, progress *is*

being made. However, this is small-scale, incremental progress—not the kind of progress that would result in the overnight elimination of global poverty.

4.3 The Foreign Aid Debate

The effect of foreign aid on economic growth is, to date, one of the most controversial debates in development economics. Economists such as Jeffrey Sachs (2005) argue that large-scale pushes in foreign aid can lead to the end of global poverty within our lifetimes. What readers may be surprised to learn is that foreign aid skepticism played an important role in prompting the rise of development RCTs. If there is one macro question that pioneers of development RCTs agreed progress had been made on, it was the *ineffectiveness* of some forms of foreign aid. As I suggest below, skepticism about foreign aid fueled the rise of development RCTs by providing the new research program an air of ethical urgency.

The case for foreign aid from high-income to low-income countries can be traced back to the post WWII Period. According to the prevailing macroeconomic theory of the time, saving and investment in capital is a key determinant of economic growth. The economic argument for foreign aid was based on the idea that developing countries do not grow due to domestic gaps in saving and investment. Large foreign aid efforts could thereby fill in saving and investment gaps and put developing countries on the path to economic growth. An influential early critique of foreign aid is Bauer (1971). Bauer (1971) argues that foreign aid is neither necessary nor sufficient for promoting economic growth. If country circumstances are not already conducive to economic growth, then foreign aid will be ineffective and wasteful. And if country circumstances are conducive to aid being effective in promoting economic growth, then foreign aid is not needed. Bauer's key insight is that foreign aid crowds out a government's incentive to become accountable to its citizens. As a result, institutional quality—one of the key determinants of

economic growth foreign aid skeptics propose—never reaches the threshold needed for economic development to take off. One can find variations of this theme throughout much of the literature expressing skepticism about foreign aid (see Ravallion, 2014).

Perhaps the most prominent foreign aid skeptic today is William Easterly. Easterly's best-selling book, *The White Man's Burden* (Easterly, 2006), provides a devastating critique of half a century of foreign aid efforts. Yet there are misconceptions about Easterly and other foreign aid skeptics' (e.g., Moyo, 2009; Deaton, 2013) criticisms that need to be cleared up to understand the connection to development RCTs.[15] There is an important distinction between *official* aid and *unofficial* aid. Official aid is either bilateral—meaning it is given from government to government—or multilateral—meaning it is allocated from government to government via a multilateral organization such as the World Bank. Unofficial aid comes from international charities and other non-government organizations (NGOs) operating in developing countries. By some estimates, unofficial aid accounts for 25-30 percent of transfers from rich countries to poor countries (Deaton, 2013, p. 176).

In *The White Man's Burden,* Easterly explicitly targets official aid efforts for not only being wasteful, but also potentially harmful. But a careful reading of *The White Man's Burden* also reveals that Easterly was enthusiastic about unofficial aid efforts. In the final chapter ("The Future of Western Assistance") of *The White Man's Burden*, Easterly lays out his vision for a path forward. Notably, Easterly makes it clear he is not against unofficial aid efforts. He writes that "Once the West is willing to aid individuals rather than governments, some conundrums that tie foreign aid up in knots are resolved" (Easterly, 2006, p. 368). One reason why Easterly argues official aid is wasteful is due to the lack of accountability mechanisms put in place to ensure aid

[15] I discuss Deaton's (2013) views in Picchio (2023).

is effectively used by recipient governments. But unlike foreign governments, NGOs operating in low-income countries are subject to different incentive structures. Starting in the early 2000s, Easterly noticed a push to rigorously evaluate unofficial aid efforts using RCTs. This trend was in part driven by the rise of large-scale philanthropic organizations, such as the Bill and Melinda Gates Foundation, taking an interest in development, and wanting to see evidence that their contributions were being put to effective use (e.g., see Gates, 2013). But the trend was also being driven by Easterly's micro-oriented colleagues, who, aware of past foreign aid failures, had begun convincing NGOs to subject their interventions to rigorous evaluation. It is interesting that, despite being a macro-oriented development economist, Easterly was one of the earliest members of his profession to praise the experimental approach to development economics.

Though it is, once again, hard to make any claims about the collective attitudes of development economists twenty years ago, there is evidence that pioneers of development RCTs harbored doubts about foreign aid effectiveness. For example, one needs to look no further than Banerjee's (2007) influential essay, *Making Aid Work*. As the title of the essay suggests, Banerjee was on Easterly's side of the foreign aid debate. Banerjee (2007) notably criticizes global institutions like the World Bank for failing to evaluate its development projects rigorously. And as noted earlier, Banerjee also singles out previous government-led anti-poverty programs, such as PROGRESA, for being insensitive to cost-effectiveness considerations. Though Banerjee labels himself an "aid optimist" in the essay, what he means is that foreign aid *can* be effective in the future. The major theme of Banerjee's essay is that foreign aid organizations need to think harder about how their money is spent. Banerjee's prescriptions should come as no surprise: use RCTs to figure out "what works" so that scare resources can be used responsibly.

It should go without saying that global poverty is one of the most significant ethical challenges of our times. But at its core, Banerjee and Easterly (and other development economists') skepticism about foreign aid are driven by an ethical concern. Their calls to make foreign aid effective are buttressed by the ethical concern that scarce resources should be used responsibly.[16] When so many people live in extreme poverty, it is imperative that foreign aid is used on effective interventions. Easterly's impassioned call to action at the end of *The White Man's Burden* is worth quoting at length:

> Put the focus back where it belongs: get the poorest people in the world such obvious goods as the vaccines, the antibiotics, the food supplements, and the improved seeds, the fertilizer, the roads, the boreholes, the water pipes, the textbooks, and the nurses. This is not making the poor dependent on handouts; it is giving the poorest people the health; nutrition; education; and other inputs that raise the payoff to their own efforts to better their lives (Easterly, 2006, p. 381).

As Easterly quickly acknowledges, however, matters are not as simple as handing out resources to the poor. If there's a key lesson that development economists learned from the failures of past foreign aid efforts, it is that simply throwing resources and opportunities at poor people will do little to reduce global poverty. A recurring theme in Banerjee and Duflo's (2011) *Poor Economics* is that even the most well-intentioned policy interventions will be ineffective unless steps are taken to understand behavioral responses to said policies.

And so, we are now in good position to see why RCTs have contributed—and can continue to contribute—to increasing the practical knowledge necessary to make successful poverty-alleviating interventions. In the final section of this introductory chapter, I highlight some of important experiments conducted by development economists.

---

[16] I expand on this theme in chapter 3.

6. RCTs in Development Economics: Some Highlights

This final section is not a comprehensive survey of all the important work done in experimental development economics (see Banerjee and Duflo, 2011; Karlan and Appel, 2011). I focus on experiments that have received the most attention and which help illustrate some of the themes discussed in previous sections and throughout the rest of the dissertation.

6.1. Deworming

One of the earliest and most well-known development RCTs is Miguel and Kremer (2004). In this experiment, Miguel and Kremer evaluated the effects of a primary school deworming program carried out by an NGO operating in Kenya. In developing countries, worm-like parasites cause infected persons to become malnourished and susceptible to chronic illness. Effective treatments for intestinal parasites do exist.[17] To conduct their evaluation, Miguel and Kremer randomly assigned seventy-five schools into three groups of twenty-five schools each. The first group of schools were given free deworming drugs while the second and third groups were used as a control. A year later, the second group of schools were given free deworming drugs while the third group remained a control. Eventually, all three groups of schools were provided with free deworming drugs.

Evaluations of deworming programs are nothing new. Perhaps unsurprisingly, Miguel and Kremer found that children who received the deworming treatment had better health outcomes. What makes Miguel and Kremer's (2004) evaluation important—and what makes it emblematic of the experimental approach to development economics—is that they (a) also focused on socioeconomic outcomes unrelated to health and (b) analyzed the cost-effectiveness of the intervention. With respect to (a), Miguel and Kremer found that the deworming program

[17] In chapter 2, I discuss ethical issues with this aspect of Miguel and Kremer's experimental design, which is typical of most development RCTs.

led to a 7.5 percentage point average gain in primary school participation in treatment schools—a reduction of overall school absenteeism by at least one-quarter. As we saw with PROGRESA, the long-term effects of increased schooling should not be ignored.[18] But unlike PROGRESA, what is noteworthy about Miguel and Kremer (2004) is the remarkable cost-effectiveness of the intervention. As they note, previous evaluations of deworming programs ignored the external effects of deworming. Miguel and Kremer (2004) found that children living within 3 kilometers of treatment schools also had lower infection rates and higher school attendance. One of the central policy lessons from Miguel and Kremer (2004) is that the positive externalities associated with deworming are by themselves sufficient to not only to fully subsize deworming treatment, but perhaps even pay people to receive treatment.

Miguel and Kremer's (2004) deworming study is part of a series of development RCTs that helped inform an important policy debate: the impact of pricing on uptake of health and education services and/or products. A full review of all the experiments that contributed to this important debate is not possible here (see Kremer and Holla, 2009). But in the next chapter, I introduce another influential experiment that contributed to this debate, and which will play a key role throughout this dissertation. At the end of chapter 3, I will discuss how the experiment in question helped advance the debate about preventative health product pricing.

6.2. Microcredit

Another major contribution is Banerjee et al.'s (2015b) evaluation of a group-lending microcredit program. Up until the mid 2010s, microcredit had been generating considerable interest and hope as an effective solution to global poverty. The zenith of microcredit's popularity came in 2006, when Mohammed Yunus and Grameen Bank—the pioneers of

[18] It is worth noting that Miguel and Kremer (2004) also found that the deworming program did *not* increase test scores.

microcredit—were awarded the Nobel Peace Prize. The motivation for microcredit is simple enough: access to credit can help would-be entrepreneurs start businesses and climb out of poverty. But banks are generally reluctant to lend credit to people who lack collateral, steady employment, or a verifiable credit history. Microfinance institutions (MFIs) fill this investment gap by offering low-interest loans to the poor. As expected, MFIs report their customers earn remarkable profits. But these claims are potentially rife with selection bias. Are the profits due to microcredit, or are customers who seek microcredit already so entrepreneurial that they would succeed without access to microcredit? Before the advent of development RCTs, these kinds of difficulties made evaluating microcredit programs difficult.

To evaluate the effects of microcredit more rigorously, Banerjee et al. (2015b) collaborated with a for-profit MFI to evaluate a group-lending program in the city Hyderabad, India. 104 neighborhoods were selected for the experiment, half of which were randomly assigned to receive the treatment, which in this case was access to the services of the MFI. The other half served as a control group and did not (initially) have access to the microcredit program. While the results showed microcredit has some positive effects, the results hardly support the enthusiasm for microcredit as a solution to global poverty. Banerjee et al. (2015b) found that investment and profits for pre-existing small businesses increased (confirming suspicions of self-selection) but consumption did not significantly increase. Banerjee et al. (2015b) also report that there were no significant changes in key socioeconomic outcomes such as health, education, and women's empowerment. In fact, two years after the control group gained access to the program, Banerjee et al. (2015b) found that few significant differences between treatment and control groups.

Along with five other RCTs on microcredit, Banerjee et al. (2015b) was published in a special issue of the *American Economic Journal: Applied Economics* (see Banerjee et al., 2015a). The special issue has proven influential in the debate about microcredit. For many development economists, the main takeaway from the special issue is that microcredit has "some modestly positive, but not transformative effects" (Banerjee et al., 2015a). The lukewarm verdict on microcredit is nowadays hailed as a triumph of the experimental approach to development economics. The scientific background for 2019 Nobel Prize in economics makes it no secret that this was one of the most important ways in which RCTs contributed to our understanding of global poverty.[19]

6.3. Multifaceted Interventions

If there is a crown jewel of the experimental approach to development economics, it is without a doubt the evaluation of a multifaceted poverty-relief program published in *Science* (Banerjee, 2015c). This study was comprised of six RCTs in Ethiopia, Ghana, Honduras, Indian, Pakistan, and Peru with a total of 10,495 participants. To get a sense of the scale of this study, consider that the RAND Health Insurance Experiment had 3,958 participants (Brook et al., 1983). The study also required cooperation and coordination from various local and international NGOs, as well as MFIs in each country. The intervention contained six different components: (1) a one-time productive asset transfer; (2) regular food or cash transfers; (3) technical skills training on managing the productive assets; (4) high-frequency home visits; (5) access to savings account; (6) some health education, basic health services, and/or life-skills training. In addition, there were over 50 outcomes of interest. These outcomes fall into ten general categories: (1) consumption; (2) food security; (3) assets; (4) finance; (5) time use; (6) income and revenues; (7)

[19] However, this narrative is not accepted by all, and for pushback see Bédécarrats et al. (2021c).

physical health; (8) mental health; (9) political involvement; (10) women's decision-making. In brief, Banerjee et al. (2015c) found statistically significant impact on all 10 key outcomes.

Attempting to summarize all the details of the *Science* study would still not do it justice—what the authors accomplished is truly remarkable and unprecedented. After all, the study spanned "three continents, and different cultures, market access and structures, religions, subsistence activities, and overlap with government safety net programs." (Banerjee et al., 2015c, p. 1). As the authors note, "This diversity should give us a high level of confidence in the robustness of the impact to variations in both the context and implementation agency" (Banerjee et al., 2015c, p. 1). Yet a careful reading of the discussion section of the article reveals sober analysis. Nowhere do the authors suggest that they have definitively unlocked the key to economic development. Nor do they promise that the intervention can dramatically reduce poverty in a short period of time. Though the *Science* study is one of the "biggest" development RCTs to date, the authors are still committed to "thinking small". The authors acknowledge that there is still much more to learn about how the intervention works, and only suggest that—at this moment in time— the multifaceted intervention is a cost-effective way of improving the lives of the poor.

Chapter 2: What is the Standard of Care in Experimental Development Economics?

1. Introduction

In this chapter, I identify and address one salient problem with the use RCTs in development economics: the routine practice of using the status quo as a control when designing and conducting a development RCT. All the experiments discussed at the end of the last chapter employed this practice. But to further illustrate the practice, below I describe a paradigmatic example of the use of RCTs in development economics. I will refer to this experiment throughout this dissertation because it perfectly illustrates many of the ethical issues with development RCTs.

In 2007, Jessica Cohen and Pascaline Dupas studied the effects of subsidies on preventative healthcare products in low-income countries. To conduct their study, Cohen and Dupas (2010) selected twenty prenatal clinics in Kenya and varied the price at which they could sell insecticide-treated bed nets (ITNs) to pregnant women. ITNs are known to prevent malaria infection and as a result are highly effective in reducing maternal anemia and infant mortality; their widespread use also generates positive externalities in the form of health benefits to non-users. In the experiment, sixteen of the clinics were randomly assigned to four different groupings corresponding to different subsidy levels ranging from a full subsidy (free) to a 90% subsidy. The remaining four clinics were used as a control and did not have access to heavily subsidized bed nets.

In the next chapter, I will discuss the findings from this experiment in more detail. For present purposes, what is worth highlighting is that using the status quo as a control, such as in the experiment above, unsurprisingly causes ethical uneasiness in outside observers. Many of the

residents of low-income countries in which development economists conduct experiments lack access to important socioeconomic resources and opportunities, e.g., preventative healthcare products such as ITNs. To put it more bluntly, research participants in development RCTs live in poverty.[20] Of course, the motivation behind conducting development RCTs is often to identify effective poverty-alleviating interventions, but matters are complicated by the fact that the status quo in these low-income countries is categorically unjust. This raises the possibility that development economists are flouting some important ethical requirement by conducting experimental research in such settings—an ethical complaint worth taking seriously.

The uneasiness many may feel about this aspect of experimental development economics raises an important ethical question: What level of socioeconomic resources and opportunities are owed to research participants by development economists? At its core, this is a question about distributive justice, and in particular, distributive justice in *non-ideal circumstances.* To assist with the inquiry at hand, I will be drawing on helpful concepts and principles from the well-established literature on the ethics of medical research. Specifically, I focus on the concept of a *standard of care* and the associated *principle of standard care*.[21] As will become clear, the debate surrounding the correct interpretation of the principle of standard care in international medical research has yielded concepts and distinctions that are relevant to the question of what is owed to research participants in development RCTs. The key objective of this chapter is to determine what exactly the "standard of care" is in experimental development economics. In brief, I will argue that (barring certain circumstances) the "standard of care" in experimental

[20] See Banerjee and Duflo (2007) for a detailed analysis of the economic lives of the global poor.
[21] The term "principle of standard care" may sound odd to readers who typically think of the notion of a standard of care as an ethical concept rather than as an ethical principle. My choice of terminology follows Hawkins (2008).

development economics is the status quo level of socioeconomic resources and opportunities that research participants have access to.

The rest of this chapter is structured as follows. Section 2 introduces some background concepts and distinctions from general ethical theory. Section 3 introduces the principle of standard care and provides an overview of one associated controversy with its application in medical research. Section 4 advances an argument by elimination to establish that the "standard of care" in experimental development economics is the status quo. Specifically, this section rules out professional obligations, natural duties, and institutional obligations as providing the normative basis for a "standard of care" in experimental development economics. Section 5 introduces an important qualification based on the natural duty of rescue that is meant to address the concern that my account is overly permissive. Section 6 concludes.

2. Background Concepts and Distinctions

Before proceeding, a word of caution: while I will be arguing that development economists are permitted to use the status quo as a control when designing and conducting a development RCT, this should not be taken to suggest that no other ethical considerations bear on the permissibility of conducting development RCTs. The focus of this chapter is a specific kind of ethical concern, namely, a concern about distributive justice. Like many, I take justice to designate an ethical concern with what is owed to persons and by whom. What is owed to someone is fundamentally linked to their claim-rights or, more generally speaking, their entitlements (Wenar, 2021). Questions of distributive justice, as I understand them here, are a subclass of questions about justice which specifically deal with the level of socioeconomic resources and opportunities persons are entitled to. For present purposes then, non-ideal circumstances are those in which persons are not being fully granted what they are entitled to, and further, may not be granted

what they are fully entitled to in the foreseeable future.[22] It is worth noting that throughout, I remain agnostic on what the *exactly* correct theory of distributive justice requires but maintain that the global poor's share of socioeconomic resources and opportunities is well-below what any plausible theory would specify.[23]

In what follows, I assume that the ethical requirements of individuals fall into two general categories. First, there are *natural duties*; these are requirements that "apply to us without regard to our voluntary acts" and "have no necessary connection with institutions or social practices" (Rawls, 1971, p. 114). More simply, these are duties everyone has in virtue of being persons; they would apply to us in a (hypothetical) state of nature in which institutional arrangements do not exist. Natural duties should be contrasted with *obligations*, which do arise out of our voluntary acts. These voluntary acts include "the giving of express or tacit undertakings, such as promises and agreements" (Rawls, 1971, p. 113). Both natural duties and obligations are often, but not always, correlative with the entitlements of persons. Natural duties (or obligations) of rectification are always correlative with someone's entitlement to compensation for a past wrong. If I have a natural duty (or obligation) to rectify some past wrong, then this implies some person has a correlative entitlement to my rectification. The natural duty of beneficence is a noteworthy example of an ethical requirement that is not

---

[22] A more elaborate definition of non-ideal circumstances can be made by reference to the distinction between ideal vs. non-ideal theory found in the work of Rawls (1971, pp. 245-246; 1999, pp. 4-6). For Rawls, ideal theory operates under two assumptions: "(i) all relevant agents comply with the demands of justice applying to them; and (ii) natural and historical conditions are favourable – i.e., society is sufficiently economically and socially developed to realize justice" (Valentini, 2012, p. 655). Non-ideal circumstances are those in which (i) and/or (ii) do not obtain. Another important approach to non-ideal theory has been advocated for by Mills (2005, 2008, 2009). On the account of non-ideal theorizing Mills envisions, the relevant normative project is "the adjudication of competing policies for redressing social injustice" (Mills, 2009, p. 182). I discuss corrective justice in section 4.4 on global institutional obligations.

[23] This is in line with engaging with another approach to non-ideal theory known as *anticipatory* theory. On this approach, "non-ideal theory has to make assumptions about the minimum requirements that any *plausible* and *complete* ideal theory of justice will include. In this vein, it can define targets for practical action *before* a complete ideal has been worked out, even in outline" (Sreenivasan, 2007, p. 221). Also see Sen (2009).

correlative with any specific person's entitlements. While we may be required promote the good of others in virtue of being persons, this does not mean any specific person is entitled to our beneficence.

An important ethical requirement that I will set aside going forward is the natural duty of non-maleficence. I take it for granted that all persons are required to not (deliberately) harm or injure others, and all persons have a correlative entitlement to not be (deliberately) harmed or injured without good reason. This should not be taken to suggest that development economists should be unconcerned with potential harms to research participants. Development economists are clearly bound by the natural duty of non-maleficence and RCT design should reflect this. I set aside concerns about potential harms to research participants to focus on the main issue at hand, namely, concerns about distributive justice. It is worth stressing, however, that questions about what is owed to research participants do have considerable bearing on whether they are harmed or not. Hawkins (2006) notably argues that certain cases of *positive obligation flouting* count as harms. For example, if a physician fails to recommend a known effective treatment for a serious medical condition, her patient has a plausible ethical (and legal) basis for claiming they were harmed. If development economists are routinely failing to comply with positive ethical requirements, then they may also be harming research participants.

3. The Principle of Standard Care

3.1 Therapeutic Obligations

The ethical requirements of physicians have traditionally been regarded as obligations since these ethical requirements stem from the voluntary act of choosing a particular profession. As a result, the ethical requirements of physicians can be further regarded as *professional obligations*. One well-recognized professional obligation physicians have to their patients is the *therapeutic*

obligation: a physician should recommend a treatment *T* for some condition *C* if and only if *T* is accepted by the medical community as effective in treating *C*. The physician's therapeutic obligation provides the basis for the concept of a standard of care and in turn the principle of standard care in medical research ethics.[24] A succinct and influential statement of the principle first appeared in the World Medical Association's (WMA) 1975 Declaration of Helsinki. Paragraph II.3 reads: "In any medical study, every patient—including those in the control group, if any—should be assured of the best proven diagnostic and therapeutic method" (WMA, 1975). Despite the many flaws with this formulation of the principle of standard care (see Levine, 1999), the motivation for the principle is intuitively appealing: researchers should not withhold effective medical care from research participants, and thereby knowingly make them worse-off than they would have otherwise been had they not partaken in a study, for the sake of obtaining valuable medical knowledge.[25] In other words, the principle of standard care specifies that researchers conducting a study *owe* research participants some pre-specified level of medical resources and attention. Further, the principle of standard care specifies that research participants are *entitled* to this pre-specified level of medical care. Providing research participants with anything deemed *ex ante* inferior to the standard of care for some medical condition is therefore wrong.

3.2 The Standard of Care Debate

Despite the intuitive appeal of the principle of standard care, its interpretation and application have been controversial. These controversies stem from the use of placebo-controlled trials

[24] This is not to say that all commentators agree that physicians are bound by their therapeutic obligations in research contexts. Miller and Brody (2002, 2003, 2007) have vigorously argued against this position. Also see Hawkins (2006), Kukla (2007), and MacKay (2014) on this issue. For the purposes of this dissertation, I avoid taking a stance on issues surrounding the philosophical foundations of medical research ethics.

[25] The most recent version of The Declaration of Helsinki (WMA, 2013) still uses the language of "best proven intervention(s)" but now allows for exemptions to the principle of standard care.

(PCTs) instead of active-controlled trials (ACTs) in medical research.[26] As the name suggests, in a PCT the control arm receives a placebo, i.e., some type of medically inert substance. In an ACT, the control arm receives a therapeutic intervention of some sort. If the study is designed in line with the principle of standard care, then the control arm in an ACT receives the "best proven diagnostic and therapeutic method." There is no denying that being assigned to the control arm of a development RCT is methodologically different than receiving a placebo in a medical study.[27] However, both development RCTs and (medical) PCTs raise similar ethical questions about what researchers owe research participants, and so it is worth reviewing one of these controversies here.

The most well-known controversy involving the use of PCTs is the short-course AZT trials of the 1990s.[28] This controversy arose in the context of international rather than domestic medical research and is in many ways analogous to the kind of research done by development economists. The goal of discussing the AZT trials is not to render an ethical verdict on the case, which many have attempted.[29] Rather, it is important to review the details of the case because, as we will see, a host of additional helpful concepts and distinctions emerged from the debate over whether the AZT trials were ethically justifiable or not.

The details of the controversy are as follows. In the 1990s, sixteen RCTs were designed and conducted in eleven countries: Burkina Faso, the Dominican Republic, Ethiopia, Ivory

---

[26] The controversy over the use of PCTs in high-income countries was sparked by Rothman and Michels (1994). See Temple and Ellenberg (2000) and Emmanuel and Miller (2001) for commentary.

[27] One obvious difference is that development RCTs are conducted in the field. Another important methodological difference is that medical trials are often double-blind, meaning that neither researcher nor research subject knows which arm of the trial they belong to until the study is over. Abramowicz and Szafarz (2020) discuss some of the ethical implications of this methodological difference in the context of experimental development economics.

[28] See Hawkins and Emmanuel (2008b) for a more thorough overview of the controversy surrounding the short-course AZT trials.

[29] There are simply too many articles discussing the short-course AZT trials to list here. The most influential literature is discussed throughout the rest of this section.

Coast, Kenya, Malawi, South Africa, Tanzania, Thailand, Uganda, and Zimbabwe. The motivation behind these trials was to address the problem of maternal-fetal HIV transmission in developing countries. However, as early as 1994, a standard of care treatment for maternal-fetal HIV transmission had been established in high-income countries such as the U.S. This treatment plan, known as the 076 regimen, was expensive, lengthy, and difficult to administer. At the time, the 076 regimen cost $1,000 per woman and involved large quantities of AZT (the trade name of the drug zidovudine) to be administered in an elaborate schedule over a minimum of 12 weeks starting in the second trimester of pregnancy. Additionally, the average annual health budgets of the countries in question were $10 per person. The goal of the short-course AZT trials was to determine whether a simpler, less expensive version of the 076 regimen would be effective in reducing maternal-fetal HIV transmission.

The controversy surrounding the AZT trials stemmed from the fact that fifteen of the trials used a placebo-control when a proven treatment already existed, i.e., the 076 regimen. The most well-known critics of the AZT trials were Angell (1997) and Lurie and Wolfe (1997), who argued that the use of a placebo-control in the AZT trials invoked an ethically unacceptable double standard. According to these critics, the placebo-controlled AZT trials would have clearly been unethical if they had been conducted in a high-income country where the 076 regimen had been established as the standard of care, and therefore the same trial should be deemed unethical in the developing world.

An initial way to understand the basis for the ensuing "standard of care debate" was an ambiguity in the *relevant reference point* (London, 2000) for the principle of standard care: "When Helsinki calls for the 'best proven therapeutic method' does it mean the best therapy available in the world? Or does it mean the standard that prevails in the country in which the trial

is conducted?" (Levine, 1998, p. 6). Critics of the AZT trials could be seen as advocating for the former, *global* interpretation of the principle of standard care whereas defenders of the AZT trials could be seen as advocating for the latter, *local* interpretation. Further, we can connect these two positions to more general debates about distributive justice by following Emmanuel (2012) in characterizing the global interpretation as reflecting a broad commitment to cosmopolitanism (e.g., Beitz, 1979) and the local interpretation as reflecting a broad commitment to statism (e.g., Rawls, 1999).

One important consideration in favor of the local interpretation is that medical research should address the unique health needs of communities in the developing world and that, due to conditions of fiscal scarcity, applying a global standard of care would hamper this ethically important objective (Varmus and Satcher, 1997).[30] Development economists will likely be sympathetic to this consideration and the local interpretation of the principle of standard care. After all, the RCTs designed by development economists are typically meant to address problems that are unique to the developing world, most notably extreme poverty. There would be little point to conducting a development RCT on a prospective socioeconomic intervention if the governments of low-income countries could, given present circumstances, provide research participants with the global "standard of care" for education, healthcare, and finance found in high-income countries.

Yet when it comes to the correct interpretation of the principle of standard care, London (2000) has introduced a further, more fundamental distinction between the *de facto* standard of care and the *de jure* standard of care. According to a *de facto* interpretation of the principle of

[30] Crouch and Arras (1998) advanced this consideration in favor of a local interpretation despite ultimately concluding that the AZT trials were impermissible. At the time, they incorrectly believed that the short-course regimen was not going to be made available to the populations in question.

standard care, the "best proven treatment" is determined by *actual* medical practice. The *de facto* interpretation can then be combined with a local or global reference point. Applied to the AZT trials, critics could be seen as arguing that the use of placebo-controls was wrong because actual medical practice for preventing maternal-infant HIV transmissions in high-income countries was the 076 regimen; this would be a global *de facto* interpretation of the principle of standard care. On a local *de facto* interpretation, the use of placebo-controls was permissible because there was effectively no established medical practice for dealing with maternal-infant HIV transmission in the countries in which the AZT trials took place. Hence, administering a placebo did not place the control group below the standard of care.

Instead of focusing on actual medical practice, a *de jure* interpretation of the principle of the standard of care looks to what established medical practice *should* be. Like before, the *de jure* interpretation can be combined with a local or global reference point, though as we see below, this is less important when the *de jure* interpretation is grounded in the therapeutic obligations of physicians. According to London (2000), critics of the AZT trials such as Angell (1997) could more charitably be interpreted as advancing a *de jure* interpretation of the principle of standard of care. Even if the local *de facto* standard of care for some medical condition is virtually non-existent in the developing world, research participants in low-income countries may still be entitled to more than the local *de facto* standard of care by researchers in virtue of their therapeutic obligation. As London (2000, p. 399) notes, "the *de jure* standard is founded upon the researcher's obligation to ensure that subjects of medical trials are not knowingly exposed to foreseeable and preventable harms." Since the researchers involved in the AZT trials *knew* that the best proven treatment for reducing maternal-fetal HIV transmission was the 076 regimen, they violated the principle of standard care *because* they violated their therapeutic

obligation. These researchers therefore acted wrongly regardless of whether they were members of the relevant host communities or the larger global medical community.

Not all commentators (including London, 2000) agree with the characterization of the *de jure* interpretation presented above. For some commentators, whether an intervention should be made accessible to a population "depends on a complex combination of factors including economic factors, narrowly medical facts, social support systems, local preferences and values, and much more" (Kukla, 2007, p. 178). But for the principle of standard of care to consider these broader, practical considerations one then needs to give up the traditional basis for the principle, i.e., the physician's therapeutic obligation, and find a new basis.[31] On this broader interpretation of the principle of standard care, physicians conducting medical research will perhaps routinely violate their therapeutic obligation by conducting a PCT in the developing world, but they would not necessarily be violating the local *de jure* interpretation of the principle of standard care—and it is the latter which is ultimately supposed to be the more important ethical consideration.

3.3 Taking Stock

Settling the debate over the correct interpretation of the principle of standard care in international medical research is not my main concern here. The goal of reviewing the standard of care debate is only to introduce a framework for thinking about what research participants are owed in development RCTs. It is possible to quickly rule out any straightforward global interpretation of principle of standard care as a viable ethical principle in the context of experimental development economics. For one, it is not clear what the global *de facto* "standard of care" is, given inequities in access to healthcare, education, and finance in high-income countries (Kukla, 2007). Further, while a global *de jure* interpretation may reflect commitment to a theory of

[31] MacKay (2015) offers one compelling basis by appealing to the institutional obligations of researchers. I discuss institutional obligations further below.

distributive justice worth aspiring towards, insisting on such an interpretation would be a paradigm example of letting an ideal become the enemy of the good—or more specifically in this case, potentially obtaining socially valuable evidence relevant to economic development.

This leaves the local *de facto* and local *de jure* interpretations as the two viable candidates, though I will circle back to discussion of a global *de jure* position further on. In the next section I turn to the question of whether the local *de jure* standard of care—understood in the context of experimental development economics as an unspecified level of socioeconomic resources and opportunities owed to research participants by development economists—is anything over and above the status quo, i.e., the local *de facto* standard of care. The key challenge to establishing a local *de jure* interpretation is identifying an ethical requirement that can serve as the basis for a principle of standard of care in experimental development economics. Below, I cast doubt on the possibility of meeting this challenge. Ultimately, I will conclude that in experimental development economics, the local *de jure* standard of care is the status quo, i.e., the local *de facto* standard of care.

4. Grounding the Principle of Standard Care

4.1 Professional Obligations

In what follows, I proceed by making an argument by elimination to determine whether development economists owe their research participants more than the status quo, i.e., the local *de facto* standard of care. Above, we saw how the professional obligations of physicians can be invoked as the basis for a *de jure* interpretation of the principle of standard care. While development economists may have some professional obligations in virtue of being members of the class of professionally trained economists (e.g., obligations not to engage in fraudulent

research practices), there is no currently recognized analogue to the physician's therapeutic obligation for economists.

Some may also be quick to point out that there is also no historically significant undertaking equivalent to the Hippocratic oath on which to base a professional code of ethics for economists. But this suggests that an ancient undertaking is an adequate justification for the professional obligations of physicians. More plausibly, one could maintain that the professional obligations of physicians are a type of contractual *role obligation,* i.e., the kind of obligation one incurs in virtue of voluntarily occupying an institutionally specified social role (Hardimon, 1994). In the case of physicians, professional obligations such as the therapeutic obligation can perhaps be justified by reference to the important role healers play in a society (Hawkins, 2006). However, it is not clear what substantive professional obligations would follow from the role economists play in society. Given the heterogeneity of the economics profession, the prospects of creating a substantive professional code of ethics based on role obligations seems dim (cf. DeMartino, 2011). It is hard to imagine what substantive professional obligations a microeconomist constructing and exploring theoretical models could come to be bound by. And it would be puzzling to insist that the microeconomic theorist, working in the confines of her office, could have a *professional* obligation to provide anyone with access to socioeconomic resources and opportunities.

4.2 Natural Duties

None of the remarks above should be taken to suggest that non-physician researchers conducting RCTs (which includes development economists) are not bound by any principles of ethical research design. As Kukla (2007) has pointed out, an ethically problematic medical trial is not automatically rendered permissible by substituting physician-researchers with non-physicians.

Using this insight, Kukla has attempted to ground principles of ethical research design in general ethical requirements of justice and respect for persons, or natural duties more simply. Kukla introduces a "Minimum Standard" Principle (MSP) which is meant to extend the principle of standard care beyond the medical context. On MSP, "researchers should not run studies unless, *to the best of their knowledge*, every trial arm receives care that is at least as good as the local *de jure* standard of care" (Kukla, 2007, p. 178). Kukla is extending the notion of care to include more than just medical treatment, and by invoking the concept of a *de jure* standard of care, Kukla is suggesting that there is a level of socioeconomic resources and opportunities that research participants are entitled to *given* the cultural and material context in which the research takes place. Going forward, the concept of "standard of care" will take on this broader meaning.

Kukla is right that research participants are owed at least some minimum level of socioeconomic resources or opportunities; such a claim is consistent with most (if not all) plausible theories of distributive justice. But while Kukla's proposal is promising in that it bypasses the traditional grounding of research ethics on professional obligations, it rests on a controversial premise. MSP implies that researchers *qua* persons have duties to provide research participants with what they are owed; it would thereby forbid the routine use of the status quo as a control in experimental development economics. But as MacKay (2015) notes in response to Kukla, "most political philosophers claim that it is the responsibility of *institutions*—not individuals—to provide citizens with what they are owed, whether access to healthcare, income, opportunities etc." (MacKay, 2015, p. 8).[32] Of course, as both Kukla and MacKay acknowledge, researchers are ethically required to not interfere with the entitlements research participants

[32] Two notable exceptions are Cohen (1997) and Murphy (1998). Both defend the view that individuals have positive requirements of distributive justice.

presently have access to (if any), but this is uncontroversial and ultimately consistent with using the status quo as a control.

In Kukla's defense, however, it is worth pointing out that MSP is meant to apply in circumstances in which people are not receiving what they are owed precisely because institutions do not comply with the demands of justice or simply do not exist. What MSP can perhaps be taken to suggest is that in these circumstances, researchers have an ethical requirement to provide research participants with care that is at least as good as the local *de jure* standard of care—whatever it may be. This is still a controversial position to maintain as there is no settled account of what positive ethical requirements persons have in non-ideal circumstances. One possible way to ground MSP's demands on researchers would be by appeal to the second component of Rawls's natural duty of justice, which requires one to further just arrangements when they do not yet exist (Rawls, 1971, p. 115).[33] Political philosophers often appeal to this duty in non-ideal circumstances since it can be "invoked as a 'normative bridge' translating the demands of justice applying to complex institutional agents—such as the state—into responsibilities falling on individuals" (Valentini, 2021, p. 46). Such a strategy seems especially fitting in the case of experimental development economics since Rawls holds, quite plausibly, that a certain level of economic development is a necessary precondition for the emergence of just institutions (Rawls, 1999, pp. 106-111).[34] However, the exact demands of this duty are

[33] The first component of the natural duty of justice requires us to comply with the demands of just institutions.
[34] It is worth highlighting that the natural duty of justice may require development economists *abstain* from interfering with the status quo all together. This is because, if Deaton (2013) is correct, interventions from outsiders crowd out the incentives of governments to tax their citizens and become responsive to their needs, which thereby hampers the development of local institutions. But as Ravallion (2014) argues in response to Deaton, there are also circumstances where outside interference will help local institutions develop. See Temkin (2022) for recent philosophical commentary on Deaton's Worry and Picchio (2023) for evaluation of Temkin's commentary on Deaton's Worry.

unclear given societal variations in historical circumstances and, more importantly, lack of full compliance with the natural duty of justice on behalf of others.

Addressing the concern above is the chief task of non-ideal theory understood as *partial compliance theory*. Three types of answers are typically offered in response to the question of what positive requirements persons have under conditions of partial compliance: do your fair share and nothing more; do more than your fair share by picking up the slack of others; and do less than your fair share so long as you reasonably expect non-compliance from others (Miller, 2011). While I do not mean to suggest that development economists are moral saints, it does not seem like a stretch to maintain that, even by running RCTs in which the status quo is used as a control, development economists are doing more than their fair share of the natural duty to bring about just arrangements. Not only are they directing socioeconomic resources and opportunities to low-income countries, but they are also devoting their talents and abilities—which could be put to other uses—towards figuring out how to effectively reduce global poverty, and in turn, bring about more just arrangements. Consequently, it is hard to maintain that development economists would owe the control arm more than the status quo even granting that there is a natural duty to promote just arrangements when they do not yet exist.

Another possible basis for MSP is the natural duty of beneficence. Valentini (2021) has recently argued that it is often more appropriate to appeal to this duty instead of the natural duty of justice in conditions of widespread non-compliance. Recall, however, that the natural duty of beneficence is not correlative with any specific person's entitlements. Consequently, appealing to the natural duty of beneficence would not settle the question of what development economists owe their research participants. Still, one may ask whether the natural duty of beneficence could establish that development economists are required to provide more than the status quo level of

resources and opportunities to research participants. However, there is also a strong case to be made that the demands imposed by the natural duty of beneficence are also subject to considerations of partial compliance (Murphy, 2000). Consequently, the same points regarding the natural duty of justice apply *mutatis mutandis*: development economists are quite plausibly already doing more than their fair share of the natural duty of beneficence by conducting an RCT in which the status quo is used as a control.[35]

Despite these remarks, there is no denying that MSP is motivated by a legitimate ethical concern. There are perhaps some circumstances so unjust, or so dire, that it would be wrong to be complicit with even if there is something potentially valuable to learn. Further on, I revisit the idea of a minimal standard of care and provide some basis for the concern motivating it by appealing to the natural duty of rescue.

4.3 Domestic Institutional Obligations

MacKay (2018, 2020) offers a much less contentious foundation for a principle of standard care in his work on *policy* RCTs, i.e., RCTs on prospective socioeconomic interventions authorized or conducted by local, state, and federal governments or their respective agencies. While some development RCTs are also policy RCTs, not all development RCTs are—a point I return to below. MacKay's approach rests on a key premise: governments are the kind of institutions that have obligations to pursue justice-related outcomes for their citizens. Consequently, governments are required to provide citizens with access to healthcare, education, and finance so that they

[35] One may question whether development economists conducting experiments really are doing more than their fair share once we factor in the personal material benefits accrued through conducting their research. The development economists mentioned at the outset did, after all, win a Nobel prize, which comes attached with considerable material benefits. The first thing to point out is that not all development economists conducting experiments enjoy the material benefits of the three Nobel prize winners. Still, publishing articles based on experimental results leads to material gains via promotion and advancement in the economics discipline. Should the material gains from each experiment enter our assessment of what a researcher's fair share is? Absolutely, but even if the exact calculations could reliably be made, it is not clear if this consideration will pose a significant challenge to the argument above.

may achieve these target outcomes. An analogue of the principle of standard care can as a result be grounded in the obligations that governments have to residents within their territory. MacKay argues that for any justice-related outcome a government has a duty to realize, governments possess an obligation to implement the policy that is "(1) evidence-based, (2) consistent with people's rights, and (3) consistent with the realization of other target outcomes (MacKay, 2020, p. 323). MacKay calls this the Best proven, Morally and Practically attainable and sustainable (BPA) policy. A policy counts as a BPA policy if and only if "(1) it is consistent with residents' rights and (2) it can be implemented for an appropriate period of time given a just system of resource procurement and allocation" (MacKay, 2020, p. 324).

The BPA policy can be seen as an analogue of the local *de jure* standard of care for policy RCTs (MacKay, 2018, p. 62, fn. 28). This raises the question: Do development economists have an obligation to ensure that the control arm receives the level of resources and opportunities specified by the BPA policy? One possible way of establishing this result is by appealing to the *institutional obligations* development economists incur in virtue of their status as *government-authorized investigators* (GAIs) (MacKay, 2018, p. 64). The rationale behind this approach is that GAIs have an institutional obligation to act only in ways that a government *may* permissibly authorize. This is crucially different from saying that GAIs have institutional obligations to act only in ways that a government *actually* authorizes. The latter has clear counterexamples (e.g., Nazi medical experiments). Yet the former offers an institutional answer to the question of what development economists owe their research participants. Since a government may not permissibly deny its residents what is owed to them for the sake of potentially learning something valuable, it follows that GAIs are not allowed to do so either. It

would further follow that GAIs are required to provide research participants with the level of socioeconomic resources and opportunities specified by the BPA policy.

Unfortunately, this proposal has serious difficulties in the context of experimental research in development economics. Most obviously, not all development RCTs are conducted or sponsored by the government of a host country. Many development RCTs are conducted or sponsored by NGOs and private for-profit businesses (e.g., microfinance banks). This raises a host separate ethical and political issues which I address in chapters 4 and 5. But the main point is that because NGOs and for-profit businesses are private entities, it is not clear what distributive justice-based institutional obligations their affiliates may come to be bound by.[36] NGOs and private businesses may, of course, have duties of justice and beneficence. But the same points about individuals discussed above should apply here as well. Regardless, it is still worth considering whether development economists working as GAIs have institutional obligations to provide the control arm with more than the status quo.

The notion of a BPA policy will undoubtedly strike development economists as extravagant due to the non-ideal circumstances in which their research takes place. This complaint is not entirely misguided; MacKay's framework is perhaps best suited for policy RCTs run or sponsored by the governments of high-income countries. While it would be inappropriate to characterize the circumstances of present high-income countries as *ideal*, the circumstances are certainly *less* non-ideal in that historical circumstances are favorable. As a result, some institutions and policies meant to promote justice-related outcomes are in place. Further, there is a case to be made that the policies in high-income countries *can* meet MacKay's criteria for BPA policies. However, when institutions and policies are in place in low-income

[36] In the context of international medical research, MacKay (2015) has argued that unlike government research agencies, private businesses do not have institutional obligations.

countries, it is hard to maintain that they fully allot residents what they are entitled to even by local standards. One could argue that since the notion of BPA policy includes feasibility constraints, it factors in the non-ideal circumstances of the developing world (resource constraints, state capacity, etc.) into the definition of best proven and *attainable* policy. However, granting this premise would simply show that the routine use of the status quo as a control is permissible on MacKay's framework. But this is not a position I wish to defend to establish my main conclusion.

The obvious problem with adopting the BPA policy as the standard of care in development economics is that it would effectively render all past and future policy RCTs conducted by development economists as unethical. It is worth emphasizing that part of what makes the developing world non-ideal, and what motivates the need for the kind of research done by development economists, is that governments (and perhaps other global institutions—see below) are not providing their residents with the level of socioeconomic resources and opportunities they are fully entitled to. Again, one does not need to identify BPA policies to know this. Consider, for example, the routine mismanagement of public resources in low-income countries due to corruption.[37] Clearly, such widespread public corruption is not consistent with residents receiving what they are fully owed as a matter of distributive justice. For this reason, it is not tenable to suggest that the BPA policy can be the standard of care in experimental development economics. Doing so would impose institutional obligations to provide BPA policy-levels of socioeconomic resources and opportunities, which, as a result, would mean that development economists would be too overburdened with picking up the slack of institutional non-compliers to ever conduct an RCT permissibly.

[37] See Svensson (2005) for discussion of the strong relationship between country income and corruption as well as a general overview of the economic literature on corruption.

MacKay acknowledges the point above and offers one way of reconciling his account of institutional obligations with non-ideal circumstances. When conducting a policy RCT in which participants are exposed to policies inferior to the BPA policy (which I am suggesting is routine practice in the case of experimental development economics), MacKay suggests that "GAIs commit a *pro tanto* wrong against participants, but this wrong is outweighed by competing considerations, namely, the value of the research" (MacKay, 2018, p. 65). While I am sympathetic to MacKay's proposal for non-ideal circumstances, it is worth noting that a similar strategy is also potentially available to *pragmatic cosmopolitans* (Emmanuel, 2012) who would reject a local *de jure* interpretation of the standard of care in experimental development economics. These cosmopolitans could insist on a global interpretation of the BPA policy and could similarly introduce a defeasibility condition to their prospective view. In other words, they too could maintain that in non-ideal circumstances a development RCT can go forward if it is sufficiently valuable despite research participants being denied what they are owed as a matter of distributive justice.

The proposal suggested above offers cosmopolitans a way of reconciling their philosophical commitments with pragmatic considerations. In the next section I cast doubt on whether, currently, one could establish that there are institutional obligations to provide research participants with the level of socioeconomic resources and opportunities specified by a *global* BPA policy—whatever it may be.

4.4 Global Institutional Obligations

Could appealing to global rather than domestic institutional obligations establish that the control arm be provided more than the status quo? Here, the thought would be that intergovernmental organizations (IGOs) such as the United Nations (U.N.) or the World Bank and International

Monetary Fund (IMF) have duties of distributive justice, and affiliation with such institutions would confer positive obligations on development economists. The most direct way a development economist could come to be affiliated with an IGO is through the IGO funding their research, but I grant that there are a variety of indirect ways that one could come to have an institutional affiliation with an IGO and very little hinges on this point.

As discussed above, one can maintain that development economists associated with, for example, the U.N. commit a *pro tanto* wrong against research participants by exposing them to the status quo when the development RCT in question is sufficiently valuable. But there is a more fundamental problem with this proposal. Even if we could identify the correct principles of global distributive justice and the BPA policies they imply, there are currently no global institutions capable of acting on them. This is because, unlike states, no currently existing global institutions have the requisite coercive capacities that give rise to demands of distributive justice (Blake, 2001; Risse, 2005; Nagel, 2005). Consequently, it is difficult to make the case that development economists—at this moment in time—could incur global institutional obligations of distributive justice. Unlike with MacKay's proposal, one could not thereby maintain that IGO-authorized development economists commit a *pro tanto* wrong against research participants—at least on distributive justice grounds (another basis for a *pro tanto* wrong is discussed below). None of this is to deny that IGOs have some coercive capacities.[38] However, to borrow a phrase from Stiglitz (2003, p. 22), the situation is currently one of "global governance without global government."

While IGOs may be unlike states in their coercive capacities, and in turn their obligations of distributive justice, there may be considerations of corrective justice that can provide the basis

[38] See, for example, Stiglitz (2003) for discussion of how IGOs such as the International Monetary Fund (IMF) and World Trade Organization (WTO) constrain the sovereignty of developing countries.

for the global institutional obligations of researchers. There is a case to be made that the IGOs that (partially) comprise the global political and economic order cause global poverty and thereby harm the residents of developing countries.[39] One could thereby argue that development economists affiliated with IGOs have obligations of rectification which make use of the status quo as a control impermissible.

Perhaps the most salient episode on which to base claims of corrective justice on is the era of structural adjustment loans (SALs) from roughly 1980 to 1999. During this period, the IMF provided loans to developing countries facing economic turmoil on the condition that they make crucial "structural" changes to their economic policies. Among the most deplored structural changes the IMF demanded were fiscal austerity measures. Cuts to social spending on health and education due to these measures could very well have harmed residents of low-income countries, as numerous scholars, advocates, and activists have claimed. Since these global institutions are (at least partly) responsible for the current unjust status quo, one can introduce the claim that development economists affiliated with IGOs like the World Bank and IMF are doing something wrong in using the status quo as a control. This is because the members of the control group are in fact owed something from these institutions, namely, rectification for the harms due to the structural adjustment era in their country. The problem, however, is establishing the appropriate counterfactual baseline for these harm-claims. In the case of SALs, the task is plagued by concerns about selection bias since the countries taking out these loans were already on a negative economic trajectory. As Easterly (2009, p. 423) notes, complaints about structural adjustment era are "based on correlations between SALs and outcomes that are the equivalent of the negative correlation between admission to an emergency

[39] See Pogge (2002) for extensive discussion of this consideration and its implications for debates about global justice.

room and a person's health, with the implication that the emergency room is bad for your health."

Setting causal identification problems aside, it's not implausible to think that global institutions have unjustly harmed the global poor and thus have obligations of rectification. However, there remains the possibility that using the status quo as a control in a development RCT will generate knowledge that will allow global institutions to discharge their obligations of rectification more efficiently and effectively. This consideration should not be overlooked since efficiency and effectiveness are at least among the most important factors institutions should consider when discharging their duties of corrective justice. And moreover, this consideration is also consistent with the possibility that IGO-authorized researchers commit a *pro tanto* wrong when they use the status quo as a control.

4.5 Taking Stock

I have taken steps to establish the difficulties in maintaining that the standard of care in experimental development economics is anything above the local *de facto* level of socioeconomic resources and opportunities. My argument has been negative; I have ruled out natural duties, professional obligations, and institutional obligations as providing the basis for a *de jure* conception of the standard of care in experimental development economics. I believe this argument justifies the use of the status quo as a control in many of the research contexts development economists have, to date, been interested in. However, in the next section I address the concern that my account is overly permissive in that it would allow experiments that are highly objectionable to go forward.

5. The Natural Duty of Rescue

In this penultimate section, I try to provide some basis for the idea of a minimum standard of care by appealing to the natural duty of rescue.[40] Following Miller (2020), I take the natural duty of rescue to be correlative with a would-be victim's entitlement to be rescued. The natural duty of rescue may require the distribution of live-saving resources in some scenarios; however, it is not strictly speaking a duty of distributive justice as I understand it here. My goal in introducing the natural duty of rescue is only to offer an uncontroversial basis for complaints about development RCTs.[41] As such, I distance myself from attempts by some commentators, most notably Singer (1972), to give the natural duty of rescue a wider, more demanding scope.

On what I am calling an uncontroversial understanding of the natural duty of rescue, physical proximity to an emergency is (contra Singer) necessary for the duty to be triggered.[42] The demands imposed by the natural duty of rescue cannot also be so burdensome that they require innumerable personal sacrifices on behalf of persons. To use a well-worn example, if while walking to work I notice a child drowning in a shallow pond that I can save without significant risk to myself, the natural duty of rescue requires that I provide this stranger with my assistance, and the stranger in question is presumably entitled to my assistance as well. Someone

---

[40] This approach is similar in some ways to Hawkins' (2006) reliance on Good Samaritan obligations in her analysis of the use of PCTs in international medical research. Besides a difference in terminology, I am using the duty of rescue to establish a minimum, whereas Hawkins combines the duty of rescue with the duties of distress avoidance and gratitude to argue that the use of placebo-controls is unethical in some circumstances. For a criticism of Hawkins' approach, see MacKay and Rulli (2017).

[41] Rulli and Millum (2016) are correct to point out that appeals to the natural duty of rescue in the research ethics literature are underdeveloped, and my proposal is no exception. These commentators distinguish between the *duty of easy rescue* and *the rule of rescue* and proceed to criticize appeals to both. The former requires that one rescue others when it involves minimal cost to oneself. The latter requires that the would-be victim be identifiable. The uncontroversial natural duty of rescue I am proposing can be seen as a synthesis of the two, which Rulli and Millum do not consider.

[42] Miller (2020) has most recently defended this position. Also see Kamm (2000).

on the other side of town, or in another country all together, does not have a duty to rescue this stranger because they lack physical proximity to the emergency.

The idea here is that by conducting their research in low-income countries, development economists place themselves in a special position to act on the natural duty of rescue. Their non-experimental colleagues working in the confines of their offices are not in this special position, at least not frequently. What this suggest is that by designing and conducting an RCT in which the control group is left in an emergency scenario, development economists would be failing to act on their natural duty of rescue, and the experiment would therefore be ethically wrong to carry out if there are sufficient life-saving resources to go around. One can imagine, for example, a development RCT designed to take place during a famine or some other type of emergency. In such a case, it would be wrong for development economists to not do everything in their power to address the severity of the situation because, as persons, they are duty-bound to provide aid in such emergency situations.[43] Similar considerations apply to situations where an emergency develops over the course of an experiment. In such a case, development economists would be duty-bound to halt the experiment and direct all their resources and attention to not only helping participants in all trial arms, but also members of the community from which research participants are drawn.

One noteworthy objection to what I am proposing is that the developing world can be characterized as a "constant emergency situation" (Hawkins, 2006). In response, I should reemphasize that my goal in introducing the natural duty of rescue as giving rise to a minimum

[43] Like Hawkins (2006), I do not deny that the natural duty of rescue is defeasible. One can conceive of an emergency scenario where the importance of evaluating a life-saving intervention is necessary because it could help persons in similar emergency situations in the future. In such a case, perhaps the natural duty of rescue is outweighed by other ethical considerations. However, trying to figure out how to effectively act on the natural duty of rescue may ultimately be consistent with complying with the duty. The details of this proposal need to be carefully worked out, but I do acknowledge that if the emergency situation stretches across time or is likely repeat itself in different places, then it may be permissible conduct a development RCT in such extreme circumstances.

standard of care is only to provide an uncontroversial philosophical basis for ethical complaints about development RCTs. Adopting a broad understanding of an emergency would prevent us from ever articulating such complaints since all development RCTs would effectively be considered unethical. While I cannot provide a satisfactory account of what constitutes an emergency here, I suggest instead that whether some prospective development RCT falls below this minimum standard, or whether development economists have failed in some instances to act on their natural duty of rescue, are matters that should be settled on a case-by-case basis. To reiterate, the natural duty of rescue can provide an uncontroversial philosophical basis for such case-by-case assessments.

6. Conclusion

To close, I should emphasize that I have only argued that the routine practice of using the status quo as a control is ethically justifiable. If there *is* something deeply unethical about the routine use of the status quo as a control in development RCTs, the wrong-making feature does not stem from considerations of distributive justice.

Note that there are still other ethical requirements that need to be met (and formulated) for development RCTs to go forward. In addition to (but not limited to) securing genuine informed consent, avoiding exploitation, demonstrating social value, and selecting research participants fairly, development economists may also be required to meet an analogue of the *principle of clinical equipoise* (Freedman, 1987, 1990), which requires the medical community to be in a collective state of uncertainty with respect to the relative therapeutic benefits of each arm of a prospective medical trial. As with the principle of standard care, the principle of clinical equipoise has traditionally been grounded in the therapeutic obligations of physicians. Further, there is virtually no uncertainty with respect to the therapeutic value of some of interventions

development economists evaluate (e.g., deworming drugs). This suggests that an extension of this well-known ethical requirement in medical research will not be so straightforwardly carried over to development economics but could possibly come to play an important role in the ethics of development RCTs as some commentators have suggested. In the following chapter, I explore whether this is the case.

Chapter 3: What Role Should Equipoise Play in Experimental Development Economics?

1. Introduction

In the previous chapter, I argued that development economists may permissibly use an unjust status quo as a control for a development RCT. The basis for this conclusion is that development economists do not have an individual ethical requirement to provide research participants with the socioeconomic resources and opportunities they are entitled to. In this chapter, I address a related concern: the fact that development economists routinely evaluate interventions known to be beneficial to their individual recipients.

At first glance, it is not clear why anyone would think that evaluating interventions known to benefit individual recipients poses any sort of ethical challenge. Before explaining why this practice requires ethical scrutiny, consider, once again, the study conducted by Cohen and Dupas (2010) on the effects of subsidies on preventative healthcare products in low-income countries. In the experiment, there was no mystery surrounding the therapeutic benefits of ITNs. It is a clear example of development economists evaluating an intervention known to benefit individual recipients. Yet Ziliak and Teather-Posadas (2016) harp on this very practice to criticize experimental development economics on ethical grounds. They argue that development economists have borrowed well-established experimental methods from the biomedical sciences and ignored the stringent ethical requirements that accompany their use. Specifically, they indict development economists for ignoring a well-known concept from medical research ethics: *equipoise*. In this chapter, I take steps towards assessing the accuracy of this complaint by focusing on the role that equipoise should play in experimental development economics.

In its most general guise, equipoise is a state of uncertainty about the relative merits of a set of interventions (London, 2020). Many physicians and research ethicists believe that some form of equipoise is a necessary component of ethical medical trial design. The most well-known formulation of this requirement (to be discussed in more detail below) is the *principle of clinical equipoise* (Freedman, 1987, 1990). The principle states that, for a clinical trial to proceed in an ethically defensible manner, the expert medical community needs to be in a state of collective uncertainty with respect to the relative therapeutic benefits of each trial arm.

Though the concept of equipoise has its roots in medical research, Baele (2013)—one of the first to extensively comment on the ethics of development RCTs— has suggested that equipoise could play an important role in experimental development economics. This suggestion has been recently echoed by commentators such as Abramowicz and Szafarz (2020).[44] Yet, for reasons that I make clear below, there are considerable difficulties with importing the concept of equipoise to experimental development economics. Nevertheless, I identify a form of equipoise that can play a role in experimental development economics. I refer to this form of equipoise as *social equipoise* following Petticrew et al.'s (2013) call for such a concept to be developed. However, as I ultimately argue, any plausible version of a *principle of social equipoise* will do little to help advance discussions of ethics in experimental development economics. At best, considerations of social equipoise can provide a principled basis for ethical complaints about the *social value* of an experiment. Considerations of social equipoise can also show how methodological debates concerning the scientific value of development RCTs have an ethical dimension as well.[45]

---

[44] Deaton (2010) and Ravallion (2020) also briefly discuss equipoise in passing.

[45] See Banerjee and Duflo (2009), Barrett and Carter (2010), Basu (2014), Cartwright (2010), Cartwright and Hardie (2012), Deaton (2010), Deaton and Cartwright (2018), Harrison (2011), Heckman (2020), Kremer and Holla (2009),

The rest of this chapter is structured as follows. Section 2 provides a brief overview of the history of equipoise and the ethical problem it is meant to resolve in medical research. Section 3 highlights the difficulties with importing the concept of equipoise to experimental development economics. Section 4 rejects equipoise as a tool for risk-benefit assessment in experimental development economics. Section 5 incorporates an additional methodological complication with importing equipoise to experimental development economics. Section 6 discusses MacKay's (2018, 2020) recent work on the ethics of public policy research and its relevance to experimental development economics Section 7 introduces the link between equipoise and value and culminates with a formulation of the principle of social equipoise. Section 8 provides a philosophical grounding for the principle of social equipoise. Section 9 applies the principle of social equipoise to Cohen and Dupas's (2010) experiment. Section 10 concludes.

2. A Brief History of Equipoise

Recall that in contexts entirely divorced from medical research, physicians have ethical (and legal) obligations to their patients. These obligations can be seen as a contractual role obligation, i.e., an obligation someone has in virtue of voluntarily occupying an institutionally specified social role (Hardimon, 1994). One of the most well-recognized ethical obligations physicians incur is the *therapeutic obligation*. Consider an alternative formulation than the one presented in the previous chapter:

> *Therapeutic Obligation*: A physician should not recommend for a patient therapy such that, given present medical knowledge, the hypothesis that the particular therapy is inferior to some other therapy is more probable than the opposite hypothesis (Marquis, 1983, p. 42).[46]

Ravallion (2020), Rodrik (2009). Also see Boone and Johnson (2009, pp. 61-64) for discussion of some methodological differences between clinical trials and development RCTs.

[46] Some commentators will refer to this ethical requirement as the duty of personal care.

The physician's therapeutic obligation is key to understanding the initial motivation behind the development of equipoise as an ethical requirement for clinical trials. As the use of RCTs in medical research became more widespread, there was growing concern that the new experimental methodology was in tension with the therapeutic obligations of physicians. Throughout, I refer to this (alleged) tension as the *RCT Dilemma:*

> *RCT Dilemma*: If a physician *P* has reason to believe that new Therapy A is better than another Therapy B, then *P* cannot permissibly enroll or advise patients to enroll in a clinical trial of A versus B because, ethically, *P* is obligated to recommend A to each new patient with a need for one of these therapies.[47]

Early commentators worried that physicians conducting clinical trials routinely face this dilemma because physicians often develop "treatment preferences" (Schafer, 1982) based on incomplete scientific evidence, personal experience, and mere hunches. If a physician preferring Therapy A to Therapy B enrolls patients in a trial comparing A to B, the physician would then (supposedly) violate her therapeutic obligation and the clinical trial would therefore be unethical.

The concept of equipoise was first introduced by Fried (1974) as a way of resolving the RCT Dilemma.[48] Fried reasoned that the only way a physician could permissibly enroll patients into a clinical trial without violating her therapeutic obligation was if she was genuinely uncertain about the relative therapeutic benefits of each trial arm. The distinguishing feature of Fried's version of equipoise—or *individual* equipoise—is that the uncertainty necessary for permissible randomization had to be in the mind of the individual physician. Because Fried did not elaborate on what he meant by genuine uncertainty, commentators interpreted him as suggesting that physicians enrolling patients in a clinical trial need to maintain that the

[47] The first commentators to make note of this dilemma appear to be Shaw and Chalmers (1970, p. 487), from which my formulation of the RCT Dilemma is drawn.
[48] Equipoise was not the only solution proposed. See Marquis (1983), Gifford (1986), and Freedman (1987) for discussions of other proposed solutions to the RCT Dilemma.

probability of Therapy A being superior to Therapy B is exactly 50 percent (Chalmers, 1978).[49] This made equipoise seem like an unlikely state to be in given the propensity for physicians to form treatment preferences. Even if a physician found herself in such a state due to the lack of any available evidence, critics complained that equipoise was fragile and would easily be disturbed as soon as a clinical trial commenced and evidence favoring one trial arm started accruing (Marquis, 1983; Gifford, 1986; Freedman, 1987). This would be a problem because clinical trials would need to be prematurely stopped before any statistically significant results could be established.[50]

The difficulties with Fried's version of equipoise led Freedman (1987) to propose clinical equipoise as a necessary requirement of ethical clinical trial design. Per London (2020), Freedman's clinical equipoise is meant to solve the same ethical problem as individual equipoise (cf. Miller and Weijer, 2003). However, clinical equipoise differs from individual equipoise in that it requires uncertainty to obtain in the expert medical community rather than the mind of the individual physician. Interestingly, one of the central motivations for this shift is an appeal to the social dimensions of scientific knowledge (Longino, 1990). Freedman (1987, p. 144) writes that medicine "is social rather than individual in nature" and its advancement relies on "progressive consensus within the medical and research communities." Crucially, this means that for clinical equipoise to obtain, the individual members of the relevant medical community do *not* each need to each be in a state of equipoise (though this is certainly a possibility). Instead, for there to be genuine uncertainty about the relative therapeutic benefits of Therapy A to Therapy B, there can

[49] Miller and Weijer (2003) provide a more charitable interpretation of Fried.
[50] One practical solution to this problem that has now become standard practice in medical research is to conceal interim analyses from investigators, research participants, and funders, and only allow an independent Data Monitoring Committee having access to interim results. This Data Monitoring Committee will be contractually obligated to withhold interim results unless certain statistical thresholds are met.

exist "an honest, professional disagreement among expert clinicians about the preferred treatment" (Freedman, 1987, p. 144). Because establishing clinical equipoise involves disagreement rather than indifference, it allows physicians to have treatment preferences and allows for a clinical trial to proceed until sufficient evidence is accumulated to resolve uncertainty in the expert medical community.

Having specified the relevant sense of uncertainty, we can turn to Freedman's influential formulation of the principle of clinical equipoise:

> *Principle of Clinical Equipoise*: at the start of a trial, (1) there must be a state of clinical equipoise regarding the merits of the regimens to be tested, and (2) the trial must be designed in such a way as to make it reasonable to expect that, if it is successfully concluded, clinical equipoise will be disturbed (Freedman, 1987, p. 144; numbers added).

As a preview of what is to come below, note that I have separated the principle of clinical equipoise into two components. This will be important when it comes to discussing one of the difficulties with importing clinical equipoise to experimental development economics. While there is no lack of professional disagreement among social scientists, social science is unlike biomedical science in that there is often a lack of consensus on how such disagreements can be resolved. This is one aspect of research in the social sciences that has been ignored in discussions of extending equipoise beyond the clinical context.

3. Difficulties with Equipoise in Experimental Development Economics

Having outlined the motivation behind the origins of equipoise, it is time to turn to the first and obvious difficulty with importing the most familiar version of equipoise to experimental development economics: any straightforward application of the principle of clinical equipoise to experimental development economics would render the great majority of development RCTs unethical if the principle is considered a necessary condition for permissible experimentation. As

noted at the outset, Ziliak and Teather-Posadas (2016) have harped on this very point to criticize development RCTs. As this section reveals, this worry is misguided.

Unlike with a new treatment for a serious medical condition, it is true that there is not much mystery surrounding the therapeutic benefits of an insecticide-treated bed net (e.g., Cohen and Dupas, 2010) or a deworming tablet (e.g., Miguel and Kremer, 2004). The point generalizes to interventions that are not healthcare related such as cash transfers and access to other valuable resources and opportunities. We can claim these things are typically good for people without appealing to a substantive theory of well-being, much less a scientific experiment (Hausman, 2011). If this is right, then to abide by a direct analogue of the principle of clinical equipoise, development economists would have to maintain—with a straight face—that there is no way of knowing whether some obviously beneficial intervention really is—in fact—welfare enhancing without first conducting an RCT.[51] This would undeniably put development economists in an extremely awkward position.[52]

To avoid unnecessary confusion, it is important that development economists conducting RCTs make clear that the goal of an experiment is not to determine that an intervention provides benefits to individual recipients (though this is a foreseeable consequence). Rather, for a large class of development RCTs, the goal of an experiment is to determine *how* effective an intervention is *for* achieving a specific socioeconomic outcome *in* a population of interest. All outcomes of interest may ultimately be grounded in a concern with human welfare, but this is not the same as saying that "we just do not know" whether giving someone a valuable resource is

[51] It is important to acknowledge that appealing to intuition or common sense is not always a reliable way of knowing whether some intervention is beneficial. There are plenty examples throughout the history of medicine of interventions thought to be "obviously beneficial" that turned out to be harmful. But this point can be easily abused and used to reach absurd conclusions, e.g., an RCT on the effectiveness of parachutes.

[52] Abramowicz and Szafarz (2020, p. 287) document at least one instance of an economist willing to put himself in this awkward position to justify an RCT.

good for them. One may still wonder why the goal of an experiment stated above cannot serve as the basis for a principle of equipoise in experimental development economics. For example, on London's (2001) account of equipoise in international human-subjects research, clinical equipoise is ultimately concerned with the *net* therapeutic benefit of a set of interventions *relativized* to a particular health problem *in* a specific treatment setting. Kukla's (2007) principle of equipoise, which I discuss further on, incorporates this kind of consideration in a way that is relevant to experimental development economics.

An objection to the argument above is that it only shows that development economists need to elevate their ethical standards to the level found in medical research. If development RCTs routinely violate clinical equipoise, then it is development economists who need to adapt—not a well-known ethical requirement. But this objection overlooks a more serious difficulty: economists do not stand in a professional relation to research participants and therefore do not have a positive obligation to provide "care" in virtue of the social role they occupy. Development economists still have a negative duty to not harm research participants (or anyone, for that matter), but as I explain below, this is orthogonal to considerations of equipoise. The more important point is this: because development economists lack the relevant role obligations, they do not face an analogue of the RCT Dilemma, which provides the motivation for the concept of equipoise in the first place.

It is worth highlighting that Miller and Brody (2003, 2007), two of the most notable critics of clinical equipoise, argue that the concept is irrelevant to the ethics of clinical research precisely because they see the RCT Dilemma as a pseudo-problem to begin with.[53] For these critics, the ends of medical practice differ from the ends of medical research. The former is

[53] Chiong (2006) rejects this argument but still argues against the importance of clinical equipoise. See Gifford (2000, 2007), Veatch (2002, 2007) and Miller and Joffe (2011) for additional critiques of clinical equipoise.

concerned with providing individualized treatment to patients and the latter with discovering generalizable medical knowledge for the benefit of society. As a result, physicians conducting medical research (arguably) do not stand in a therapeutic relationship to research participants. Hence, the ethical principles regulating medical practice, such as the requirement that physicians provide the best-known treatment for some medical condition, are inappropriate in the context of medical research.[54] While assessing this argument is beyond the scope of this chapter, it is important to acknowledge the possibility that equipoise has no role to play in experimental development economics, because it has no role to play in medical research either.

4. First, Do No Harm

Risk-benefit assessment is one of the most important ethical considerations in medical research. It is the primary means by which the injunction "do not harm" is operationalized in human-subjects research (National Commission for the Protection of Human Subjects of Biomedical and Behavioral Research, 1978). As mentioned above, development economists have a negative duty to not harm anyone. But it is also important to stress that strict adherence to the rule "do not harm" is untenable as documents such as the *Belmont Report* make clear: learning what is harmful or beneficial sometimes requires exposing research participants to some risks (National Commission for the Protection of Human Subjects of Biomedical and Behavioral Research, 1978). A key function of Institutional Review Boards (IRBs) is to make risk-benefit assessments and determine if enrollment in a clinical trial subjects a research participant to permissible amounts of risk.

[54] It is also worth noting that not all healthcare related research is conducted by physicians, which complicates any attempt at philosophically grounding research ethics on the role obligations of physicians. Intuitively speaking, an ethically dubious clinical trial is not rendered permissible by substituting physician-researchers with non-physician-researchers (Kukla, 2007).

Commentators such as Weijer (2000) and Weijer and Miller (2004) go a step further and argue that clinical equipoise is essential to their approach to risk-benefit assessment, which they call "component analysis". Risk-benefit assessment should undisputedly be *the* highest priority in determining the permissibility of a development RCT. Some may be led to think that some form of equipoise can play a role in making these important risk-benefit assessments so that the "do not harm" principle is met. It is important to establish that, when it comes to development RCTs, considerations of risk-benefit assessment are orthogonal to equipoise. As I argue below, equipoise has no role to play in making risk-benefit assessments in experimental development economics.

Component analysis rests on the contested distinction between therapeutic and non-therapeutic research.[55] It is hardly worth diving into the controversy and taking a stance on the issue since the distinction will prove irrelevant to experimental development economics. For present purposes, I will accept the distinction as advocates of component analysis suggest. Per Weijer (2000, p. 359), in therapeutic research there is therapeutic intent, i.e., the intention to provide some individualized therapeutic benefit to enrolled participants. In these contexts, researchers occupy the dual role of being both physicians and scientists. Non-therapeutic research involves evaluating interventions that are not intended to benefit individual participants. The purpose of such research may be to answer a scientific question, for example. Advocates of component analysis maintain that both kinds of research are permissible under different circumstances. This should already raise questions if one accepts—as advocates of clinical equipoise do—that medical research ethics should be governed by the norms of medical care (Miller and Brody, 2007, p. 157).

[55] See Levine (1988, pp. 8-10) for discussion and references.

On component analysis, there are two different standards of risk-benefit assessment. If a clinical trial falls under the category of therapeutic research, then the proposed study must pass the test of clinical equipoise to meet requirements of risk-benefit assessment. The reasoning should be familiar at this stage: the physician's therapeutic obligation requires such a test. Violations of clinical equipoise harm research participants by denying them what they are entitled to, viz., the best-known treatment. Clinical equipoise is supposed to ensure *ex ante* that research participants partaking in a study are not made worse-off than they would have been had they received standard treatment outside a clinical trial. It is for this reason that advocates of component analysis believe that a clinical trial conducted for therapeutic purposes would subject participants to an unfavorable risk-benefit ratio if clinical equipoise is not established. Advocates of component analysis also add an important addition to the familiar clinical equipoise requirement: IRBs are the entities that should determine whether a proposed study meets the principle of clinical equipoise. By ensuring clinical equipoise is met, IRBs comply with one of their key functions.

Advocates of component analysis also acknowledge that clinical trials may be conducted for different purposes, hence the distinction between therapeutic and non-therapeutic research. On component analysis, when research is conducted for non-therapeutic purposes, *clinical equipoise is completely disregarded*. Instead, on the non-therapeutic track of component analysis, IRBs are tasked with determining whether risks are minimal and reasonable: "The IRB must determine that nontherapeutic risks are, first, minimized consistent with sound scientific design and, second, reasonable in relation to the knowledge that may be gained from the study" (Weijer and Miller, 2004, p. 571). In other words, when research is conducted for non-

therapeutic purposes, advocates of component analysis acknowledge that equipoise is orthogonal to questions of risk-benefit assessment.

As already suggested in section 3, development economists do not have obligations analogous to those of physicians because they do not stand in a professional relationship to research participants. Without casting doubt on the merits of component analysis as a framework for risk-benefit assessment (see Miller, 2012), and without having to rehash earlier arguments, it is safe to conclude that equipoise provides an ill-suited basis for making risk-benefit assessments in experimental development economics. Even advocates of component analysis, who defend the importance of clinical equipoise, would likely acknowledge this point. And in fact, component analysis is not completely irrelevant to experimental development economics. Component analysis provides helpful criteria to follow when performing a risk-benefit assessment for a prospective development RCT: risks should be minimal and reasonable.

5. Methodological Considerations

Methodological considerations also pose an obstacle to making equipoise relevant to experimental development economics. As this section reveals, any form of equipoise in experimental development economics will have to incorporate something like the second component of Freedman's formulation of the principle of clinical equipoise. Otherwise, equipoise in experimental development economics will turn out to have no ethical bite.

RCTs are regarded as the "gold standard" when it comes to evaluating the effectiveness of some medical intervention (cf. Cartwright, 2007; Cartwright and Munro, 2010). While one clinical trial may not provide enough evidence to completely resolve all (reasonable) professional disagreement, there is agreement in the medical community that a series of clinical trials *is* often the best way to resolve the professional disagreement in question. This is another

way of saying that the results of a well-designed clinical trial can disrupt the medical community's state of equipoise with respect to the merits of some biomedical intervention. Freedman tacitly assumed this methodological consensus when devising the principle of clinical equipoise. The principle of clinical equipoise would thereby deem a poorly designed clinical trial as unethical since it would do little to disrupt the medical community's state of equipoise with respect to a disputed clinical intervention. Though this is intuitive enough, it is not clear how this important aspect of the principle of clinical equipoise derives from the therapeutic obligations of physicians—a point I return to further below.

Unlike in biomedical science, there is much less consensus among development economists about the scientific value of RCTs in their discipline (see anthologies such as Cohen and Easterly 2009a; Ogden, 2017; and Bédécarrats et al., 2020a). The absence of methodological consensus poses a serious difficulty for those who think that equipoise should play a role in experimental development economics: How can an RCT resolve any sort of professional disagreement if not all members of the relevant professional community believe that RCTs *can resolve* professional disagreements in the first place? It would be hasty to simply ignore this problem by dismissing this lack of methodological consensus in development economics as another instance of the reluctance economists have historically had towards experimentation (see Guala, 2005, pp. 2-3). As the discussion below reveals, the worries about the scientific value of RCTs in development economics are worth taking seriously.

For evidence of the disagreement about the scientific value of RCTs, consider the well-known methodological criticisms raised by Angus Deaton.[56] For present purposes, I will restrict my attention to Deaton's (2010) discussion of the problem of external validity (or

[56] See Ogden (2020) for a survey of the various critiques of RCTs in development economics and a discussion of how development economists conducting RCTs have changed their research practices to address these critiques.

generalizability).[57] While an RCT does—in principle—provide an unbiased estimate of the mean causal effect of an intervention, matters are not so simple in the social scientific context. Skeptics often argue that an RCT is a "black box" method of causal inference: "A treatment is administered, an outcome is observed, with no need for any understanding of what is going on in between and *why* a treatment produces its outcome" (Reiss, 2013, p. 205). The criticism in question stems from the fact that an RCT alone does not provide much in the way of a scientific explanation. Specifying the relevant sense of "scientific explanation" is important. When it comes to the social sciences, one leading contender is causal mechanistic explanation (Elster, 2015). The lack of explanatory evidence generated by an RCT significantly contributes to the problem of external validity since drawing inferences beyond the immediate test population typically requires knowledge of how the treatment produced the outcome of interest. Deaton (2010) notably argues that evidence from development RCTs will be of little use in furthering our understanding of economic development, and thereby making successful interventions, if efforts are not taken to also identify the casual mechanisms that generate the effects uncovered through field experimentation. More concerning is that policymakers can misapply findings from an RCT if they do not take the problems such as external validity seriously (Cartwright and Hardie, 2012).

It is worth acknowledging that external validity is also a problem for clinical trials, though the problem is often ameliorated by a corpus of background biomedical knowledge (Backmann, 2017). While a corpus of undisputed social scientific knowledge is harder to come

[57] Of course, internal validity matters as well. There is little point in worrying about external validity if internal validity is compromised. And as Deaton and Cartwright (2018) suggest, the distinction between internal and external validity may ultimately be misguided. I focus on external validity because it has been at the center of methodological debates surrounding development RCTs and for good reason: discovering generalizable knowledge is an oft-cited goal of scientific research.

by (Reiss, 2019), Deaton suggests that the problem of external validity in development RCTs can at least be addressed by closer engagement with economic theory:

> It is certainly not always obvious how to combine theory with experiments. Indeed, much of the interest in RCTs…comes from a deep skepticism of economic theory, and impatience with its ability to deliver structures that seem at all helpful in interpreting reality. Applied and theoretical economists seem to be further apart now than at any period in the last quarter century. Yet failure to reintegrate is hardly an option because without it there is no chance of long-term scientific progress or of maintaining and extending the results of experimentation. RCTs that are not theoretically guided are unlikely to have more than local validity, a warning that applies equally to nonexperimental work (Deaton, 2010, p. 450).

In addition to closer engagement with economic theory, Deaton also suggests that experiment design should incorporate previous empirical findings, particularly from behavioral economics. On Deaton's view, experiments guided by theory and previous empirical findings are more likely to generate the knowledge of causal mechanisms necessary to make successful socioeconomic interventions. Going forward, I will take these two considerations to be necessary (but not sufficient) components of sound experiment design.

While I have no illusions that methodological disputes in the social sciences can be easily settled with appeals to ethical considerations, there is an important takeaway when it comes to the role equipoise should play in experimental development economics: any form of equipoise in experimental development economics will have to incorporate something like the second component of the principle of clinical equipoise and require that an experiment be capable of resolving—or at least making steps towards resolving—some sort of collective uncertainty in the expert development community. This uncertainty can take the form of professional disagreement but could also stem from a lack of evidence on a research question of interest. Without such a component, any version of equipoise in experimental development economics will prove to be a toothless ethical requirement since, as already noted, there is an abundance of professional

disagreement among social scientists and often a lack of consensus on how to resolve such disagreements. For this reason, experimental design needs to consider both economic theory and previous empirical findings so that results from an experiment have some chance of disturbing equipoise in the expert development community. A successful account of the role equipoise should play in experimental development economics will give experimentalists an ethical reason to address the methodological concerns of critics.

6. Policy Equipoise

Commentators such as Baele (2013), Abramowicz and Szafarz (2020), and Asiedu et al. (2021) have all acknowledged some combination of the difficulties (highlighted above) with importing clinical equipoise to experimental development economics. Instead of rejecting equipoise all together, Asiedu et al. (2021) have recently moved the discussion forward by suggesting that MacKay's (2018, 2020) notion of *policy equipoise* should play a role in experimental development economics. While development RCTs may generate policy relevant evidence, the concept of policy equipoise is of limited relevance to experimental development economics, as I explain below.

MacKay has developed the concept of policy equipoise in response to the increasing use of RCTs for public policy research (henceforth, policy RCTs) by local, state, or federal governments and their respective research agencies. Recall that in MacKay's framework, the physician's therapeutic obligation is replaced by a government's duty to promote justice-related outcomes for its citizens. MacKay argues that for any justice-related outcome a government has a duty to realize, governments possess an obligation to implement the policy that is "(1) evidence-based, (2) consistent with people's rights, and (3) consistent with the realization of other target outcomes" (MacKay, 2020, p. 323). MacKay calls this policy the Best Proven,

morally and practically Attainable and sustainable (BPA) policy. A policy counts as a BPA policy if and only if "(1) it is consistent with residents' rights and (2) it can be implemented for an appropriate period of time given a just system of resource procurement and allocation" (MacKay, 2020, p. 324).

Because governments are required to implement BPA policies, governments and their respective research agencies are also required to treat citizens fairly by not subjecting anyone to policies known to be inferior to the BPA policy. This generates an analogue of the RCT Dilemma for government-authorized researchers conducting a policy RCT, and MacKay's notion of policy equipoise is *one* possible way to justify the fairness of randomization in policy research contexts. It is important to note that MacKay (2020) only introduces policy equipoise as a justification of the fairness of randomization and is *not* proposing "all things considered" criteria for when a policy experiment is ethically permissible.

Since much policy research often takes place under non-ideal conditions, i.e., conditions in which BPA policies are not in place, MacKay (2020) in fact proposes *two* principles of policy equipoise. I distinguish the two principles by whether they apply to ideal vs. non-ideal circumstances:

> *Principle of Policy Equipoise (Ideal)*: 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 (MacKay, 2020, p. 329).
>
> *Principle of Policy Equipoise (Non-Ideal)*: 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 (MacKay, 2020, p. 333).

Given the non-ideal circumstances that characterize the developing world, some will be inclined to think that the first policy equipoise principle will never apply to experimental development economics. However, as suggested in the previous chapter, one could argue that since the notion

of BPA policy includes feasibility constraints, it factors in the non-ideal circumstances of the developing world (resource constraints, state capacity, etc.) into the definition of best proven and *attainable* policy. But since MacKay provides his two principles in anticipation of difficulties with correctly interpreting the notion of a BPA policy, little hinges on this point.

Though both principles can be relevant to experimental development economics, they will have limited application for two reasons. First, MacKay's (2020, p. 321) considered view is that policy equipoise is a sufficient condition for permissible randomization—not a necessary condition as MacKay (2018) initially suggested. As MacKay (2020) and MacKay and Cohn (2023) discuss in more detail, there are circumstances where governments have good reasons to randomize without being in a state of policy equipoise. For example, a government may want to evaluate a BPA Policy against an intervention known to be superior in realizing target outcomes, but that is not practically attainable or sustainable. A government may want to do this for estimating cost-effectiveness or to gain a better understanding of the causal mechanism which generates the outcome of interest. MacKay's (2020) framework provides criteria for when randomization is permissible in scenarios like these without making any reference to policy equipoise.

Second, notice that not all development RCTs are policy RCTs according to MacKay's framework. Many development RCTs are sponsored or conducted by non-public entities such as privately funded non-governmental organizations (NGOs) and private for-profit corporations (PFPCs) such as microfinance banks (e.g., Banerjee et al., 2015) and utility companies (e.g., Devoto et al., 2012). Though private entities are not ethical freelancers, it is uncontroversial to maintain that private entities have different ethical duties than governments—a point MacKay acknowledges, and which serves as the basis for the two principles of policy equipoise.

None of this is to suggest that policy equipoise will never apply to experimental development economics in some policy research contexts. Policy equipoise may be used to justify the fairness of randomization in the case of "embedded experiments" (Drèze, 2022), where development economists work closely with a government agency to evaluate a policy. However, if some version of equipoise is to play a similar role in experimental development economics as it does in medical research, it should be (1) a necessary condition instead of a sufficient condition for permissible randomization; and (2) general enough to apply to a larger share of development RCTs—not just embedded experiments.

7. Social Equipoise

7.1 *Equipoise as Evidence of Social Value*

In a widely cited article on the ethical requirements of medical research, Emmanuel et al. (2000) identify *value* as the first requirement of ethical medical research. These commentators maintain, quite plausibly, that medical research is valuable when it is directed at an intervention that could lead to improvements in health or well-being. The second ethical requirement they list is scientific validity since, for medical research to be valuable, it needs to be conducted in a methodologically rigorous manner. What is noteworthy about Emmanuel et al.'s (2000) framework is that it does not rely on the therapeutic obligations of physicians. Instead, the primary ethical reason given for why medical research needs to be valuable and scientifically valid is that finite resources should be used responsibly.[58] What I take Emmanuel et al. (2000) to mean by "responsible use of finite resources" is that researchers should be sensitive to (ethical) opportunity costs. Considering that estimates for conducting a development RCT range

[58] Emmanuel et al. (2000) also cite avoidance of exploitation as a secondary reason. Exploitation is a notoriously difficult concept to analyze and on which to get clarity, so I avoid discussing it here. See Hawkins and Emmanuel (2008a) for an anthology devoted to the topic of exploitation in international medical research. The issue of exploitation is one that a complete account of the ethics of development RCTs will have to address.

anywhere between $500,000 to $1,500,000 (Bédécarrats et al., 2020b, p. 24), it should be uncontroversial to suggest that this ethical reason also carries over to research in experimental development economics.

Despite not relying on the therapeutic obligations of physicians, Emmanuel et al. (2000) do briefly suggest an entirely new role for the principle of clinical equipoise. Instead of providing a solution to the RCT Dilemma, clinical equipoise can play the dual role of indicating that research is scientifically valid and in turn valuable. As Emmanuel et al. (2000, p. 2704) argue, "If there exists a consensus about what is the better treatment, there is no null hypothesis, and the research is invalid." With respect to value, they add that "without clinical equipoise, research that compares therapies is unlikely to be of value because the research will not contribute to increasing knowledge about the best therapy…" (Emmanuel et al., 2000, p. 2704). Note here that these considerations make good on the promise of a clearer explanation for why a poorly designed clinical trial is unethical.

What I take Emmanuel et al. (2000) to be suggesting is that clinical equipoise plays an *evidentiary* role in medical research. I want to suggest that equipoise can play a similar role in experimental development economics: equipoise can play the role of indicating when an experiment has *social* (as opposed to medical) value. A full analysis of what constitutes social value is well-beyond the scope of this dissertation. However, I take it that it is unobjectionable to maintain that, like making improvements in health, promoting development-related outcomes is socially valuable. To borrow a concept from John Rawls (2005), an "overlapping consensus" of

plausible normative frameworks that generate this conclusion obviates the need to single out any one specific framework as providing adequate justification.[59]

Because development economists (and their funders) have obligations to use finite resources responsibly and can comply with such obligations by conducting (or sponsoring) socially valuable research (see Pierson and Millum, 2018), below I argue that the concept of *social equipoise* can play a role in ensuring that these obligations are complied with. Though social equipoise clearly takes inspiration from medical research ethics, the term "social equipoise" is not (to my knowledge) commonly used. As made clear below, social equipoise is not meant to resolve any sort of dilemma, nor does it require researchers to awkwardly maintain agnosticism with respect to the therapeutic value or individual welfare benefits of some interventions. Instead, social equipoise can play the role of indicating when a development RCT is socially valuable.

7.2 *Three Desiderata for a Principle of Social Equipoise*

The connection between equipoise and social value has, to date, been most thoroughly explored by Kukla (2007). As Kukla (2007, p. 173) makes clear, "My version of the principle of equipoise does not focus on equipoise with respect to the relative expected outcomes of trial arms, but rather on equipoise concerning the social value of the intervention being tested." Though Kukla does not use the term "social equipoise," their formulation of the principle of equipoise is a notable first attempt at articulating such a principle:

> *Kukla's Principle of Equipoise*: In order to begin or continue human subjects research, one must be in a state of equipoise with respect to whether or the extent to which the intervention being tested *should* be made accessible to the population that falls under the scope of the research (Kukla, 2007, p. 180).

[59] This is not to suggest that development economists should ignore insights from philosophical commentators (e.g., Sen, 1999), but rather that they do not need to patiently wait for philosophical controversies to be fully settled to engage in socially valuable research.

One important feature of Kukla's principle of equipoise is that it is *not* grounded in the therapeutic obligations of physicians nor any special obligations for that matter. Instead, Kukla's account is supposed to be grounded in general ethical requirements of justice and respect for persons. In the next section, I specifically argue that the duty to respect persons provides the normative basis for the principle of social equipoise. However, Kukla's principle is not without complications. Below, I respond to these complications by proposing three desiderata that any principle of social equipoise should meet if it is to apply to experimental development economics.

The first two desiderata concern the locus of uncertainty for social equipoise. Because Kukla seems to assume that researchers *qua* persons have duties to provide others with what they are owed as a matter of justice, there is undue concern with equipoise obtaining in the mind of the individual researcher, thus indicating that all researchers face an analogue of the RCT Dilemma with respect to their duties to provide others what they are due. There are two issues with this. The first issue, already discussed in the previous chapter, is that Kukla assumes that individuals are required to provide others with what they are owed as a matter of justice. The second issue is that Kukla's principle of equipoise suffers from the same fatal deficiency as individual equipoise, namely, that it is incredibly fragile. It is extremely unlikely that development economists (or any social scientists) are genuinely indifferent about whether some intervention is more socially valuable than the status quo. This is especially true if development economists are engaging in sound experiment design by consulting economic theory and previous empirical findings. Consequently, one can expect development economists to form judgments (akin to treatment preferences) about whether some socioeconomic intervention they

are investigating is socially valuable. If they thought otherwise, they would not bother designing an experiment in the first place.

The first desideratum to any principle of social equipoise is that the locus of uncertainty should not be the individual. The second desideratum is that the locus of uncertainty should be in the relevant expert community. As with clinical equipoise, there can be honest, professional disagreement in the expert development community. But this honest, professional disagreement should not be confined to just the experimental subset of the expert development community. To fully acknowledge the social dimensions of scientific inquiry as the principle of clinical equipoise does, development economists need to design their experiments in a way that can result in social equipoise being disturbed; this means attempting to design experiments in a way that could potentially convince non-experimental colleagues that the outcome of an RCT could have some bearing on a research question of interest. As already stressed above, experiment design needs to incorporate both economic theory and previous empirical findings so that the results from the experiment have some chance of disturbing social equipoise in the wider expert development community.[60]

The final desideratum I propose concerns the necessary and sufficient conditions for establishing that an experiment is socially valuable. Kukla claims that social equipoise requires that there be uncertainty about whether an intervention should be made accessible to a population of interest. Though this may be sufficient for establishing social value, it is not necessary. There are multiple ways in which a development RCT can be socially valuable. In addition to using RCTs for impact evaluation, Banerjee and Duflo (2009) also promote using RCTs to estimate

[60] Specifying the relevant criteria for membership into the expert development community is a challenge that I cannot address here. The challenge may prove to be an ineliminable difficulty with the concept of social equipoise and may ultimately provide a decisive reason for keeping equipoise out of experimental development economics all together. Notice, however, that this is a difficulty that the principle of clinical equipoise also faces (Veatch, 2007).

key parameters in economic models used for policymaking—a presumably socially valuable research activity. Taking cue from Deaton and Cartwright's (2018) discussion of the scientific value of RCTs, Morduch (2020) has recently emphasized that there are two distinct ways in which RCTs have been used by development economists. The first use of RCTs is for the familiar purpose of evaluating the impact of socioeconomic interventions. The use of "what works" RCTs (as Morduch calls them) for impact evaluation corresponds well with Kukla's criteria for establishing social value.[61] But Morduch also identifies a more promising use of RCTs in development economics; these are "how and why" RCTs. "How and why" RCTs can be used to answer research questions that are "exploratory, theory-driven, and motivated by the desire to understand economic possibilities and constraints" (Morduch, 2020, p. 109). As such, "how and why" RCTs can be used to gain valuable insights into the nature of economic contracts, behaviors, and institutions.

"How and why" RCTs often make use of short-term limited-scale programs and do not necessarily have bearing on whether the treatment being administered should be made accessible to a wider population. Still, what is learned through such experiments could be socially valuable. Using an RCT to confirm the prediction of a theory or provide a counterexample to a general theoretical proposition could have important downstream consequences (Deaton and Cartwright, 2018). If social equipoise is to be relevant to experimental development economics, the concept should not rule out a "how and why" RCT in which the treatment arm receives a highly beneficial intervention that cannot be made accessible to the population from which the sample is drawn. For example, consider an experiment that briefly provides the treatment arm with very

[61] It is also worth highlighting that "what works" RCTs are the development RCTs that have been the central focus of the methodological criticisms discussed earlier.

generous unconditional cash transfers that are not sustainable long term.[62] Such an experiment could still prove socially valuable and be permissible to conduct so long as additional ethical criteria are met.[63]

7.3 Formulation of the Principle of Social Equipoise

With these desiderata in mind, I propose the following principle as a necessary condition for ethical experiment design in development economics:

> *Principle of Social Equipoise:* At the start of an experiment, (1) there must be disagreement or doubt (i.e., equipoise) in the expert development community with respect to whether or the extent to which the intervention being tested has bearing on a research question of interest and (2) the experiment must be designed in such a way as to make it reasonable to expect that, if it is successfully concluded, something will be learned: there will be some effect on the extent of disagreement or doubt (i.e., equipoise will be disturbed).

Given the nature of social science, condition (2) is harder to obtain than in medical research. But insisting on a strict reading of condition (2) would also make social equipoise an impossible standard to meet—some flexibility is required. However, the appropriate degree of flexibility should not give development economists license to ignore previous empirical findings or engage in a-theoretical experiment design.

A reasonable complaint about the principle of social equipoise that I have formulated above is that it is not much of an ethical principle but rather a statement of sound methodological research practice. In fact, some readers will associate social equipoise with the concept of scientific relevance and will disregard social equipoise as trivial or superfluous. As warned at the outset, this is unavoidable. While the principle of social equipoise does little to advance the state of ethical discourse surrounding development RCTs, it at least provides a principled basis for

---

[62] See Morduch (2020) for analysis of actual examples of "how and why" RCTs.

[63] There is still plenty of work to be done on this front. But some criteria discussed by Emmanuel et al. (2000) relevant here include favorable risk-benefit ratio and fair subject selection. MacKay's (2020) criteria for permissible randomization will also be worth considering.

complaints about the social value of an experiment. Social equipoise also explains why methodological considerations are ethically relevant to experimental development economics. The methodological criticisms of development RCTs can as a result be interpreted as complaints about the social value of such experiments.

8. Grounding the Principle of Social Equipoise

Kukla (2007, p. 184) claims that their version of equipoise follows from the "basic principle that the research enterprise gives investigators no license to compromise citizen's moral entitlements to justice, respect, and welfare protection." Of course, persons have a negative duty to not interfere with anyone's entitlements to justice, respect, and welfare protection. But if researchers *qua* persons do not have a positive duty to provide others what they are owed as a matter of justice, then it is unclear why researchers (and this includes development economists) have a duty to abide by the principle of social equipoise. After all, how can a negative duty provide the normative basis for a positive ethical principle?

To resolve this difficulty, it is helpful to revisit Emmanuel et al.'s (2000) ethical framework for medical research. These commentators suggest that the main reason why value and scientific validity are key ethical considerations is that finite resources should be used responsibly. This reason presumably stems from a more fundamental ethical requirement of some sort. Identifying this ethical requirement is crucial for providing the principle of social equipoise with an adequate philosophical foundation.

8.1 The Natural Duty of Beneficence

To ensure wide applicability, it is best to remain committed to Kukla's original intention of grounding the principle of social equipoise in ethical requirements everyone has in virtue of being persons, or natural duties for short.[64]

The first candidate duty I will consider as a possible philosophical basis for the principle of social equipoise is the natural duty of beneficence. The natural duty of beneficence is a positive duty that requires persons to promote the good of others. There is no denying that there is considerable controversy with respect to the demands of abiding by the natural duty of beneficence. While I cannot settle these controversies here, I take it as uncontroversial that all persons are—to some extent—required to promote the good of others and that this duty can be discharged in a variety of ways by a variety of different socially situated persons. Since development economists are not exempt from the natural duty of beneficence, this opens one possible normative basis for the principle of social equipoise: development economists comply with the natural duty of beneficence when they follow the principle of social equipoise. The natural duty of beneficence would further explain why researchers are required to use finite resources responsibly. By not following the principle of social equipoise and engaging in socially valueless research, development economists are squandering resources that could be put to more beneficent uses. In other words, development economists would fail to consider the (ethical) opportunity costs of pursuing an experiment.

While the natural duty of beneficence provides an initially attractive normative basis for the principle of social equipoise, there are two drawbacks to relying on it. The first is that the

[64] See Pierson and Millum (2018) for a discussion of funders obligations in the context of health research. Notably, Pierson and Millum argue that funders have an obligation to maximize social value. The principle of social equipoise could also be seen as extending to funders of development RCTs as well.

natural duty of beneficence is often considered to be an imperfect duty. Unlike perfect duties, the natural duty of beneficence is not correlative with any specific person's claim-rights or entitlements. Though the natural duty of beneficence requires us to promote the good of others, it does not specify that any specific person is entitled to our beneficence. This creates a problem for the principle of social equipoise. One appealing feature of grounding the principle of clinical equipoise in the therapeutic obligations of physicians is that such grounding implies that a violation of the principle directly wrongs research participants. On the current proposal, a development economist who violates the principle of social equipoise has perhaps done something wrong, but they have not wronged any specific individual. If possible, it would be better to base the principle of social equipoise on a perfect duty so that research participants have a basis to claim they were wronged by being enrolled in a socially valueless experiment.

The first drawback to the natural duty of beneficence being the basis for the principle of social equipoise naturally leads to a second drawback, namely, that beneficence may be an inappropriate ethical consideration to appeal to when it comes to questions of global poverty. Consider what Nelson Mandela has claimed in the past: "overcoming poverty is not a gesture of charity. It is an act of justice. It is the protection of a fundamental human right, the right to dignity and a decent life" (quoted in Valentini, 2021, p. 48). Valentini (2021) takes Mandela's assertion as reflecting a commitment to a popular sentiment in political discourse; she interprets Mandela as saying that "assistance is something the well-off *owe to* the world's poor, not a 'mere' matter of beneficence" (Valentini, 2021, p. 48).[65] If possible, the principle of social equipoise should reflect this concern, and the only way to do so is by identifying a perfect duty on which to base it.

[65] Valentini ultimately argues against this view.

8.2 Respect for Persons

A more plausible basis for the principle of social equipoise is the natural duty of respect for persons. The duty to respect persons originates in the Kantian tradition in moral philosophy, but a full-fledged commitment to Kantianism is not needed for its endorsement (Millum and Bromwich, 2020). The importance of respecting persons has, for example, been enshrined in foundational documents in research ethics such as the *Belmont Report* (National Commission for the Protection of Human Subjects of Biomedical and Behavioral Research, 1978). Following Millum and Bromwich (2020), I maintain that respect for persons is best understood as a form of recognition respect for rights-holders. In research ethics, the importance of respecting persons has traditionally provided the normative basis for obtaining informed consent from autonomous research participants (discussed in more detail in the following chapter). However, I propose that respecting persons can play an additional role when it comes to research in development economics. Specifically, I propose that respecting persons provides the normative basis for why finite resources should be used responsibly.

To understand the conceptual link between respecting persons and the responsible use of finite resources, one first needs to refer to the unmet claims that the global poor have to socioeconomic resources and opportunities. As I argued in the previous chapter, these are not claims that the global poor have against development economists or other private actors. So, I cannot maintain here that that the global poor's claim-rights are correlative with an individual duty of distributive justice that private actors have. I instead propose that using finite resources to fund valueless research fails to take these unmet claims seriously, and therefore fails to show adequate respect to the claimholders in question, i.e., the global poor. Consider an expensive experiment that violates the principle of social equipoise and is therefore likely to have no social

value. Assuming other relevant ethical criteria is met (e.g., there is no risk of harm to participants), the chief wrong with going through with the experiment is that it would fail to acknowledge the unmet claim-rights of research participants. Instead of going through with the experiment, the funds could instead be used in a way that mitigates—and thereby acknowledges—the wrongs suffered by those whose claim-rights are unmet and will continue to go unmet. By failing to abide by the principle of social equipoise, development economists would therefore be violating their natural duty of respect for persons.[66]

9. Cohen and Dupas's Experiment Revisited

There is no denying that the foregoing discussion of equipoise has been abstract. In closing, I revisit the experiment discussed at the outset and argue that it met the principle of social equipoise. By reviewing the experiment, I want to suggest the following: (1) social equipoise has already been playing *some* role in experimental development economics; and (2) Cohen and Dupas (2010) provides a helpful model for abiding by the principle of social equipoise. In anticipation of any potential misinterpretations of my position, I am *not* concluding that the principle of social equipoise is being routinely complied with in experimental development economics based on the analysis of one case study.

Hearing the details of Cohen and Dupas's (2010) experiment undoubtedly makes outside observers uneasy. Why not provide free ITNs to as many people as possible if they are known to be therapeutically effective against malaria and known to generate positive externalities? As Cohen and Dupas make clear at the start of their article, prior to their experiment there was well-

[66] There is also a basis for saying that enrolling participants into a study that is socially valueless exploits them. Short of providing a full account of exploitation, one can rely on Siegel's (2008) Kantian theory of exploitation to make this argument. Siegel proposes that a sufficient (though not necessary) condition for A exploiting B is that A demeans or degrades B. Since failing to treat someone with the respect they are due as persons is just another way of saying that they have been demeaned or degraded, it follows that engaging in valueless research demeans or degrades someone. Therefore, violations of the principle of social equipoise are—in one sense—exploitative.

known professional disagreement over whether health products which generate positive externalities should be given away for free or whether a positive subsidized price should be charged. The debate was undoubtedly a question about which intervention promoted a development-related outcome (reducing malaria transmission) more effectively. What is noteworthy is that this debate had two prominent development economists advocating for conflicting positions. Sachs (2005) was in favor of free distribution and relied on standard economic theory to argue that making ITNs free would lead to higher uptake and use. Easterly (2006) opposed free distribution and cited empirical findings from behavioral economics to suggest that fully subsidizing ITNs can have the opposite effect. Easterly reasoned that charging a positive price could induce people to use ITNs due to psychological sunk-cost effects (Thaler, 1980). Also noteworthy is that this debate had policy ramifications, though crucially neither of the agencies involved were government entities (at least, initially). The World Health Organization (WHO)—an agency of the United Nations—had sided with Sachs while Population Services International—a nonprofit global health organization—had sided with Easterly.[67]

Prior to experimentation, it was clear that there was disagreement in the expert development community about the social value of full versus partial subsidization of ITNs. What's more, since this disagreement was based on a clash between standard economic theory and previous empirical findings, a well-designed RCT could generate evidence relevant to the debate. This is not to suggest that evidence from one experiment could definitively settle the debate as some may have initially been inclined to think.[68] Though Cohen and Dupas do not use the term "social equipoise" anywhere in their article, it is clear that, prior to their experiment,

[67] See Kremer and Holla (2009) for an overview of the more general debate over how pricing affects take-up of education and health services and products, and how evidence RCTs has contributed to this policy debate.
[68] See Rodrik (2009, pp. 28-39) for a more detailed analysis and commentary. As Rodrik makes clear, Cohen and Dupas's experiment is not immune from external validity critiques. Cf. Kremer and Holla (2009).

social equipoise existed; this is something the article reporting their findings makes clear from the first page.

As it turns out, Cohen and Dupas ultimately found that free distribution of ITNs led to both higher uptake and usage than partial subsidization (or no subsidization whatsoever). Yet here my optimism in social equipoise playing a significant role in experimental development economics is tempered by the continued lack of methodological consensus surrounding the use of RCTs. Though I believe that I have made a plausible case for the claim that the principle of social equipoise was met in Cohen and Dupas's experiment, Ogden (2020) points out that the evidence from their experiment has ultimately proven more convincing to non-economists than actual economists. While there is no denying Cohen and Dupas's findings have played a key role in subsequent policy decisions by NGOs operating in low-income countries as well as the policy decisions of governments,[69] similar conclusions about the subsequent beliefs of development economists are harder to draw. Even in what seems like a clear case in which the principle of social equipoise was met, there will still be detractors in the expert development community who insist that very little of value can be learned through a development RCT (e.g., Pritchett, 2020).

10. Conclusion

Some may be disappointed to learn that equipoise has little bearing on the ethical questions that development RCTs pose. It is still worth emphasizing that considerations of social equipoise can provide observers and practitioners of experimental development economics with a principled basis for ethical complaints of a certain type—namely, ethical complaints about the social value of an experiment. Though perhaps it is obvious that sound research practices are necessary

[69] For example, GiveWell directed $100 million to free ITN distribution after Cohen and Dupas's experiment (Ogden, 2020, p. 127). Cohen and Easterly (2009b, p. 18) also report that the experiment played a key role in the Kenyan government's eventual decision to fully subsidize ITNs for pregnant women.

component of ethical research, the principle of social equipoise can play a role in ensuring that this is the case.

Development economists should still grapple with the ethical questions posed by their new experimental methodology. However, equipoise will be of little help with this task, and further attempts to import the concept to experimental development economics may ultimately lead to more confusion than clarity.

Chapter 4: Who Has the Right to Cluster Randomize?

1. Introduction

In many development RCTs, randomization occurs at the group level rather than the individual level. Instead of randomizing individuals into treatment and control arms, development economists will randomize entire neighborhoods or villages. For example, in Cohen and Dupas's (2010) study on the effects of price subsidies on ITN use and uptake, different villages—not individuals—were randomized into different subsidy levels. This technique is known as *cluster randomization*. Because cluster randomization takes place at the group level, it is often not possible to obtain individual consent from participants. And in cases where participants do provide individual consent, there are always non-participants within clusters who do not consent to experimentation in their cluster. In many cases, individuals may not even be aware that they are part of (or not part of) an experiment (Barret and Carter, 2010).

The ethical problem cluster randomization poses should be immediately apparent: individual consent is easily the most known and well-regarded requirement of ethical research involving human-subjects. The first principle in the *Nuremberg Code* famously states that "the voluntary consent of the human subject is absolutely essential" (in Levine, 1988, p. 425). Though cluster randomization offers methodological advantages (discussed below), some may be inclined to think that an important ethical principle should not be neglected for methodological reasons. Baele (2013, p. 24) suggests that development economists' methodological preference for cluster randomized experiments offers them a convenient way to work around individual consent requirements, which he believes should not be neglected. But development economists are not alone in employing this experimental technique. Cluster randomization is also employed in health and public policy research, not just in low-income countries, but in high-income

countries as well. For this reason, I first focus on general ethical issues with cluster randomized controlled trials (cluster RCTs) before turning my attention back to experimental development economics.

As this chapter reveals, there are reasons to be suspicious of the claim that individual consent is an absolute deontic constraint on research involving human-subjects. U.S. Federal Regulations governing human-subjects research provide exemptions to individual consent requirements. Using these guidelines as a point of departure, Gelinas, Wertheimer, and Miller (Gelinas et al., 2016) develop a normative framework for assessing when and why health research without consent is permissible. Building off this framework and applying it to public policy experiments, MacKay and Chakrabarti (2019) have recently argued that government investigators do not need to secure participant's consent if certain conditions are met. The main task of this chapter is to extend a similar analysis to cluster RCTs performed by private non-government actors such as philanthropic foundations, international aid groups, non-profit research firms, and university-affiliated researchers. Research involving cluster randomization is often sponsored or conducted by such private actors, and this is especially true of policy research conducted in the developing world. My concern in this chapter is whether these private actors can permissibly cluster randomize without individual consent.

In this chapter, I argue that cluster randomization without consent can be permissible. In the next chapter, I expand my analysis of research without consent beyond the traditional concerns of research ethicists. As I argue there, what I refer to as "privatized policy experimentation" raises a deeper worry concerning political legitimacy. Taken together, these two chapters suggest that what is really alarming about research without consent is not that it

violates well-known principles from research ethics, but rather, that it threatens research participant's political rights.

The rest of this chapter is structured as follows. In section 2, I provide an overview of cluster RCTs. I discuss the methodological advantages of cluster randomized experiments and provide examples of cluster RCTs to help guide the ensuing discussion. In section 3, I review the main contributions to the literature on research without consent. In section 4, I discuss the importance of ensuring risks are minimal in contexts where there is no individual consent. In section 5, I extend the key insight from Gelinas et al.'s (2016) and MacKay and Chakrabarti (2019) to propose what I call the "authorization condition" for research without consent. In section 6, I use the authorization condition to analyze the experiments discussed in section 2. Lastly, in section 7, I turn my attention back to experimental development economics and supplement my framework for the non-ideal conditions that characterize the developing world. Section 8 concludes.

2. Why Cluster Randomize?

Cluster RCTs differ from the more familiar randomized controlled trials (RCTs) used in clinical research in one important respect: the unit of randomization is a social entity, e.g., a hospital, school, village, or neighborhood. In the health research context, cluster RCTs are used for knowledge translation research and quality improvement research (McRae et al., 2011). Public health researchers also use cluster RCTs to evaluate the effectiveness of social, behavioral, and community-level interventions (Sim and Dawson, 2012). And as mentioned at the outset, cluster RCTs are used by development economists to evaluate the effectiveness of interventions on questions relating to socioeconomic outcomes in low-income countries.

In addition to the methodological advantages that randomized experiments confer, two methodological reasons are used to justify the use of a cluster RCT over another study design (Edwards et al., 1999). The first reason is the nature of the intervention being evaluated. The delivery mechanism for some interventions precludes the possibility of assigning individuals into treatment and control arms. For example, consider a study that evaluates the effects of fluoride in the water supply. While an individual can evade the intervention by drinking bottled water, they cannot strictly speaking opt out of the experiment without moving to a different cluster. Researchers cannot ensure that the fluoride in the water supply only reaches households that consent to the intervention. For many interventions of interest, it is simply not possible administer a treatment in an individuated manner.

The second methodological reason is to avoid statistical biases. In addition to reducing selection bias, cluster RCTs also allow researchers to reduce the possibility of contamination bias. Contamination bias occurs when individuals not assigned to the treatment arm of an intervention are nevertheless exposed to the intervention via interaction with those who have been assigned to the treatment arm. The possibility of these "leaks" can cast doubt upon the scientific validity of a study.

The salient ethical problem with cluster RCTs is that they make obtaining individual consent impracticable. But it is important to note that not all cluster RCTs make individual consent completely impracticable. In some cluster RCTs, participants can consent to having their data collected. However, in most cluster RCTs, individuals do not consent to being randomized into treatment and control arms, meaning that their environment is manipulated without their having any say in the matter. This is unlike a clinical trial where individuals provide consent for both (a) randomization into treatment and control arms and (b) data collection.

We can track this difference by introducing a distinction between individual-cluster randomization and cluster-cluster randomization (Edwards et al., 1999). The latter makes any form of individual consent impracticable, viz., consent is neither obtained for randomization nor for data collection. As an example of a cluster-cluster RCT, consider a recent study on public health messaging by Breza et al. (2021). Though study was conducted in the U.S., what is noteworthy is that Abhijit Banerjee and Esther Duflo—two pioneers of development RCTs—are co-authors. It is worth looking closely at the details of this cluster RCT. If the past is any guide to the future, then other researchers will want to emulate Banerjee and Duflo's success going forward.

During the height of the COVID-19 epidemic, Breza and colleagues were interested in the effects of health professionals promoting preventative health behaviors on social media. Their interest was spurred by the Kaiser Family Foundation, which had sponsored a large project where doctors recorded videos to provide information about COVID-19 vaccinations. To research the question of interest, Breza and colleagues conducted a large-scale public health advertising campaign on Facebook. The campaign consisted of short videos by health professionals urging users to stay home for the Thanksgiving and Christmas holidays. 820 counties in 13 states were selected, and then zip codes within those counties were assigned into treatment or control arms. The treatment was exposure to short videos by public health experts. The first outcome of interest was holiday travel and was measured with publicly available Facebook mobility data. The second outcome of interest was COVID-19 infections recorded at zip code level, again publicly available. Breza and colleagues found that the social media messages led to statistically significant reductions in holiday travel and COVID-19 infections, and so there is a case to be made that the research was socially valuable. However, this study

clearly demonstrates the problems with cluster RCTs: Facebook users involved in the study did not consent to be randomized or have their data collected for the purposes of a scientific study.

Unlike cluster-cluster randomization, individual-cluster randomization involves randomizing clusters but delivering interventions individually, which thereby allows for some for some form of individual consent to take place. This may lead one to believe that individual-cluster RCTs are ethically innocuous and less cause for concern. But unlike clinical trials, individual consent in an individual-cluster RCT typically takes place *after* randomization has already occurred (McRae et al., 2011). What participants consent to is data collection, not randomization into treatment and control arms. A prime example of an individual-cluster RCT is the Cohen and Dupas (2010) study. Recall how in Cohen and Dupas's (2010), only pregnant women were sold subsidized ITNs, with the rate of subsidy determined randomly by village. The pregnant women who purchased ITNs and agreed to be involved in the study did provide individual consent to be tracked and have their hemoglobin levels recorded for data collection purposes. However, none of the pregnant women involved in the study consented to randomizing the price of ITNs by location. These women had their environments deliberately manipulated without having any say in the matter.

What both cluster-cluster and individual-cluster randomization have in common then, is that there is no individual consent for randomization into treatment and control arms. Most of my focus in this chapter is with the fact that cluster RCTs pose a challenge to obtaining consent for randomization. For the most part, I will set aside concerns about when data collection requires individual consent because this question has already received attention by research ethicists and poses no unique problems to cluster RCTs.

3. Research Without Consent

3.1. Early Contributions

Despite the known use of cluster RCTs by public health and health policy researchers, very little has been written on the permissibility of cluster randomization. This is somewhat surprising given the importance that individual consent played in the development of medical research ethics. In what is (to my knowledge) the earliest discussion of ethical issues with cluster randomization from a public health perspective, Edwards et al. (1999) suggest that individual consent can be replaced by the consent of an appropriate "guardian". Examples of guardians include "the chief executive of a hospital, managing partner of a primary care practice, or head teacher" (Edwards et al., 1999, p. 1408). Sim and Dawson (2012) worry that so-called guardians will not always represent the interests of all members of a cluster and instead argue that the permissibility of a cluster RCT should be determined by an institutional review board (IRB), which will often be in a better position to identify the interests of different stakeholders.

Neither of these approaches offer a satisfying answer to the challenge that cluster randomization poses to individual consent. Sim and Dawson (2012) are correct to insist that IRBs should expand the scope of their risk-benefit assessment in a cluster RCT. After all, a key function of an IRB is to make risk-benefit assessments and ensure risks to research participants are minimal and reasonable (Weijer, 2000; Weijer and Miller, 2004). However, in most research contexts needing IRB-approval, the IRB performs this key function *before* individual consent is obtained.[70] Sim and Dawson's proposal fails to address the real issue at stake because, as we will see below, they misidentify the main role that individual consent plays in human-subjects research.

[70] Wertheimer (2010, ch. 2) notably argues that the prevailing regulatory and ethical framework for human-subjects research is paternalistic despite all the emphasis placed on the importance of consent and autonomy.

Can the guardianship model advocated for by Edwards et al. (1999) provide the basis for group consent as a replacement for individual consent *after* an IRB approves a cluster RCT? While this approach seems initially attractive, there are a lot of details left unspecified. Edwards et al. (1999) quickly suggest that representative capacity could be determined democratically. However, democratically electing representatives would make the group consent process just as cumbersome as obtaining individual consent. Further, the "guardians" which Edwards et al. (1999) identify are typically not appointed via democratic processes. This is not to suggest such appointments are illegitimate. On the contrary, institutions like hospitals and schools are permitted to make leadership appointments non-democratically. As the discussion below reveals, questions of individual consent fall to the wayside once we consider that "guardians" have institutionally specified rights to authorize certain interventions.

3.2. The Common Rule Guidelines

As mentioned at the outset, U.S. Federal Regulations governing human-subjects research—otherwise known as "the Common Rule"—allow for waiver of individual consent requirements if certain conditions are met. These requirements provide a helpful point of departure for the ensuing discussion. I follow McRae et al.'s (2011, p. 6) helpful formulation of The Common Rule's requirements:

1) The research involves no more than minimal risk to the subjects;
2) The waiver or alteration will not adversely affect the rights and welfare of the subjects;
3) The research could not practicably be carried out without waiver or alteration; and,
4) Whenever appropriate, the subjects will be provided with additional pertinent information after participation.

Most of the ensuing discussion will focus on the second condition above. This is where the philosophical action lies since it is unclear what rights (if any) are at stake when individual consent requirements are waived. This is not to suggest that the remaining three conditions are

unimportant. All four conditions are ethically important, but it is the second condition that requires the most philosophical scrutiny.

3.3 Health Research Without Consent

Gelinas et al.'s (2016) framework for research without consent is, to date, the most promising approach to justifying the use of cluster randomization. Gelinas et al.'s (2016) analysis focuses on healthcare research, but as we will see further below it can also be generalized to policy research more broadly construed.[71] Gelinas et al.'s framework abandons any appeal to proxy consent. Instead, Gelinas et al. begin with an analysis of the function that individual consent plays in human-subjects research. For Gelinas et al., consent is meant to first and foremost (a) protect a research participant's autonomy rights but also (b) promote and protect a research participant's welfare interests.[72] Though (a) and (b) may conflict in certain cases, for Gelinas et al. the main point is that "the most basic function of consent is to waive rights of control, allowing others to interact with us in ways that would otherwise be wrong" (Gelinas et al., 2016, p. 36).

The analysis of individual consent which Gelinas and colleagues offer is significant because it opens the possibility that cluster randomization is permissible when it neither violates autonomy rights nor threatens the welfare interests of any individual research participants (or non-participants). Like Gelinas and colleagues, I maintain that there is no general right to *not* be included in research without consent. Positing such a right would be implausible since it would

[71] Cluster randomization is only one example of research without consent which Gelinas et al. (2016) analyze. Gelinas and colleagues also analyze research involving stored biological samples and simulated patient studies, which is not my concern here.

[72] Edwards et al. (1999) and Sim and Dawson (2012) fail to address the issue at hand because they entirely focus on (2) instead of (1).

rule out uncontroversial study designs, e.g., observational research done in public spaces or with anonymous data (Gelinas et al., 2016, p. 36).

Though Gelinas and colleagues propose two separate criteria for when research without consent is permissible, for now I only focus on the first. Gelinas et al. propose the following two conditions which are jointly sufficient for making research without consent permissible: "(1) the research stands to infringe on no right of the participants and (2) it is impracticable to obtain consent" (Gelinas et al., 2016, p. 35). Before elaborating on the individual rights condition (1), its important get clear on the impracticability condition (2). Gelinas et al. specify that by "impracticable" they mean "either that obtaining consent imposes (prohibitively) high costs of time or economic resources on researchers or that obtaining consent threatens to undermine the scientific validity of the research" (Gelinas et al., 2016, p. 36). As Gelinas et al. emphasize, it is still important to obtain consent in cases where it is practicable *even if* failure to do so violates no one's rights. This is because doing so promotes public trust and transparency in the research process—an important social good (Gelinas et al., 2016, p. 36). Whereas the justification of Gelinas et al.'s first condition is rights-based, it is important to note that, on their framework, the justification of the second condition (impracticability condition) is instrumentalist (MacKay and Chakrabarti, 2019, p. 190).

For the purposes of analyzing cluster randomization without consent, what is important to focus on is Gelinas et al.'s individual rights condition. Consider Gelinas et al.'s analysis of quality improvement (QI) studies undertaken at the institutional level. An institutional level QI study will evaluate the effectiveness of an intervention on healthcare delivery or patient outcomes in an individual healthcare facility or healthcare system. QI studies are often done without individual consent. But Gelinas and colleagues argue that because institutions such as

hospitals and healthcare systems have certain rights of control, QI studies without consent are permissible when they evaluate interventions within their realm of control. Like Gelinas et al., I will not specify the exact boundaries of where a healthcare institution's rights of control end and where a patient's right to exercise autonomy over their healthcare decisions begin. Denying that healthcare institutions have a legitimate sphere of control would be highly implausible as evidenced by the example below, which also illustrates Gelinas et al.'s approach to research without consent:

> institutions have a right to require that clinicians wash their hands with a particular type of disinfectant soap, without consulting or obtaining consent from patients for this practice. Given that this is a practice within the decisional authority of the institution, so long as the institution ensures a standard level of care with regard to hand washing and sanitation, institutions should have the right to conduct research comparing two types of disinfectant soap for use by clinicians without patient consent. The rights of patients will not be infringed by the conduct of such QI research without consent (Gelinas et al., 2016, p. 37).

Notice that Gelinas et al. also invoke the notion of a standard of care to justify QI studies without consent. Some may take this to imply a third condition for research without consent that Gelinas et al. do not make explicit. However, this consideration illustrates the flexibility of their framework nicely. Providing patients with substandard care would presumably infringe on the claim-rights of research participants, which would thereby make condition (1) unfulfillable. However, going forward the main kind of rights-violations that will be of interest are autonomy rights-violations since individual consent is meant to protect such rights.

With respect to the present inquiry, the main lesson from Gelinas et al.'s framework can be summarized as the following general principle:

> *Intervention Without Consent*: In a research context, *X* may permissibly administer intervention *I* in population *P* if *X* possesses the right to implement intervention *I* in *P* in a non-research context.

In the example above, the hospital has the right to decide which disinfectant soap a clinician uses in a non-research context. Therefore, the QI study on disinfectant soap is permissible. The autonomy rights of clinicians and patients are not affected since they lack a right to determine which disinfectant soap a hospital uses.

3.4 Public Policy Research Without Consent

MacKay and Chakrabarti (2019) have extended Gelinas et al.'s framework to policy research more broadly construed. The key insight to MacKay and Chakrabarti's analysis is that *governments* are also institutions with rights of control. Specifically, governments have rights of control over *policy* decisions within a given territory. As MacKay and Chakrabarti point out, it would be puzzling to maintain that a government has the right to subject residents to a new policy but nevertheless insist that a pilot program of that policy is impermissible to carry out.[73] Using this insight, MacKay and Chakrabarti propose two jointly sufficient conditions for permissible policy experimentation without consent: "(1) the government institution conducting or authorizing the experiment possesses a right to rule over the spheres of policy targeted by the research; and (2) does not involve the violations of participant's autonomy rights" (MacKay and Chakrabarti, 2019, p. 188).

Notably, MacKay and Chakrabarti dispense with Gelinas et al.'s impracticability condition when policy research is conducted in *ideal* conditions. However, MacKay and Chakrabarti do argue that the impracticability condition is necessary in non-ideal conditions (MacKay and Chakrabarti, 2019, pp. 195-196). This consideration is especially relevant in the context of experimental development economics (section 6), so I will revisit it further on. Before

[73] Governments still have duties to treat residents fairly when evaluating the policy in question. This aspect of policy experimentation is taken up more thoroughly in MacKay (2020).

doing so, it is important to first focus on conditions (1) and (2) as they apply to policy experiments in ideal circumstances.

Justifying condition (1) should be relatively uncontroversial. The notion of a government's "right to rule" is fundamentally connected to whether a government and its actions are legitimate. However, what *exactly* makes a government legitimate *is* undeniably a controversial topic—it is arguably the central question of modern political philosophy. But unless one is a philosophical anarchist (e.g., Wolff, 1970) or holds a theory of legitimacy that no real-world governments currently meet (e.g., Simmons, 1999, pp. 769-770), one can maintain that governments (and their various branches and agencies) possess the right to enact certain policies so long they are within their realm of control. What actions are within a government's realm of control is, of course, determined (in part) by its resident's autonomy rights, which are reflected in both conditions (1) and (2). Like MacKay and Chakrabarti, as well as Gelinas et al. (2016) in the healthcare context, I will not specify a full account of where a government's right to rule ends and a person's autonomy rights begin. Doing so would constitute a full-fledged theory of political legitimacy, which is well-beyond the scope of this chapter, and which is superfluous if one already holds that real-world liberal democratic governments have the right to enact *some* policies. However, in discussing experiments in developing countries, it will be important to introduce a theory of *minimal legitimacy*, which I do in section 6.

Though it is not possible to specify a full list of autonomy rights here, there are some worth explicitly mentioning. This is also important to make sense of condition (2). Here, I appeal to MacKay's (2015) concept of *generic liberalism*, i.e., a family of claims that "(1) most, if not all, liberal political theories support, and (2) most, if not all, liberal democracies institutionalize or aspire to institutionalize" (MacKay, 2015, p. 3). One important autonomy right that has

already been discussed is the right to make certain decisions about one's own healthcare. This right presumably stems from a more fundamental right to bodily autonomy. Though its implausible to maintain that the right to bodily autonomy is absolute,[74] the right to bodily autonomy is undeniably a staple of generic liberalism.[75] Research requiring the drawing of blood samples, as in the Cohen and Dupas's (2010) experiment, clearly require individual consent regardless of institutional context. In all but the most extreme cases then, it should be uncontroversial to maintain that it is impermissible to conduct research without consent when an intervention or data collection process involves a bodily intrusion. Such research would clearly not be within the realm of control of any government or social institution. Other uncontroversial autonomy rights include the right to basic freedoms such as freedom of expression, association, and religion. It should be safe to assume that any intervention or data collection process that would somehow interfere with these familiar autonomy rights would be impermissible to carry out.

4. Minimal Risk of Harm

MacKay and Chakrabarti's framework for policy research suffers from one deficiency. Recall that the basic function of informed consent is to (a) protect a research participant's autonomy rights but also (b) promote and protect a research participant's welfare interests. MacKay and Chakrabarti largely ignore (b) by only focusing on autonomy rights. Gelinas et al. (2016) similarly gloss over welfare considerations despite acknowledging that individual consent also plays a role (in addition to the IRB) in protecting and promoting an individual's welfare interests. Individual consent does this by ensuring that research does not expose participants to risks to

[74] Though perhaps a controversial (and timely) subject, many believe that governments have the right to impose vaccine mandates on citizens, which involves a bodily intrusion.
[75] It is also worth noting that political philosophers such as Fabre (2006) have challenged the primacy of the right to bodily autonomy in liberal thought.

their welfare that they are not willing to bear. It is for this reason that the Common Rule's conditions for waiver of individual consent emphasize that research involve no more than minimal risk. MacKay and Chakrabarti's framework for policy research without consent can thereby be supplemented with one additional condition: (3) the research involves minimal risk of harm to participants and non-participants.

One could argue that governments only have the right to implement policies that impose minimal risk on its residents, thereby making condition (3) redundant. While this may be true, MacKay and Chakrabarti's first condition still does not consider the risks that participants may be exposed to as part of an experiment that are independent of the policy intervention being evaluated. For example, data collection—even if it does not involve violations of anyone's autonomy rights—could still involve exposing research participants to non-trivial risks if the data were to fall into the wrong hands. There could also be unintentional but foreseeable harms that come about through research participant's interactions with non-participants in their cluster. Imagine, for example, that a study on unconditional cash grants causes non-recipients to become jealous, which then culminates in violent conflict. Scenarios such as these are well within the realm of possibility. For this reason, it is important to make condition (3) explicit.

What constitutes a minimal risk? Here, I appeal to the Common Rule: "the probability and magnitude of harm or discomfort anticipated in the research are not greater in and of themselves than those ordinarily encountered in daily life or during the performance of routine physical or psychological examinations or tests" (U.S. Department of Health and Human Services, 2018). Research that involves minimal risk should therefore be understood as comparable to the risks one assumes in conducting their daily affairs.[76] Risks beyond this

[76] The minimal risk standard needs much more elaboration, which I do not have space to do here. There is an ambiguity with interpreting what counts as "minimal risk" paralleling the discussion of the standard of care in a

threshold are not minimal. Exposing participants to such risks without their consent—even in MacKay and Chakrabarti's first two conditions are met—would be wrong.

One may object that this requirement is unnecessary for the same reason that I criticized Sim and Dawson's (2012) proposal: IRBs are already tasked with ensuring that risks are minimal. However, this does not obviate the need to make the requirement explicit. After all, when it comes to approving a cluster RCT, an IRB should also be tasked with determining whether any aspect of a research protocol violates the autonomy rights of cluster members. It is also worth emphasizing that IRBs are not infallible. Researchers still have ethical duties to minimize risks even after their research protocol receives IRB approval.

5. The Authorization Condition

MacKay and Chakrabarti's (2019) framework for policy research without consent has important implications for policy experiments conducted by private actors. MacKay and Chakrabarti's definition of a "government policy experiment" includes experiments conducted by private actors on behalf of a government, which may lead one to think that there is nothing more to add when it comes to private actors conducting policy experiments. But a significant number of policy experiments are also sponsored by philanthropic foundations and then designed and conducted by private non-profit research firms. For example, the Abdul Latif Jameel Poverty Action Lab (J-PAL), one of the non-profit research firms most known for conducting policy experiments, is largely funded by private foundations.[77] It would be misleading to subsume these

---

previous chapter. For example, the risks one assumes in their daily life in a high-income country are vastly different than the risks someone assumes in a country marred by violent conflict. Even more vexing is whether there is some sort of *de jure* interpretation of minimal risk indexed to one geographic location, i.e., what amount of risk *ought to be* commensurate with daily life. The key obstacle to developing any *de jure* interpretation of minimal risk is respecting people's differing risk preferences.

[77] J-PAL's website lists the following as major donors: Arnold Ventures, Co-Impact, Community Jameel, Echidna Giving, The Bill and Melinda Gates Foundation, Google.org, The William and Flora Hewlett Foundation, King Philanthropies, The John D. and Catherine T. MacArthur Foundation, The Douglas B. Marshall Jr. Family Foundation, Omidyar Network, and The Alfred P. Sloan Foundation. J-PAL also lists funding from government

experiments under the heading of "government policy experiments" since these experiments are not necessarily proposed by or conducted on behalf of a government.

While MacKay and Chakrabarti (2019) do not make this explicit, I believe it follows from their framework that private actors such as philanthropic foundations, international aid groups, non-profit research firms, and university-affiliated researchers cannot evaluate public policy interventions without first securing appropriate government authorization. After all, governments, and only governments, possess the right to rule over public policy matters. Additionally, there are interventions which only a government may have the right to implement, or which a private actor may lack an institutionally specified right to implement. This suggests that in some contexts, private actors do not have the right to cluster randomize without appropriate authorization. The Breza et al. (2021) study, which I revisit below, may be one example.

When it comes to privatized policy experimentation, I propose the following as a modification to the first condition in MacKay and Chakrabarti's (2019) framework:

> *Authorization Condition*: A private actor may permissibly conduct a cluster RCT on an intervention it does not have an institutionally specified right to implement if the private actor has authorization from the appropriate government agency or social institution.

I derive the authorization condition from two general principles. First, we need the general ethical principle implicit in Gelinas et al.'s (2016) framework, reproduced again below:

> *Intervention Without Consent*: In a research context, *X* may permissibly administer intervention *I* in population *P* if *X* possesses the right to implement intervention *I* in *P* in a non-research context.

And then we need a principle allowing for rights to be transferred.

> *Transfer of Right:* if *X* has a right to *Ø*, then *X* has a right to authorize *Y* to *Ø* on *X's* behalf.

---

agencies such as the Australian Department of Foreign Affairs and Trade and the UK Foreign, Commonwealth & Development Office (J-PAL, 2022).

The authorization condition is simply an application of *Intervention Without Consent* combined with *Transfer of Right.*

Combined with the MacKay and Chakrabarti's original second condition regarding data collection, as well as my proposed minimal risk condition, these three conditions provide a framework for privatized policy research. Like in MacKay and Chakrabarti's (2019), the three conditions are jointly sufficient for making privatized policy research without consent permissible. Private actors may thereby permissibly cluster randomize without consent when all three conditions are met. But note that the proposed framework is only a formal framework. It says nothing substantive about what rights private actors possess. Whether a particular instance of cluster randomization without consent is permissible cannot be settled by the authorization condition alone. To apply the authorization condition (and the framework), one needs to first identify whether a private actor possesses the relevant rights in non-research contexts.

Some may even be inclined to think that the authorization condition is all there is to say about policy experiments conducted by private actors. While *Intervention Without Consent* is plausible, *Transfer of Right* may be contentious. For example, many philosophers accept that the autonomy rights of some individuals are inalienable, meaning that they cannot be waived or transferred. One many similarly suspect that this inalienability thesis extends to various dimensions of the right to rule which governments possess.[78] In the next chapter, I explain in detail why research without consent threatens the *political* rights of participants even when the authorization condition is met. However, I ultimately argue that this challenge to the authorization condition can be answered.

[78] See Cordelli (2020) for a book-length treatment of this position. I draw on Cordelli (2020) in the next chapter.

To summarize, when it comes to privatized policy research, I am proposing three conditions which are jointly sufficient to make cluster randomization without consent permissible:

1) The private actor has the (a) right to implement the intervention of interest or (b) authorization from the appropriate government agency or social institution to implement the intervention.
2) The data collection does not involve the violations of participant's autonomy rights.
3) The research involves minimal risk of harm to participants and non-participants.

Before turning to the main challenge to the authorization condition in the next chapter, I return to the cluster RCTs discussed earlier. The two cluster RCTs introduced earlier show how the authorization condition can be applied. In the penultimate section of this chapter, I expand my framework to experimental development economics and non-ideal conditions, more generally.

6. Applying the Authorization Condition

6.1. Did Breza and Colleagues do Something Wrong?

Breza et al. (2021) is a perfect case study for the authorization condition. This is because Breza and colleagues (arguably) met conditions (2) and (3), so their experiment allows us to test my proposed framework with a controlled real-world example. If there is something wrong with what Breza and colleagues did, I suggest that it is not because of how data was collected or that it exposed participants to risk of harm.

Before turning to the authorization condition, I quickly review why Breza and colleagues met conditions (2) and (3). First, consider the data collection process. Breza and colleagues had access to two sources of data: Facebook mobility data and county-level COVID-19 infection data. Both sources of data are publicly available. Further, both data sources contain anonymous data, i.e., data that is aggregated without identifiers and can never be linked to an individual. Per the Common Rule, research with anonymous data does not constitute human-subjects research,

and therefore does not require individual consent.[79] One could still object that the Common Rule does little to protect the rights of Facebook users have over "their" data, but there is a strong case to be made that we do not "own" our online data in the first place (Goodrich, 2021).[80] It is therefore difficult to maintain that Breza et al. (2021) did anything wrong on data collection grounds.

Breza and colleagues also needed (and received) IRB approval for their study. This means that the minimal risk condition was likely met. But Breza and colleagues took additional steps to ensure that they complied with minimal risk requirements. A careful reading of their study reveals that they withheld the treatment from some counties out of concern for adverse effects:

> Given the growing polarization between the presidential election and the inauguration, and considering a small number of extremely negative and politically charged comments posted during the Thanksgiving campaign, the research team was concerned that the messaging campaign might have adverse unintended effects in very rural, heavily Republican-leaning counties (Breza et al., 2021, p. 1624).

Of course, it is hard to know for sure whether Breza and colleagues seriously harmed anyone with their intervention. But this seems unlikely given their own diligence and IRB oversight. Again, if there is something objectionable about Breza et al. (2021), it is not due to violations of the data collection (2) or minimal risk (3) conditions.

There is still a case to be made that Breza and colleagues did something wrong because they violated the authorization condition. Consider that nowhere in Breza et al. (2021) do authors provide any indication that the Centers for Disease Control (CDC)—the U.S.'s federal public health agency—was aware of their experiment. The authors also do not provide any indication

---

[79] There is an important distinction between anonymous data and anonymized data. Anonymized data can be traced back to an identifiable individual and is considered human-subjects research per the Common Rule.

[80] There is also a case to be made that in making such data available as part of its Data for Good initiative (https://dataforgood.facebook.com/dfg/), Facebook is complying with its ethical obligations (see Goodrich, 2021).

that local public health agencies in treatment clusters were aware of the intervention. If Breza and colleagues did something wrong, I suggest that it is because they did not get authorization from the CDC and the relevant local health public agencies. The case against Breza and colleagues crucially hinges on whether public health messaging is within the legitimate sphere of control of private actors.

Appealing to the *Intervention Without Consent* does not settle the matter. Of course, Breza and colleagues could have purchased ads on Facebook promoting public health behavior without any intention of conducting an experiment. There is indication that the Kaiser Family Foundation had already done something similar. But the key question is whether private actors *should* have a right to engage in public health messaging in non-research contexts. At the time of the study, the CDC advocated that people stay home for the holidays, thereby indicating that the intervention was congruent with the CDC guidelines. To defend themselves, Breza and colleagues could appeal to the following principle:

> *Congruency Condition*: if a legitimate authority has a right to promote outcome *X*, then private actors also have a right to promote outcome *X* without authorization.

Determining whether Breza and colleagues violated the authorization condition ultimately depends on whether the congruency condition holds. I assume that, if Breza and colleagues contradicted the CDC's public health guidelines by instead showing videos of doctors encouraging holiday travel, then they would not have been within their rights to implement the intervention irrespective of whether they were conducting an experiment or not. This point highlights why the authorization condition matters: there may be some interventions that only a government has the legitimate authority to implement or authorize. Whether public health messaging is an example of such an intervention is exactly what is at stake here.

Settling the question about public health messaging by fully assessing the congruency condition would take us too far field. To do so, we would have to determine whether real-world governments (and their agencies) simply possess the relevant rights of control over certain interventions, or whether real-world governments (and their agencies) meet certain procedural criteria that ensures that the rights and welfare of subjects are not adversely affected by the exercise of authority. If we could determine such a criterion, we would be in a better position to assess when the congruency condition holds. This is because we would be able to assess when private actors share with governments whatever it is that permissibly allows governments to intervene without adversely affecting anyone's rights or welfare.

Even if public health messaging is not within any private actors' sphere of control (which would mean that Breza and colleagues violated research participant's autonomy rights), it may turn out that Breza and colleagues were all-things-considered justified in conducting their study. Minor rights violations can sometimes be justified and suggesting otherwise is highly implausible. Gelinas et al. (2016) incorporate this consideration by adding additional (previously not discussed) criteria for research without consent to their framework. The following two conditions are also jointly sufficient for establishing the permissibility of research without consent on Gelinas et al.'s framework: "(1) research that does infringe participant rights can be justified if the gravity of the rights infringement is minor and outweighed by the expected social value of the research and (2) if it is impracticable to obtain consent" (Gelinas et al., 2016, p. 36). Condition (2) is easily met in Breza et al. (2021). But there is also a strong case to be made that condition (1) was also met. Breza and colleagues discovered that a relatively inexpensive and non-intrusive intervention (short video ads on social media) had impact on enough people's behavior that it led to a desirable public health outcome (lower COVID-19 infection rates).

Though perhaps Breza and colleagues committed a minor rights infringement, public health agencies are now equipped with evidence of how to promote preventative health behavior in a cost-effective manner.

6.2. Cohen and Dupas's Experiment, Revisited Once More

Recall that Cohen and Dupas's (2010) study was a based on an individual-cluster RCT. Research participants provided consent to be tracked and to have their hemoglobin levels recorded, but none consented to having the price of ITNs randomized. Further, there were pregnant women who purchased ITNs in treatment facilities (meaning they were affected by the intervention) who did *not* have their data collected. These women may have not even known that the price of ITNs was being manipulated for research purposes.

Did Cohen and Dupas do something wrong by varying the price of ITNs without research participant consent? To answer this question, we need to look at some more details of their experiment design. Cohen and Dupas (2010, p. 11) report that there were 70 health clinics in the region they were interested in studying. These 70 health clinics were a mixture of public government run health clinics and private NGO-run health clinics. One of the criteria which Cohen and Dupas used to narrow their sample was public status. Ultimately, Cohen and Dupas only included 20 health centers in their sample. The 20 health clinics that were randomized into different subsidy levels were all public, i.e., managed and operated by the Kenyan government—specifically the Kenyan Ministry of Health.

As university-affiliated researchers, Cohen and Dupas do not possess rights of control over preventative health product pricing in Kenya, much less any country. However, the Kenyan Ministry of Health (presumably) does. And if one looks at Cohen and Dupas (2010) carefully, there is evidence that the researchers complied with the authorization condition. In a footnote,

Cohen and Dupas thank the Kenyan Ministry of Health and its staff for their collaboration, thereby indicating they complied with the authorization condition.

If the Kenyan Ministry of Health possesses rights of control over ITN-pricing in Kenya, then Cohen and Dupas did not violate any research participant's autonomy rights by cluster randomizing without their consent. It would be convenient to end the analysis of cluster RCTs here, but this would be hasty. Two issues remain. The first is that Cohen and Dupas's experiment influenced the Kenyan government's decision to fully subsidize ITNs. In the following chapter, I explain why this is concerning and offer a response. The second reason is that Cohen and Dupas's experiment, like all development RCTs, took place in a developing country. Just how advocating for economic policies that work in high-income countries can prove disastrous if implemented in low-income countries (see Rodrik, 2015), we must also be careful not to mechanically import ethical criteria that applies to research in high-income countries to research in low-income countries. I address this concern throughout the rest of this chapter.

7. Authorization in the Developing World

7.1. Authoritarian Development

Policy experiments sponsored and conducted by private actors are characteristic of the experimental approach to development economics. However, the development context poses complications for the authorization condition and policy research more generally. The main complication that must be addressed is the historical connection between recent Western assistance and authoritarian approaches to economic development. Easterly (2013) has extensively chronicled how Western government aid agencies, policymakers, philanthropists, and development experts have funded or been complicit with the actions of authoritarian regimes in the developing world. Western support for these authoritarian regimes has been prompted by

what Easterly forcefully argues is the mistaken notion that authoritarianism is necessary for the promotion of economic development.[81] The figures behind these regimes can hardly be described as benevolent autocrats. Easterly extensively documents the extensive human rights violations perpetrated by leaders of countries receiving foreign aid and technical expertise from the West. Because I have argued that private actors may be authorized to conduct policy experiments without consent, it is important to address the issue of authorization from an illegitimate government. As I will argue, authorization from an authoritarian regime should always be considered invalid. However, I will also argue that in some circumstances it *could be* all-things-considered permissible to conduct a policy experiment within a territory governed by an authoritarian regime.

7.2. Minimal Legitimacy

Historically, the Rockefeller Foundation was a key source of private funding for development policy research in pre-Mao China. As Easterly (2013, pp. 47, 55, 64-65) documents, in the 1920s and 30s the Rockefeller Foundation paved the way for the authoritarian approach to development by funding policy research in China. Easterly suggests that the Gates Foundation—arguably the biggest funder of development policy research today—plays a similar role in supporting authoritarian development as the Rockefeller foundation. However, Easterly provide no documentation that the Gates Foundation has financed authoritarian regimes. Easterly (2013, pp. 123-127) does provide clear evidence that Bill Gates has publicly praised the policies of autocratic rulers, which is not to downplay the severity of Easterly's criticism of Gates. I use this episode involving Gates to discuss the importance of ensuring that policy experiments are authorized by minimally legitimate governments.

[81] Sen (1999) advances similar arguments against the lure of authoritarian development.

In an article for the *Wall Street Journal* entitled "My Plan to Fix the World's Biggest Problems", Gates (2013) provides high praise for the Ethiopian government's recent progress on the problem of child mortality. Among the reasons cited for the Ethiopian government's success is adoption of a community-based health program modeled after a successful program in the Indian state of Kerala. Gates uses Ethiopia's success in reducing child mortality to argue that commitment to good measurement and careful analysis of field data are essential to fixing the problems that plague the developing world. What Gates does not mention anywhere is that, at the time the article was written, Ethiopia was led by Meles Zenawi—an autocratic ruler. Easterly (2013, pp. 155-158) documents some of Meles's transgressions, most notably a sham democratic election held in 2005 and its aftermath. When student demonstrators took to the streets to protest what international poll watchers regarded as a manipulated vote count, government security forces opened fire on them, resulting in more than 100 deaths. Meles's political opponents were subsequently denied access to food or thrown in jail.

Suppose that—contrary to fact—development economists had been interested in evaluating Ethiopia's community-based health program with a cluster RCT. And suppose further that these development economists had the authorization and full cooperation of the Ethiopian government to do so. The immediate ethical reason for why it would be wrong for researchers to carry out this experiment is that it would make them complicit with the actions of an authoritarian regime—a compromise of any researcher's ethical integrity. But something needs to be said about *why* complicity with an authoritarian regime constitutes a breach of researcher's ethical integrity. It is also important to address the worry that preceding framework for policy research without consent could systematically permit such experiments to go forward.

Implementing community-based health programs is within the sphere of control which governments possess. After all, a similar program had been implemented in Kerala without complaints. All other things being equal, it is safe to say that a randomized evaluation of the program would have been permissible to carry out in India. Yet the reason why a randomized evaluation of the program would be permissible to carry out in India but not Ethiopia has something to do with the government authorizing the experiment. While the hypothetical experiment would not violate the autonomy rights of anyone if authorized by a legitimate government, the same cannot be said of an identical experiment authorized by an authoritarian regime.

While the autonomy rights are important, there are even more fundamental rights at stake in the development context. Following Christiano (2011), I maintain that there is a human right to democratic governance. This right provides the basis for the claim that research without consent authorized by an authoritarian regime is invalid. If such an experiment were to proceed, researchers would be doing something wrong even if granted government authorization.

The idea that there is a human right to democracy is controversial.[82] However, a strong instrumentalist case for such a right has been made by Christiano (2011). Citing plausible empirical considerations, Christiano argues that democracies have a better track record of effectively protecting what he calls *rights of personal integrity*. Such rights are the most fundamental human rights and include "the right to not be tortured, the right to not be arbitrarily imprisoned, and the rights to not be murdered or disappeared by the state" (Christiano, 2011, p.

[82] Rawls (1999, pp. 71-81) notably argues that there is no such human right. Christiano (2011) also identifies two commonly cited arguments against positing a human right to democracy. The first is that a human right to democracy would conflict with the rights persons have to collective self-determination. The second is based on the observation that new democracies often violate the human rights of citizens.

145). The human right to democracy derives from the contingent fact that rights of personal integrity are routinely violated in non-democratic societies.

Christiano's argument for a human right to democracy implies that currently existing governments have duties to adopt or maintain democratic institutions. But since it would be unduly demanding to require developing countries to meet the standards of democratic governance found in high-income countries, Christiano carefully notes that his argument only establishes the right to a *minimally egalitarian democracy*. Christiano's notion of a minimally egalitarian democracy can provide the basis for a criterion of minimal legitimacy. As I suggest below, when it comes to policy experimentation without consent, only authorization from a government that meets such criteria can be considered valid.

The central idea motivating minimally egalitarian democracy is that all persons should be able to equally partake in the collective decision making of their political society. Christiano identifies three conditions for minimally egalitarian democracy, which are worth reproducing in their entirety.

> (1) Persons have formally equal votes that are effective in the aggregate in determining who is in power, the normal result of which is a high level of participation of the populace in the electoral process.
> (2) Persons have equal opportunities to run for office, to determine the agenda of decision making, and to influence the processes of deliberation. Individuals are free to organize political parties and interest group associations without legal impediment or fear of serious violence, and they are free to abandon their previous political associations. They have freedom of expression at least regarding political matters. In such a society, there is normally robust competition among parties and a variety of political parties that have significant presence in the legislature.
> (3) Such a society also acts in accordance with the rule of law and supports an independent judiciary that acts as a check on executive power (Christiano, 2011, p. 146).

Christiano emphasizes that a minimally egalitarian society need not be fully just, nor will it fully live up to the ideals of democracy. Importantly, Christiano also suggests that such a society need

not be robustly liberal in the sense that familiar autonomy rights (which he calls liberal rights) unconnected to democratic processes will be fully protected. However, this should not give development economists license to conduct research that violates the unprotected autonomy rights of individuals in developing countries. Following the tradition going back to John Locke's *Second Treatise of Government* (Locke, 1980), I understand autonomy rights as existing pre-institutionally. Researchers in developing countries still have duties to not violate such rights even if they are not recognized by law or convention.

There are governments in the developing world that clearly do not meet the standards of minimal legitimacy necessary for private actors to conduct permissible policy experimentation without consent. This raises an important question: Are development economists conducting experiments in countries run by illegitimate governments? It would be impossible to answer this question without an extensive review of every experiment conducted in the past twenty years. However, the experiments conducted by the most prominent development economists have been concentrated in two regions: Kenya and India. Cohen and Dupas's (2010) previously discussed experiment took place in Kenya along with other well-known studies such as Miguel and Kremer (2004) evaluation of a school deworming program. Many of Banerjee and Duflo's experiments have also taken place in Kenya and as well as India (see Banerjee and Duflo, 2011). In lieu of an exhaustive review, I focus on these two countries.

India is widely regarded as the world's most populous and multicultural democracy and should not pose any significant challenges. However, Kenya is more complicated. In recent history Kenya's governance has been autocratic, and periods of democracy in the 1990s and 2000s have been marred by ethnic violence and accusations of electoral fraud (Central Intelligence Agency, 2022). As of 2010, Kenya adopted a new constitution which introduced

substantive checks and balances on executive power. In 2013, Kenyans elected a new president. However, a 2017 election was once again marred by accusations of fraud and subsequent violence (Central Intelligence Agency, 2022). Have development economists working in Kenya obtained legitimate authorization to conduct their research then? Kenya is clearly an imperfect democracy, but a country that aspires towards democracy, nonetheless. If governments such as Kenya's are deemed illegitimate then the bar for minimal egalitarian democracy will be too high for many developing countries to meet, thereby defeating the rationale for such a criterion in the first place. Authorization from a government such as Kenya's should be considered valid because Kenya's government meets the threshold for minimal legitimacy.

7.3. Defeasibility Condition

Some will find the insistence on authorization from minimally egalitarian democratic governments too restrictive. Consider once again the hypothetical randomized evaluation of Ethiopia's community-based health program. Some will find that such an experiment could be permissible to carry out despite being authorized by an illegitimate government. The rival ethical consideration at stake here would be the expected social value of the experiment. Conducting the experiment could potentially lead to improvements in human well-being within Ethiopia, and perhaps other developing areas. Insisting on valid authorization would, at least in some circumstances, prevent development economists from discovering socially valuable interventions. These interventions would do nothing to further the aims of justice but could still improve the lives of those living in the most oppressive circumstances—an ethical consideration that should not be overlooked.

A familiar way of justifying exemptions to the authorization condition (as it applies in the developing world) is to acknowledge that the research violates participant's rights, but also argue

that such rights violations only constitute a *pro tanto* wrong against research participants. Gelinas et al.'s (2016) previously discussed framework could potentially be seen as accommodating this consideration. Recall that their second criteria for research without consent allows minor rights infringements when the expected value of research is high. As noted earlier, I agree with Gelinas and colleagues that rights violations can be sometimes justified. However, I do not think violations of a person's right to democracy, which on the current framework serve as a proxy for more fundamental rights of personal integrity, should ever be considered a minor rights infringement.

Instead of relying on considerations of *pro tanto* wrongdoing, I propose the following defeasibility condition as an alternative solution:

> *Defeasibility Condition*: If a legitimate government would authorize a policy experiment without individual consent, then researchers may permissibly carry out the same experiment in a territory governed by an illegitimate government when the expected social value of the research is high.

The basis for the defeasibility condition is the independent reason researchers would have to carry out the experiment due to its high social value. This independent reason provides the grounds for the researcher's claim to legitimate authority in a region where one is lacking.[83] The fact that a legitimate government would authorize the experiment is essential because it provides evidence that researchers have the relevant independent reason, and therefore claims to legitimacy. The basis for the defeasibility condition therefore implies that if researchers could conduct the high social value experiment without the authorization or cooperation from an illegitimate government, then they would be permitted to do so. This is because an illegitimate government's authorization adds nothing to the ethical permissibility of conducting a policy

[83] This consideration is very much influenced but not committed to Raz's (1986) instrumentalist theory of legitimacy.

experiment. Hence, the proposed defeasibility condition still accommodates the intuition that authorization from an authoritarian regime should never be considered valid.

The defeasibility condition could potentially justify a randomized evaluation of Ethiopia's community-based health program. This is because, as mentioned before, a similar program had already been implemented in the Indian state of Kerala. If the regional government of Kerala permissibly conducted a randomized evaluation of their community-based program, then the proposed defeasibility condition suggests that development economists would have been allowed to implement a similar research protocol in Ethiopia irrespective of the Ethiopian government's authorization. Of course, researchers often need government authorization and cooperation for practical purposes, e.g., to conduct an experiment without interference from security forces. But the main point is that this authorization carries no ethical weight.

There is an important caveat that needs to be made with respect to the proposed defeasibility condition. The defeasibility condition provides an important exemption to the authorization condition in cases where research has high expected social value. However, the social value calculus should not just focus on what can be immediately learned. Easterly's (2013) criticism of previous Western development experts should not be disregarded. The social value calculus needs to also consider the possibility that conducting an experiment may legitimate an authoritarian regime. It may be true that there is something of immense social value to be learned in an experiment conducted in a territory ruled by an autocrat. But the value of this knowledge also needs to be traded-off against the potential legitimating effects that accompany complicity with the actions of authoritarian regime.[84]

[84] See Temkin (2022) for further discussion of the tension between consequentialist and deontological criteria when deciding how to respond to those most in need.

7.4. Accountability Condition

An important refinement made to MacKay and Chakrabarti's (2019) framework was that research without consent may be permissible if there is a minimal risk of harm involved to members of a cluster. Even if careful steps are taken to ensure this condition is complied with, there are bound to be unforeseeable harms to members of a cluster that are the result of conducting a policy experiment. For example, imagine again an experiment on cash transfers results in jealously among non-participants, which thereby leads to conflict and violence. Even if researchers and IRB members take careful steps to ensure members of a cluster are only exposed to minimal risks, there is still the possibility of significant but unforeseeable harms to members of a cluster.

One may argue that this is a consideration in high-income countries as well. But when policy experiments in liberal democratic countries are conducted by government actors (or in close collaboration with government actors) there are accountability mechanisms in place. Most notably, those negatively affected by a policy experiment can make their voices be heard at the ballot box. And if the harm is severe, there are legal institutions in place that negatively affected parties can make use of to seek compensatory justice. Similar considerations apply to policy experiments conducted by private actors. Philanthropic foundations and non-profit research firms are still legal entities that can be easily identified and held accountable for harms either in criminal or civil court.

Matters are more complicated in the developing world. The accountability mechanisms that residents of high-income countries can use to seek compensation for harms may simply not exist. Legal institutions may be corrupt or simply too weak to properly hold foreign-based private actors accountable for harms. It is for this reason that international aid agencies and the

non-profit research firms they hire need to identify means of making themselves accountable. It would be too facile to simply insist that private actors should devote resources to investigating and subsequently compensating individuals harmed by a policy experiment. Of course, this is true. But private actors also need to take steps *before* an experiment goes forward to ensure that they can compensate those negatively affected by an experiment. In other words, private actors need to meet an *accountability condition*:

> *Accountability Condition*: Before an experiment goes forward, private actors must take steps to ensure they can compensate anyone negatively affected by a policy experiment.

To be clear, the accountability condition is an ethical requirement, not a legal requirement. It is necessary to impose such an ethical requirement on private actors precisely because the legal institutions that *should* hold private actors accountable are not in place.

While determining which self-imposed mechanisms would be successful is largely an empirical matter, one possible way for private actors to comply with the accountability condition is by requiring them to take out "experiment insurance". The idea is that private actors should voluntarily agree to pay a premium to a non-profit third party tasked with investigating and compensating those negatively affected by a policy experiment. While the details of such a proposal need to be carefully worked out,[85] it at least illustrates the kind of steps private actors conducting policy experiments need to take to meet the accountability condition.

[85] One notable problem with the proposal is that it may suffer from moral hazard, thereby leading to riskier experiments taking place. This would defeat the whole point of having experiment insurance. Another problem is that it may be difficult to reliably investigate harm claims, which would incentivize participants and non-participants to make false claims in search of a payout. This could discourage potential insurers from providing the service. It could also lead to insurers charging exuberant premiums, which would then discourage experimentation and policy innovation.

7.5. Transparency Condition

Regardless of whether experiment insurance is a viable option, it is imperative that private actors conducting experiments not only make themselves accountable, but also make it known that they are accountable. In other words, private actors conducting experiments need to make themselves publicly accountable. This is not only important from the point of view of compensatory justice. As I argue below, making meeting the accountability condition in a public manner is important because it promotes trust and transparency in the policy research enterprise.

Recall that Gelinas et al.'s (2016) framework for research without consent involves an impracticability condition. The motivation behind the impracticability condition was instrumentalist: obtaining individual consent when it is practicable promotes public trust and transparency in the research enterprise—an important social good. MacKay and Chakrabarti (2019) make a plausible case that in ideal conditions the impracticability condition is redundant since, by stipulation, ideal conditions are those in which researchers comply with practices and principles designed to protect participants rights and interests, and this is publicly known to all. It is not clear if MacKay and Chakrabarti would extend this line of reasoning to privatized policy experiments in high-income countries. But regardless, it is important at this stage to emphasize the importance of the impracticability condition now that our inquiry is embedded in non-ideal circumstances.

Despite the importance of the impracticability condition, there is a lot more to say with respect about promoting public trust and transparency. As MacKay and Chakrabarti (2019, p. 190) also point out, the impracticability condition is just one possible way of promoting public trust and transparency in the policy research enterprise. Private actors conducting policy experiments in the developing world could meet the impracticability condition and still be met

with suspicion and distrust. In fact, private actors could meet all the preceding ethical conditions and still be met with suspicion and distrust. Such suspicion and distrust would not be entirely unreasonable given the legacy of Western colonialism and imperialism in the developing world. It is for this reason that private actors conducting research without consent need to meet a more general *transparency condition*.

> *Transparency Condition*: Private actors conducting policy experiments must take steps to promote public trust and transparency in the research enterprise.

Below I discuss some steps that private actors should take to meet the transparency condition.

First, the impracticability condition should be understood as necessary but not sufficient for meeting the more general transparency condition. Another important step was already mentioned in the discussion of accountability. Private actors conducting policy experiments have compensatory justice-based reasons to abide by the accountability condition, but they also have an instrumental reason to make it be known that they abide by the accountability condition. By making themselves publicly accountable, private actors also take steps to promote public trust and transparency.

Though obtaining individual consent may be impracticable in many policy experimentation contexts, it is still possible to promote public trust and transparency by taking measures to inform members of a cluster that an experiment is taking place. McRae et al. (2011) suggest members of a cluster can still be made aware of an experiment via flyers or letters.[86] Another important practice that McRae et al. (2011) suggest is identifying a contact person that can answer questions about the experiment. Adopting such measures would presumably be sufficiently low-cost that it would not make an experiment too expensive to carry out. And

[86] However, McRae et al. (2011) fail to consider that this proposal may lead to contamination bias.

importantly, it would make members of a cluster feel that their lives are not being arbitrarily interfered with, even when such feelings are not warranted.

MacKay and Chakrabarti (2019, p. 190) also suggest that making experimental results available to the public is another means of promoting public trust and transparency. Asiedu et al. (2021) point out that providing feedback to research participants and communities reduces the likelihood that members will feel exploited by researchers. But as Asiedu et al. (2021) also note, caution must also be taken when making results of an experiment publicly available. Misunderstanding experimental results can cause distress to research participants and communities, and this could cause potential research participants and communities to become reluctant to partake in additional experiments in the future.

The measures discussed above are by no means exhaustive. But this does not mean that private actors conducting policy experiments can neglect the transparency condition. Researchers conducting policy experiments in the developing world cannot afford to lose the public's trust and should take all reasonable steps in making their aims transparent.

8. Conclusion

Though cluster RCTs often lack individual consent, this does not mean that these kinds of experiments are impermissible to carry out. Though I have argued that policy experiments conducted by private actors raise additional ethical challenges, I believe these challenges have been answered. To summarize, I have argued that in high-income countries, private actors may permissibly conduct a cluster RCT without consent if the following conditions are met:

1) The private actor has the (a) right to implement the intervention of interest or (b) authorization from the appropriate government agency or social institution to implement the intervention.
2) The data collection does not involve the violations of participant's autonomy rights.
3) The research involves minimal risk of harm to participants and non-participants.

When conducting research in the developing world, I have suggested three additional conditions for privatized policy experimentation without consent: (1) the minimal legitimacy condition; (2) the accountability condition; and (3) the transparency condition. Combined with the framework for privatized policy research without consent in high-income countries, these conditions are jointly sufficient in establishing the permissibility of research without consent in developing countries. The use of cluster RCTs therefore poses no significant ethical challenge to experimental development economics.

Chapter 5: Policy Experiments, Philanthropy, and Public Goods

1. Introduction

In the previous chapter, I argued that, if certain conditions are met, private actors may permissibly conduct research without individual consent. In this chapter, I motivate and respond to a worry with my proposed authorization condition—the crux of my framework for research without consent by private actors. The authorization condition is once again reproduced below:

> *Authorization Condition*: A private actor may permissibly conduct a cluster RCT on an intervention it does not have an institutionally specified right to implement if the private actor has authorization from the appropriate government agency or social institution.

Recall that I arrived at the authorization condition by combining two more general ethical principles:

> *Intervention Without Consent*: In a research context, *X* may permissibly administer intervention *I* in population *P* if *X* possesses the right to implement intervention *I* in *P* in a non-research context.
>
> *Transfer of Right:* if *X* has a right to *Ø*, then *X* has a right to authorize *Y* to *Ø* on *X's* behalf.

At its core, the focus of this chapter is *Transfer of Right*. Specifically, I focus on *Transfer of Right* as applied to a government's right to rule over policy matters. As we see below, the central issue with the authorization condition is that it assumes that a government (or government agency) can transfer its right to rule to private actors. The authorization condition rests on a shaky foundation without further justification of this key premise.

Though the starting point for this chapter (and the previous) is the worry that policy experiments are done without individual consent, this chapter goes beyond traditional concerns in research ethics. What I refer to as "privatized policy research" raises a deeper worry about the political rights of research participants and the nature of political legitimacy, more generally. To

motivate this worry, I draw on Cordelli's (2020) recent work on the ethical and political problems with private entities taking up the roles of governments. One of Cordelli's central concerns with the privatization of government functions is that it grants private actors a quasi-legislative authority to determine what a person is owed—and how much they owe—as a matter of justice. I similarly argue that because policy experiments conducted by private actors have the potential to influence public policy—and thereby affect a person's entitlements—something more needs to be said about what constitutes legitimate authorization to conduct policy research without consent. What should worry us about research without consent is not that private actors are meddling in the daily affairs of research participants. Rather, what should worry us is that this research may constitute an unjustified form of political coercion.

Despite lending credence to the worry that private actors exercise power illegitimately by conducting policy experiments without consent, I ultimately argue that this worry can answered. To answer the worry, I draw on Reich's (2019) qualified defense of philanthropic foundations. Reich argues that philanthropic foundations can produce public goods that government agencies and private businesses are often incapable of producing. Because private actors—such as philanthropic foundations and non-profit research firms—play a unique role in promoting a key government function, I similarly argue that these private actors are legitimately authorized to conduct research without consent.

The rest of this chapter is structured as follows. In section 2, I summarize Cordelli's (2020) analysis of privatization. In section 3, I draw on examples that show how Cordelli's objections to privatization also extends to the authorization condition, and privatized policy research, more generally. Section 4 further motivates the worry with privatized policy research by bringing philanthropic foundations into the fray. In section 5, I outline Reich's (2019) first-

best justification for privatized policy research, and then present my second-best justification. Section 6 concludes with some avenues for future research on the ethics of development RCTs.

2. What's Wrong with Privatization?

Up until recently, philosophers have had little to say about privatization. This is surprising given that governments have increasingly turned to private actors to carry out key government functions, e.g., education, criminal punishment, military defense, adjudication, healthcare, and welfare provision. I do not mean to suggest that privatized policy research is necessarily part of this broader phenomena. Scientific and social scientific research is not typically regarded as an essential government function. After all, governments routinely fund private (and public) research universities. However, as the discussion below reveals, worries about privatization form the basis to a worry with the authorization condition.

Chiara Cordelli's *the Privatized State* (Cordelli, 2020) is the most developed philosophical account of what is wrong with private actors taking up the roles of governments. At the heart of Cordelli's account is a concern with political legitimacy, which as we saw in the last chapter is the key determinant of whether a government has a right to perform policy research without consent. What is notable about Cordelli's diagnosis of the ills of privatization is that it may imply that it not possible for private agents to conduct research without consent without violating the political rights of research participants—a possibility I expand upon in the next section.

Cordelli uses the example of Stuart—a random person in the street—to illustrate how privatization is connected to the concept of political legitimacy. Suppose Stuart were to walk up to you claim the right to determine how much compensation is owed to a third party you allegedly injured or whether a costly medical procedure you need will be publicly subsidized.

Even if Stuart is reasonable in his determinations, we would still think that Stuart has no right to make the determinations he claims he has a right to. As Cordelli puts it, "who is Stuart to determine what you owe and what you are owed?" (Cordelli, 2020, p. 7). Cordelli argues that the reason it would be wrong for Stuart—or anyone else—to act on these determinations is that Stuart lacks the appropriate moral authority to change your *normative situation*, i.e., "to impose new obligations on you or to determine what you are entitled to as a matter of justice" (Cordelli, 2020, p. 8). In other words, Stuart's claims lack legitimacy. But unlike Stuart, a (legitimate) government has the authority to determine your obligations and entitlements. This can be seen as an implication of the general right to rule that governments possess.

Cordelli likens privatization to a system of governance where private actors, such as Stuart, perform entitlement-affecting actions that lack political legitimacy. What is wrong with such a system is that the legislative authority to determine person's obligations and entitlements comes from a private *unilateral will*. For Cordelli, a democratic system of governance renders its citizens free— and is therefore legitimate—because all citizens are co-authors of law and policy. This is another way of saying that under democratic governance, law and policy are the product of an *omnilateral will*.[87]

Cordelli is well-aware that real-world democratic governments routinely authorize privatization. It would therefore seem that private actors can acquire whatever legitimacy a democratic government has via *Transfer of Right*. However, critics of privatization have long held that there are limits on which rights and powers a government agency may validly transfer to private actors. One source of such limits is what legal scholars refer to as the "non-delegation

[87] Cordelli is influenced by Immanuel Kant's political writings (Kant, 1991). However, this aspect of Kant's political thought owes a heavy debt to Jean Jacques Rousseau's *The Social Contract* (Rousseau, 1987). The notion of an omnilateral will bears considerable resemblance to Rousseau's notion of *the General Will*.

doctrine", which some take to be a feature of any constitution with a separation of powers. As the name suggests, the doctrine is taken to imply a prohibition on any of the branches of government delegating their powers to other branches or other government institutions. Legal scholars invoke the non-delegation doctrine to argue that certain forms of privatization are unconstitutional. For present purposes, what is important is not the legal argument, but whether there is adequate philosophical justification for the doctrine of non-delegation. There are at least three philosophical justifications typically invoked,[88] but here I focus on Cordelli's "collective nonalienation" argument for the doctrine of non-delegation.

Cordelli's argument rests on the importance of individual self-rule, which I interpret to be the main ethical value which autonomy rights are meant to protect. On Cordelli's framework, these autonomy rights also provide the basis for political rights. Like other philosophers, Cordelli maintains that our capacity for self-rule cannot be abdicated. It is for this reason that we cannot permissibly waive our autonomy rights and enter a slavery contract (cf. Nozick, 1974). But as Cordelli points out, abdicating our capacity for self-rule is not necessarily an instantaneous decision such as in the case of signing of a slavery contract. There are many individual actions that in the aggregate can erode our capacity for self-rule. For example, while one glass of wine does not erode the capacity for self-rule, multiple glasses of wine every day do.

For Cordelli, the capacity for self-rule extends to collectives. This collective capacity for self-rule forms the normative basis for the democratic state. Cordelli's chief insight is that, like individuals, collectives also cannot permissibly abdicate their capacity for self-rule. Per Cordelli, a collective cannot permissibly put itself in a condition in which "the preconditions of self-rule

[88] The first approach appeals to the fiduciary obligations that governments have to their citizens. The second appeals to the people's original consent for government. The third approach appeals to the importance of separation of powers. See Cordelli (2020, pp. 122-134) for an overview and criticisms of each of these philosophical justifications of the non-delegation doctrine.

are impossible or undermined to the point that the only available choice is to obey someone else" (Cordelli, 2020, p. 14). For example, Cordelli argues that a collective cannot validly abdicate its right to self-rule by consenting to be ruled by a tyrant or a colonial power. This holds even if such rulership were benevolent and would lead to more efficient outcomes than democratic self-rule.

Cordelli argues that privatization, when widespread, contributes to the erosion of a collective's capacity for self-rule. This is because, by privatizing key government functions, persons are routinely subject to the unilateral will of private actors. When subject to a private unilateral will, private actors can alter the normative situation of persons without their having any say in the matter. But this situation is morally unacceptable, and Cordelli likens such situations to the state of nature—a state we have an ethical obligation escape. For Cordelli, defending the non-delegation doctrine is of the utmost importance since, without it, citizens of democratic societies risk returning to the state of nature.

Cordelli relies on the following general principle to ground her defense of the non-delegation doctrine:

> *The Principle of Authority Transmission:* An agent (A) who acts on behalf of a principal (P) lacks the moral authority to enter into agreements with others that P him-or herself lacks the moral authority to enter into (Cordelli, 2020, p. 134).

The principle of authority transmission is plausible. In fact, the principle of authority transmission complements rather than undermines the authorization condition. To explain the wrong of privatization, Cordelli needs a second principle, which she arrives at by taking the link between the inalienable right to self-rule and democratic governance, and then combining it with the principle of authority transmission:

> *The Principle of Collective Nonalienation*: a democratic government lacks the moral authority to pass either laws or policies that, whether individually or aggregately, amount

> to the abdication of its citizen's basic capacity for democratic self-rule (Cordelli, 2020, p. 142).

Cordelli proceeds to combines this principle with various empirical premises to argue that the systemic privatization of government functions amounts to collective abdication for self-rule, which is impermissible.

Assessing the empirical component of Cordelli's argument is not my concern in this chapter. But it is worth highlighting that a virtue of Cordelli's account of the wrong of privatization is that one-off instances of privatization are like the occasional alcoholic beverage a self-ruling individual may consume—the weekly glass of wine does not pose a threat to an individual's capacity for self-rule. By analogy, one-off instances of privatization do not threaten a collective's capacity for self-rule either. As Cordelli emphasizes, the wrong of privatization comes from its aggregate effects on a collective's capacity for self-rule.

To take stock, it is worth reemphasizing that my concern in this chapter is not privatization of government functions. My goal in reviewing Cordelli's analysis of privatization is only to motivate a related worry with the authorization condition and privatized policy research, more generally. In the next section, I explain why we should take seriously the possibility that governments are illegitimately transferring private actors their right to rule over policy, and therefore their right to conduct research without individual consent.

3. Privatized Policy Research

Cordelli argues that systematic privatization is wrong because it illegitimately grants private actors the authority to alter a person's normative situation. In this section, I motivate a similar worry: privatized policy research has the capacity to alter the entitlements of research participants. In other words, the worry is that the authorization condition impermissibly grants

private actors the power to alter research participants' normative situation without their say in the matter. The authorization condition therefore licenses an unjustified form of coercion.

The best way to motivate the worry at hand is with examples. Consider the RAND Health Insurance Experiment (Brock et al., 1983) first discussed in Chapter 1. Recall the main finding from the RAND Health Insurance Experiment:

> For most people enrolled in the RAND experiment, who were typical of Americans covered by employment-abased insurance, the variation in use across the plans appeared to have minimal to no effects on health status. By contrast, for those who were both poor and sick—people who might be found among those covered by Medicaid or lacking insurance—the reduction in use was harmful, on average (Newhouse, 2004, p. 108).

As noted, it is tempting but difficult to draw any substantive causal link between the RAND Health Insurance Experiment and the continual use (and increase) of cost-sharing for medical care in the U.S. But as also noted, one notable effect of the experiment has been documented. Newhouse and Normand (2017, p. 2161) report that, to this day, the Congressional Budget Office uses the findings from the RAND experiment to estimate the value of cost-sharing when providing value estimates of proposed healthcare-related legislation.

The link between the RAND Health Insurance Experiment and the Congressional Budget Office's estimates of cost-sharing should not be ignored. There is a case to be made that the RAND Corporation was granted quasi-legislative authority to determine U.S. citizens' healthcare entitlements. In other words, the RAND Corporation has altered the normative situation of U.S. citizens via their experimental findings.

The RAND Health Insurance experiment may exemplify a more worrying trend with privatized policy research akin to Cordelli's worry about privatization. The worry is that privatized policy experimentation systematically alters the normative situation of persons in

unacceptable ways. I provide more motivation for this worry in the next section, where I establish the link between philanthropy and policy research.

There is a drawback with arguing that privatized policy experimentation *systematically* alters the normative situation of persons. Taking this route does not explain what is wrong with research without consent—the worry animating this chapter and the previous. The line of argument I pursued above would not specifically single out research participants as the ones wronged by the RAND Corporation. After all, the experiment altered the normative situation of *all* U.S. citizens—not just the research participants involved in the study.

The worry I am advancing best applies when the research sample constitutes the population of interest, e.g., an evaluation that includes all individuals (or households) eligible for a specific policy intervention at a given time. Such cases are the exception rather than the rule, so it is best to understand the worry as becoming more salient as the research sample approximates the entire population of interest. As an imperfect example, Consider Cohen and Dupas (2010) once more. Like the RAND Experiment, there is evidence that Cohen and Dupas altered the normative situation of Kenyans. Cohen and Easterly (2009, p. 18) report that the experiment played a key role in the Kenyan government's eventual decision to fully subsidize ITNs for pregnant women. Unlike the RAND Experiment, however, Cohen and Dupas's study specifically altered the normative situation of the population from which their sample was drawn rather than the normative situation of the average Kenyan citizen. Specifically, it was pregnant Kenyan women who had their normative situation altered. Even though women (and their children) in Kenya are better off due to the Kenyan government's decision, we should not lose sight of the fact that these women's entitlement to preventative healthcare was altered by Cohen and Dupas's actions. While assessing the counterfactual is difficult, it is well-within the realm of possibility

that the Kenyan government would not have fully subsidized ITNs had it not been for Cohen and Dupas's study.

To reiterate, what I am suggesting is troubling with the lack of consent in Cohen and Dupas (2010) is *not* that the researchers violated the (unspecified) autonomy rights of research participants by, for example, manipulating their environment. As I argued in the last chapter, Cohen and Dupas were authorized by the Kenyan Ministry of Health. Presumably, there should be no objections to the Kenyan Ministry of Health carrying out Cohen and Dupas's research protocol without them. What I am suggesting is of concern is that the Kenyan Ministry of Health allowed foreign researchers to conduct an experiment on a population (Kenyan women) that altered the normative situation of this population without the population's say in the matter. Of course, there are plenty of women who did not partake in the experiment who have had their normative situation altered because of Cohen and Dupas (2010). But the main point is this: without further justification, it seems Cohen and Dupas were authorized to do something that that private actors should not have legitimate authority to do, i.e., alter someone's normative situation.

Before turning to my answer to the worry at hand, I address a deflationary response to the worry itself. An objection one may put forth is that unlike in the case of privatization of government functions, private actors still do not have the direct capacity to alter research participants' normative situation. Private actors—such as Cohen and Dupas or the RAND Corporation—only provide *evidence* that governments then factor into their decision making. Because private actors are only providing evidence that governments are free to disregard, they are not really altering anyone's normative situation. For example, consider a well-known randomized evaluation of a private school voucher program in Colombia (Angrist et al., 2002).

Despite a positive evaluation of the program by leading economists, Cohen and Easterly (2009, p. 17) report that the program was discontinued and never revived nor scaled up. In this case, private actors received appropriate authorization to evaluate a policy intervention. However, the Colombian government's authorization did not directly alter anyone's normative situation because of the experimental findings. One may reasonably suspect that much policy-relevant social scientific research shares a similar fate to the Colombian school voucher study.

One response to this objection is to point out that governments have duties to implement policies that are evidence-based (MacKay, 2020). But granting full compliance with this duty would still not show that private actors are directly changing the normative situation of persons. The best one can do is establish that private actors can indirectly alter the normative situations of persons. Private actors do this by deciding which interventions to study. By deciding which interventions to study, private actors can indirectly influence (though not determine) policy agendas. Private actors can also indirectly determine policy agendas by focusing their attention on novel interventions not yet considered by policymakers. But such indirect influence should be no immediate cause for alarm. After all, philosophers may (and hope to!) provide arguments that rationally persuade policymakers to change policy. It would similarly be unfair to argue that social scientists engage in illegitimate coercion by conducting policy-relevant research.

To further motivate the concern about undue influence of private actors in policymaking, we have to say more about how much policy research is funded. In the next section, I discuss the link between privatized policy research and philanthropic foundations. I do so to provide more motivation for the worry that policy experiments affect the normative situation of persons in ways that are concerning.

4. Policy Research and Philanthropy

4.1. Introduction

In this section, I further motivate why privatized policy research requires more careful examination. Specifically, I focus on the connection between policy research and philanthropic foundations.

There is no denying that philanthropic foundations have enormous resources at their disposal. It would be a mistake to believe that these foundations are—to borrow the title of Reich's (2019) book on the topic—*Just Giving* private wealth away. Like privatization, the influence philanthropic foundations have on national politics further reinforces concerns about illegitimate exercises of power by private actors. Reich writes that the "existence and growing power of private foundations to influence public policy sits in tension with ordinary democratic expectations of the political equality of citizens" (Reich, 2019, p. 137). Notably, one of the ways in which philanthropic foundations can influence public policy is by funding policy research.

Consider the second wave of field experimentation discussed in chapter 1. MDRC, one of the non-profit research firms that pioneered the use of randomized assignment for testing social programs in the U.S, was funded by the Ford Foundation. Not only that, but Rolston and Gueron (2013, p. 429) provide details of how the Ford Foundation provided generous funding for policy experiments after federal support for program evaluation dried up in the 1980s and 1990s. In the preface to *Fighting for Reliable Evidence*, Gueron and Rolston write:

> The value of the MDRC perspective comes not from salience of the organization per se but from what it tells about game-changing role a private organization and its funders—particularly, in our case, the Ford Foundation—can play in the public policy process. Research costs money, and this story would have been very different if, at several critical moments, Ford had not financed MDRC's random assignment experiments (Gueron and Roltson, 2013, p. xiii).

Other examples are easy to come by. While the RAND Corporation is mainly known for conducting research on behalf of the U.S. government and its various agencies, U.S. state and local governments, and non-U.S. governments agencies and ministries, RAND also conducts research that goes beyond client-sponsored work. To do this, the RAND corporation relies on funding from both government sources and philanthropic foundations (RAND, 2022). And as noted in the previous chapter, the Abdul Latif Jameel Poverty Action Lab (J-PAL), the non-profit research firm headed by Abhijit Banerjee and Esther Duflo, is largely funded by private donors.[89]

On J-PAL's list is one donor that deserves particular attention: The Bill and Melinda Gates Foundation (henceforth, Gates Foundation). In the next section, I specifically focus on the Gates Foundation as a proxy for the general worry with philanthropic foundations funding policy research in developing countries.

4.2. The Gates Foundation

A significant portion of policy research in the developing world—including the work of development economists—is sponsored by the Gates Foundation. Blunt (2022) has recently argued that the Gates Foundation—and transnational philanthropy more generally—should cause more concern among political philosophers. As Blunt makes clear, he thinks the problem with transnational philosophy is that "it gives the wealthy uncontrolled power over the basic interests of other people in a way that is difficult to reconcile with the commitments to minimal autonomy that underpin much of the distributive global justice literature" (Blunt, 2022, p. 2040). Blunt is not wrong to suggest that the Gates Foundation is a powerful organization. As Blunt (2020, p.

[89] J-PAL's website lists the following as major donors: Arnold Ventures, Co-Impact, Community Jameel, Echidna Giving, The Bill and Melinda Gates Foundation, Google.org, The William and Flora Hewlett Foundation, King Philanthropies, The John D. and Catherine T. MacArthur Foundation, The Douglas B. Marshall Jr. Family Foundation, Omidyar Network, and The Alfred P. Sloan Foundation. J-PAL also lists funding from government agencies such as the Australian Department of Foreign Affairs and Trade and the UK Foreign, Commonwealth & Development Office (J-PAL, 2022).

2041) points, the Gates Foundation spends just as much as Canada and Australia do on overseas development aid—a remarkable figure.

Addressing Blunt's concerns with the Gates Foundation is important. Throughout this dissertation, I have largely ignored questions about who funds development RCTs. And if Blunt's analysis of the Gates Foundation is correct, we should be concerned that the authorization condition—and privatized policy research more generally—violates the political rights of the global poor. To see why, we need to unpack Blunt's criticisms of the Gates Foundation.

At the heart of Blunt's critique is a concern with *domination*—the presence of arbitrary power within social relation or institutions.[90] Blunt argues that the Gates Foundation's research funding activities limits freedom qua non-domination. Though Blunt never identifies who *exactly* is being dominated by the Gates Foundation, he does suggest that it is the global poor who are most affected by the Gates Foundation's power (Blunt, 2022, p. 2055). This is a consideration worth taking seriously. If the Gates Foundation dominates the global poor with its research funding activities, then governments may have obligation to not authorize Gates-funded research projects in their countries.

Blunt's argument proceeds as follows. First, he identifies four features of domination:

1. It must occur in a social relationship or social institution;
2. There must be an asymmetric power division;
3. This division of power must be sufficient to produce dependency;
4. This power must be arbitrary in at least one of two senses; its use rests on the judgment of the more powerful agent, and /or the terms of social cooperation must be incontestable by the weaker agent (Blunt, 2022, pp. 2041-2042).

[90] The concern with domination is also reflected in Cordelli (2020). When Cordelli argues that the wrong of privatization is that it subjects persons to the unilateral will of private actors, what she is ultimately concerned with is domination. This should come as no surprise given Cordelli's Kantian influences. Kant's political writings are often seen as part of the republican tradition—a tradition that has enjoyed a resurgence due to Pettit (1997).

Blunt claims these features are necessary condition for domination. But presumably, what he means is that they are sufficient for domination since his thesis is that the Gates Foundation poses a problem for political philosophy, and because his argumentative strategy involves showing that the Gates Foundation meets all four conditions. Conditions 1 and 2 are met without controversy: philanthropic activities are shaped by legal and political forces and are (arguably) a social institution (see Reich, 2019). It is also hard to maintain that the Gates Foundation is *not* extremely powerful.

In addressing Blunt, I will set aside condition 3 focus on condition 4. The real concern should not be dependency, but with dependency on arbitrary power. If one can show that the Gates Foundation is not wielding power arbitrarily, then Blunt's critique of transnational philosophy loses much of its bite. This is not to suggest that we should be unconcerned with the power of wealthy philanthropists. I use Blunt's analysis to refine what the real worry with the Gates Foundation is.

It is hard to deny that the Gates Foundation has tremendous influence on global health priorities. Blunt extensively documents this. Among Blunt's main complaints is that Gates has pushed global health programs towards high-tech, vaccine-focused initiatives. The funding activities of the Gates Foundation are without a doubt influenced by Bill and Melinda Gates's unilateral judgements about which interventions are valuable and worth investing in. And, unexpectedly, those who want to receive funding from the Gates Foundation need align their research agenda with Gates's. This suggests the Gates Foundation's power is arbitrary in the two senses specified above. But the question is whether this is sufficient to establish that the Gates Foundation's power is arbitrary. I suggest that it is not.

The Gates Foundation does not operate in a political vacuum. As Blunt acknowledges, philanthropy is a social institution shaped by political and legal forces (recall Blunt needs this to be true to meet condition 1). Assuming these political and legal forces are democratic, it is hard to maintain philanthropic foundations exercise power illegitimately. In fact, Blunt acknowledges that an objection to his argument is that philanthropic foundations "operate under significant external control to make their exercise of power non-arbitrary" (Blunt, 2022, p. 2054). Blunt's first response to this objection above is to point out that the Gates Foundation is only subject to oversight in the U.S., where it is headquartered, and where it is not subject to sufficient oversight. Blunt may be correct that philanthropic foundations need more oversight in the U.S. (see Reich, 2019). But Blunt also maintains that the real problem is that the Gates Foundation is unaccountable for its power overseas.

The problem with Blunt's second response is that he provides no evidence that the Gates Foundation operates in a political vacuum outside the U.S. It would be truly remarkable to discover that foreign governments have no control over the Gates Foundation's activities within their territory. This is not to deny that the Gates Foundation has considerable influence over global and domestic policy agendas—this is precisely the worry I am trying to motivate. The more plausible worry is that the Gates Foundation *is* being authorized to operate in developing countries, and that this authorization is illegitimate. The reason such authorization is illegitimate is because it grants a foreign actor—Bill Gates—the power to dictate public policy and alter the normative situation of persons in an undemocratic fashion. Given the tremendous power and influence the Gates Foundation has, the question becomes whether the governments of developing countries are doing something impermissible by allowing the Gates Foundation to operate within their borders. That is, are the governments developing countries allowing the

Gates foundation to dominate their residents? In the next section, I address this question. I argue that the governments of both high and low-income countries are justified in authorizing philanthropic foundations to sponsor or conduct policy research in their countries.

5. Privatized Policy Research—A Public Goods Justification

As the discussion in section 4 suggests, one of the mechanisms by which wealthy philanthropists can influence public policy is by financing policy experiments. Does this mean that the Reverend John Hayes Holmes was correct to characterize philanthropic foundations as "repugnant to the whole idea of a democratic society" (in Reich, 2019, p. 5)? Despite taking this concern seriously, Reich does argue that with the right institutional structure, philanthropic foundations can play an important role in promoting the aims of the liberal democratic state. Reich advances two related arguments to defend the institution of private philanthropy. Though Reich often uses the term "policy experimentation" in advancing his two arguments, Reich's central concern is not policy experimentation as understood here. Nevertheless, Reich's two arguments are general enough that they can be used to justify the authorization condition and address the worry that privatized policy experimentation is in tension with the political rights of research participants.

After reviewing Reich's first-best defense, I will highlight two important differences between Reich's argument and the version I am adopting. The first difference is that my version of Reich's public goods argument is a second-best justification. And the second difference is that my version rests on a familiar argument for the permissibility of government coercion.

5.1. The Pluralism Argument

Following standard economic theory, public goods as goods that are both non-rivalrous and non-excludable (Varian, 1992, ch. 23). Textbook examples of public goods include national defense and clean air. Unlike with private goods, market mechanisms are known to undersupply public

goods. Since it is not possible (or desirable) to prevent anyone from consuming a public good, private businesses have little incentive to supply a public good. And because individual consumers are often aware that they can enjoy a public good without contributing to its provision, consumers each have an individual incentive to free ride on the contributions of others absent any central authority individual. This results in public goods being undersupplied.

For many economists, one of the most important functions of government is public good provision (e.g., Samuelson, 1954; Stiglitz, 1987). Standard economic analysis of public goods suggests that a voting mechanism is a desirable non-market alternative to public good provision. However, as the economic analysis of public goods suggests, a democratic government's ability to provide public goods is subject to majoritarian constraints and the preferences of the median voter. Reich illustrates this point: "If a majority of citizens prefer police protection and a minority prefer arts funding, then politicians will vote to fund the police and not the arts" (Reich, 2019, p. 154). This means that governments will adequately supply highly demanded public goods while other public goods will be undersupplied or simply not supplied at all.

Reich argues that philanthropic foundations can help address both market and government failures in the supply of public goods. This is because philanthropic foundations are unaccountable to the logic of the market and the majoritarian constraints of democratic governance. With the right incentive structure, philanthropic foundations can fund minority, experimental, or controversial public goods. Reich goes onto argue that this is important because, in a liberal democratic society marked by conditions of pluralism, citizens will have diverse preferences with respect to which public goods should be supplied. Philanthropic foundations thereby offer a mechanism of decentralized public good provision.

Reich's pluralism argument faces a notable difficulty, which he acknowledges: it is not clear that the diverse preferences of wealthy philanthropists will be representative of the wider citizenry. While Reich's pluralism argument may not succeed with respect to promoting pluralism in a democratic society, the public goods aspect of his argument is worth homing in on for purposes of justifying privatized policy experimentation. Scientific research, and scientific knowledge more broadly, can often (not always) be characterized as a public good. This is one important reason governments are willing to fund private (and public) research universities. Privatized policy experimentation can likewise be seen as the use of private means for provision of an important public good that governments will undersupply, namely, high quality policy research. However, critics of privatization such as Cordelli (2020) may still argue that relying on non-government actors for public good provision creates an unacceptable form of dependency on the private wills of wealthy philanthropists. These critics could further argue that, by expanding their size and reach, governments could meet the demand for public goods. A second argument is therefore needed to demonstrate that it is permissible for a government to authorize a private actor to conduct a policy experiment.

5.2. The Discovery Argument

Philanthropic foundations are not only well-poised to supply public goods; they are also in a unique position to *discover* a specific kind of public good: policy innovation. Just how a philanthropic foundation's lack of accountability can be harnessed for the provision of unpopular public goods, the fact that foundations can operate on longer time horizons allows them serve as vehicles for policy experimentation (in the broad sense). The longer time horizon on which philanthropic foundations operate allows them to experiment (in the broad sense) with policies that have uncertain results. Such policies are potentially socially valuable in that they promote

human welfare or social justice. As Reich puts it, philanthropic foundations can serve as a democratic society's risk capital.

The discovery argument rests on a key supposition: democratic governments often do not know how to best promote human welfare or discharge duties of justice to their citizens. This applies to the present moment, but also into the further future, where social and political circumstances may rapidly change. Reich illustrates the point as follows:

> What kinds of policies and programs … will best promote educational opportunity and achievement? Some believe universal preschool is the answer, others a better school finance system, others better and pervasive opportunities for online learning. Examples easily multiply. What kinds of policies will best reduce recidivism rates in prisons or in substance abuse programs? Or what kinds of changes will reduce carbon emissions with the lowest cost to economic growth? (Reich, 2019, p. 159)

In other words, for many policy issues, the relevant expert community is often in a state of radical uncertainty. Reich argues that in constantly varying social conditions, its desirable to decentralize policy experimentation so that better (understood as promoting social value) policies are discovered and eventually implemented.

The key to the discovery argument is that philanthropic foundations operate on a longer time horizon than democratic governments. The institutional design of democratic governments often prevents them from being ideally suited to carry out policy experiments. As Reich points out,

> citizens in a system of democratic government tend to expect and prize tested and reliable outcomes in public policy. Elected representatives who allocate public funds to highly risky strategies to social problems—in the sense that the selected policy may fail in delivering any benefits at all—also run the risk of being punished at the ballot box (Reich, 2019, p. 161).

Reich further adds that "wasteful government spending tends to be deplored, and yet experimentation requires that some experiments fail if the approach is to deserve the label experimentation in the first place" (Reich, 2019, p. 161). This point should not be overlooked.

Critics of randomized policy experiments often complain that many experiments fail to discover any significant findings despite their exuberant costs (e.g., Pritchett, 2020). Yet in clinical drug development, where randomized experiments are the "gold standard", its well-known known that most drug development fails. By some estimates, the rate of failure is as high as 90% (Sun et al., 2022). The nature of scientific experimentation is that it involves repeated failure, and this is something the public may have a hard time accepting when such failures are subsidized by their tax dollars.

Philanthropic foundations are therefore uniquely poised to supply the public good of policy innovation. Their unaccountability and long time-horizon make them institutional oddities in a democratic society. But these non-democratic oddities can nevertheless be harnessed for public purposes. Democratic governments need philanthropy to supply policy innovation because democratic governments cannot meet demand for policy innovation by expanding its size and power.

Reich notes that the discovery argument is the stronger of the two arguments justifying the existence of philanthropic foundations. Though Reich is using the term "experimentation" in a broad sense, his argument is general enough to apply to the funding of policy research by philanthropic foundations. After all, one way in which philanthropic foundations contribute to the process of policy innovation is by funding policy experiments. Below, I argue that the discovery argument outlined above justifies privatized policy experimentation. I leverage this justification to further suggest that experiments conducted by private actors do not violate the political rights of research participants. Before turning to this argument, I first highlight an important difference between the argument I will advance and Reich's version of it.

5.3. The General-Theory of Second Best

Reich makes it clear that his argument is a "first-best" justification. By this, Reich means that his argument applies in ideal circumstances, i.e., not present circumstances. As Reich makes clear, for his argument to work there needs to be considerably more regulation of philanthropy and a more egalitarian distribution of income and wealth. But it would be a mistake to disregard Reich's public good justification as irrelevant to present circumstances. The mistake would stem from ignoring what economists call "the general theory of second best" (Lipsey and Lancaster, 1956). The general lesson from the theory of second best is that when an optimality condition cannot be satisfied, the next-best solution may involve changing variables away from the values that would otherwise be optimal.

Goodin (1995) is the first commentator to call attention to the importance of this theorem to political philosophy. A non-technical illustration of the general lesson from the theory of second best is provided below:

> Suppose that the car that I want as my first-best car has three attributes: it is a (1) new (2) silver (3) Rolls-Royce. But suppose that there is no such car available at the moment, and for some reason I really must acquire a car immediately. Hence I have to settle for second-best. The Theory of Second-Best cautions me that my second-best car will not necessarily be one that displays more rather than fewer of the same attributes as my first-best car. Thus, for example, the second-best from my point of view would probably be a (1) week-old (2) black (3) Jaguar rather than a (1) new (2) silver (3) Toyota, if those were the only two cars on offer. That is true, even though the Jaguar displays none of the same features as my first-best car, and even though the Toyota displays two out of three. That is precisely the point of the Theory of Second-Best (Goodin, 2012, p. 157).

What exactly does car shopping have to do with policy research? Even if various background institutions are unjust (i.e., optimality conditions are not met), the best next alternative may not be restricting the activities of philanthropic foundations or denying them authorization to conduct policy research. It may be the case that allowing privatized policy research—in its current guise—is a second-best solution. In other words, allowing philanthropic foundations to continue

funding policy research under unjust circumstances may be better than denying philanthropic foundations the power to do so. While philanthropic foundations may hold a degree of arbitrary power under the status quo, they can still supply an important public good: policy innovation in unjust circumstances. Of course, a trade-off needs to be made between public good provision and allowing some degree of arbitrary power to permeate a democratic society. But as I argue below, governments are within their rights to make this trade-off.

5.4. Political Legitimacy and Public Goods

Reich uses discovery argument to argue that the institution of private philanthropy can complement rather than supplant the aims of the liberal democratic state. The central aim Reich cites is the promotion of social justice. Applied to policy experimentation, Reich's reasoning can be construed as follows: Since (1) governments ought to promote the aims of social justice, and (2) authorizing private agents to sponsor and conduct policy research is one important way in which governments can promote the aims of social justice, it (supposedly) follows that (3) governments ought to authorize private agents to sponsor and conduct policy research. The main issue with this argument is that it does not address the central problem motivating the inquiry into the ethics of policy experiments: the political rights of research participants.[91] Governments can still promote the aims of social justice while violating the political rights of its citizens, though this point rests on the exact details of the correct theory of social justice. To avoid this thorny issue, its best to appeal to the political ideal that animated the preceding discussion: legitimacy.[92]

---

[91] This argument also has a notable defect. Just because one's goal is *G*, and *P* promotes *G*, it does not follow that one has a good reason to *P*. There may be moral constraints on pursuing *P*, or there may be an alternative *P** that promotes *G* more effectively.

[92] Another avenue for such justification is to use a hypothetical consent model akin to Rawls's (1971) original position. The idea here would be that idealized agents would appreciate the force of the discovery argument and hypothetically consent to the practice of privatized policy experimentation. But since appeals to hypothetical consent in research contexts are bound to be controversial, and since I have followed recent commentators in maintaining

A commonly cited justification for government coercion is that government is necessary for the provision of many public goods (cf. Schmidtz, 1991). One can trace the popularity of this argument among economists to Adam Smith's *The Wealth of Nations* (Smith, 1993). But it is really Thomas Hobbes who—in *Leviathan* (Hobbes, 1994)—can be interpreted as more generally arguing that the source of government legitimacy stems from its ability to solve collective action problems—such as the free rider problem—that result in public goods being undersupplied.[93] This function of government provides the basis for my justification of the authorization condition and privatized policy experimentation more broadly. I argue as follows:

> (1) If a government permissibly authorizes a private actor to perform one if its functions, then the private actor exercises power legitimately in performing that function.
> (2) Government power is legitimate when directed at the provision of a public good.
> (3) If a government cannot adequately supply a public good *P*, it is permitted to authorize a private actor supply public good *P*.
> (4) The government cannot adequately supply policy innovation.
> (5) The government is permitted to authorize private actors to supply policy innovation (from 3 and 4).
> (6) Therefore, private actors exercise power legitimately when authorized to supply policy innovation (from 1, 2, and 5).

The argument shows that the authorization condition is justifiable not just to all members of the public but more specifically to those directly affected by policy experimentation. As a matter of principle then, privatized policy research does not violate the political rights of persons. Privatized policy research also does not violate the political rights of research participants when the authorization condition is met. Government authorization to conduct a policy experiment should be considered legitimate authorization.

---

that there is no need to appeal to any form of individual consent in certain research contexts, it is best to provide an alternate, less controversial justification.

[93] Schmidtz (1991, p. 2) credits Hobbes with the public goods justification of government. The key public good that Schmidtz sees Hobbes as focusing in on is collective security or peace. See Hampton (1997, ch. 3) for discussion of this Hobbes-inspired approach to political legitimacy.

6. Directions for Further Research

This dissertation is only an initial inquiry into the ethics of development RCTs. There are still other ethical considerations to explore, and there undoubtedly are additional ethical considerations not yet identified. In this concluding section, I identify three avenues for further research.

6.1. When is Randomization Fair?

The fairness of randomization has received attention in the burgeoning literature on the ethics of policy experiments, which this dissertation is part of. As MacKay (2020) points out in his seminal article on the topic, randomization plays both an epistemic and distributive role in policy experiments. While I have argued in chapter 2 that development economists do not owe research participants more than the status quo level of resources and opportunities, research participants are still entitled to fair treatment. Broome (1990) introduces the influential idea that distributive fairness is concerned with the satisfaction of claims. Broome's (1990) theory of distributive fairness requires that claims are satisfied in proportion to their strength. The key feature of Broome's theory of fairness is that it explains why a chance mechanism such as a lottery is the fairest way of distributing a scarce indivisible good. If I am right to conclude that participants in development RCTs do not have claims to resources against researchers, then every participant's claim to a scare resource is of "zero" strength and (trivially) equal. Therefore, randomization is permissible on Broome's framework. Relatedly, Stone (2011) argues that decision-makers are permitted to use lotteries to distribute scarce goods when there is indeterminacy about who it should go to. Again, if no participant has a claim against researchers, then (trivially) there is indeterminacy about who the good should go to. Therefore, randomization via lottery is also permissible on Stone's (2011) framework.

MacKay (2020) and MacKay and Cohn (2023) have already extended the insights from Broome (1990) and Stone (2011) to the analysis of policy experiments. However, there is a different angle to the fairness of randomization. Randomization does not necessarily have to be done via lottery. For example, in Miguel and Kremer's (2004) deworming study, "randomization" was done by the school's order in the alphabet—not via lottery. As Deaton (2010, p. 446) argues, assignment via an arbitrary feature (such as alphabetization) has methodological drawbacks but may be reasonable alternative when genuine randomization is not possible. But the possibility of "quasi-randomization" or assignment via arbitrary feature raises a variant on "When Is Randomization Fair?" question. To randomize fairly, do researchers need to use a genuine randomizing mechanism—such as coin flip—or can they rely on arbitrary features—such as the alphabetical order in which a research participant's name appears—to create treatment and control groups? This kind of question has received some attention in the literature on fairness and lotteries (e.g., Kornhauser and Sager, 1988) but has not yet been extended to development RCTs and other policy experiments.

6.2. Deconsequentializing the Duty of Rescue

In chapter 2, I appealed to the duty of rescue as the basis for a minimal standard of care. I left my account of the duty of rescue underdeveloped, and an important avenue for future research involves doing further analysis on the demands the duty of rescue imposes in non-ideal circumstances. In recent work relating to many themes in this dissertation, Temkin (2022) attempts to deconsequentialize the duty of rescue. Through a series of thought experiments, Temkin elicits our intuitions about moral character to convincingly show that, when it comes to helping those in need, sometimes moral agents are not required to perform an action that is optimal as judged by the lights of consequentialism. One of the most important thought

experiments is *Sudden Epiphany* (Temkin, 2022, p. 67). Temkin asks us to imagine a reasonably well-off person confronted with a drowning child who, at the crucial moment of deciding whether to jump in or not, realizes he could save more lives by not jumping in and selling an expensive non-waterproof watch (which he cannot take off) and then donating the proceeds to an effective aid organization. Since the aid organization will save more lives, he lets the child in front of him drown.

Many readers will accept that letting one child in front of you drown is wrong even if this will lead to more innocent lives saved abroad. But the point of the thought experiment is not simply to pump the reader's intuitions against consequentialism. The thought experiment leads to a crucial insight: the ethical impulse to rescue people in need in both real and hypothetical scenarios should be disentangled from general duties of beneficence. Those who invoke Singer's (1972) Pond Example to defend philanthropic giving have routinely overlooked this point. Temkin argues that direct confrontation with those in need of rescue (such in Singer's Pond Example or Sudden Epiphany) generates a special obligation that can take precedence over the general duty to do good. Temkin still maintains, along with Singer, that duties of beneficence are invariant to geographic distance. But the crucial lesson is that the duty to rescue stems from virtue-based reasons, not consequentialist reasoning. As Temkin elegantly puts it, "human life would, overall, be sadder, bleaker, colder, less valuable, and less praiseworthy, if we lacked the virtues and spark of humanity that propel us to make costly dramatic rescues despite the inefficiencies of such rescues" (Temkin, 2022, p. 83).

A possibility to further explore in further work is whether Temkin's virtue-theoretic account of the duty of rescue can be invoked to develop a more robust account of how duty of rescue grounds a minimal standard of care. The key move is to strip the duty of rescue of its

consequentialist underpinnings and instead ground it in a virtue-theoretic framework. By conducting research in emergency settings, such as a famine, development economists would be compromising their ethical integrity—not necessarily (or in addition to) depriving research participants something they are owed.

6.3. Exploitation

Hawkins and Emmanuel (2008b) suggest that ethical concerns about exploitation are what ultimately drove the controversy over the short-term AZT trials, which I discussed in chapter 2. In the medical research ethics literature, exploitation is commonly cited as the primary wrong-making feature of an unethical research (e.g., Emmanuel et al., 2000; Miller and Brody, 2006). An avenue for future research is to investigate whether the literature on exploitation produced by medical research ethicists has any bearing on development RCTs. If this dissertation can serve as inductive evidence, it seems likely that previous work on exploitation may be helpful but will not directly shed light on development RCTs and policy experiments. One important difference that immediately comes to mind is that concerns about exploitation are most salient when medical research is funded or conducted by for-profit businesses. As we have seen throughout this dissertation, development RCTs and policy experiments are sponsored or conducted by governments and other non-profit actors, e.g., university-affiliated researchers and research firms.

Much of the research ethics literature on exploitation is also influenced by Wertheimer (1996). On Wertheimer's account, A exploits B when A takes unfair advantage of B. But there is recent work on exploitation that is connected to many of the themes in this final chapter. Vrousalis (2013) provides the following general definition exploitation: A exploits B if and only if A and B are embedded in a systematic relationship in which (a) A instrumentalizes (b) B's

vulnerability (c) to extract a net benefit from B. Vrousalis's (2022) account of exploitation is fundamentally linked to the concept domination, and an avenue for possible research is seeing if this domination-linked account of exploitation has any implications for the ethics of development RCTs and research ethics more generally.

REFERENCES

Abramowicz, M. and Szafarz, A. (2020). Ethics of RCTs: Should Economists Care about Equipoise? In F. Bédécarrats, I. Guérin, and F. Roubaud (eds.), *Randomized Control Trials in the Field of Development: A Critical Perspective* (pp. 280-292). New York: Oxford University Press.

Alderman, H., Das, J., and Rao, V. (2016). Conducting Ethical Economic Research: Complications from the Field. In G. DeMartino and D. McCloskey (eds.), *Oxford Handbook on Professional Economic Ethics* (pp. 402-422). New York: Oxford University Press.

Asiedu, E., Karlan, E., Lambon-Quayefio, M.P., Udry, C.R. (2021). A Call for Structured Ethics Appendices in Social Science Papers. *PNAS*, 118(29).

Angell, M. (1997). The Ethics of Clinical Research in the Third World. *The New England Journal of Medicine,* 337(12), 847-849.

Angrist, J.D., Bettinger, E., Bloom, E., King, E., and Kremer, M.R. (2002). Vouchers for private schooling in Colombia: Evidence from a randomized natural experiment. *American Economic Review*, *92*(5), 1535-1558.

Angrist, J. D., and Krueger, A.B. (2001). Instrumental variables and the search for identification: From supply and demand to natural experiments. *Journal of Economic perspectives*, *15*(4), 69-85.

Angrist, J. D., and Pischke, J. S. (2009). *Mostly harmless econometrics: An empiricist's companion*. Princeton: Princeton University Press.

Angrist, J. D., and Pischke, J. S. (2010). The credibility revolution in empirical economics: How better research design is taking the con out of econometrics. *Journal of Economic Perspectives*, *24*(2), 3-30.

Angrist, J. D., and Pischke, J. S. (2014). *Mastering Metrics: The path from cause to effect*. Princeton: Princeton University Press.

Aron-Dine, A., Einav, L., and Finkelstein, A. (2013). The RAND Health Insurance Experiment, Three Decades Later. *Journal of Economic Perspectives*, *27*(1), 197-222.

Backmann, M. (2017). What's in a Gold Standard? In Defence of Randomized Controlled Trials. *Medicine, Health Care, and Philosophy*, *20*, 513-523.

Baele, S. (2013). The Ethics of New Development Economics: Is the Experimental Approach to Development Economics Morally Wrong. *The Journal of Philosophical Economics, 7*(1), 2-42.

Banerjee, A.V. (2007). *Making Aid Work*. Cambridge, MA: The MIT Press.

Banerjee, A.V. and Duflo, E. (2007). The Economic Lives of the Poor. *Journal of Economic Perspectives, 21*(1), 141-168.

Banerjee, A.V. and Duflo, E. (2009). The Experimental Approach to Development Economics. *The Annual Review of Economics, 1*(1), 151-178.

Banerjee, A.V. and Duflo, E. (2011). *Poor Economics: A Radical Rethinking of the Way to Fight Global Poverty*. New York: PublicAffairs.

Banerjee, A.V. Karlan, D., and Zinman, J. (2015a). Six Randomized Evaluations of Microcredit: Introduction and Further Steps. *American Economic Journal: Applied Economics, 7*(1), 1-21.

Banerjee, A.V., Duflo, E., Glennerster, R., and Kinnan C. (2015b). The Miracle of Microfinance? Evidence from a Randomized Evaluation. *American Economic Journal: Applied Economics, 7*(1), 22-53.

Banerjee, A.V., Duflo, E., Goldberg, N., Karlan, D., Osei, R., Parienté, W., Shapiro, J., Thuysbaert, B. and Udry, C., (2015). A multifaceted program causes lasting progress for the very poor: Evidence from six countries. *Science*, *348*(6236), 1-16.

Barrett, C. and Carter, M., (2010). The Power and Pitfalls of Experiments in Development Economics: Some Non-Random Reflections. *Applied Economic Perspectives and Policy*, *32*(4), 515-548.

Bauer, P.T. (1971). *Dissent on Development*. Cambridge, MA: Harvard University Press.

Bédécarrats, F., Guérin, I., and Roubaud, F., (eds.) (2020a). *Randomized Control Trials in the Field of Development: A Critical Perspective*. New York: Oxford University Press.

Bédécarrats, F., Guérin, I., and Roubaud, F., (2020b). Editor's Introduction: Controversies around RCTs in Development. In F. Bédécarrats, I. Guérin, and F. Roubaud (eds.), *Randomized Control Trials in the Field of Development: A Critical Perspective* (pp. 1-28). New York: Oxford University Press.

Bédécarrats, F., Guérin, I., and Roubaud, F., (2020c). Microfinance RCTs in Development: Miracle or Mirage? In F. Bédécarrats, I. Guérin, and F. Roubaud (eds.), *Randomized Control Trials in the Field of Development: A Critical Perspective* (pp. 186-226). New York: Oxford University Press.

Basu, K. (2014). Randomisation, Causality and the Role of Reasoned Intuition. *Oxford Development Studies*, *42*(4), 455-472.

Beitz, C.R. (1979). *Political Theory and International Relations*. Princeton: Princeton University Press.

Blake, M. (2001). Distributive Justice, State Coercion, Autonomy. *Philosophy & Public Affairs*, *30*(3), 257-296.

Boone, P. and Johnson, S. (2009). Breaking out of the Pocket: Do Health Interventions Work? Which ones and in What Sense? In J. Cohen and W.R. Easterly (eds.), *What Works in Development: Thinking Big and Thinking Small* (pp. 55-83). Washington, DC: Brookings Institute Press.

Breza, E., Stanford, F.C., Alsan, M., Alsan, B., Banerjee, A.V., Chandrasekhar, A.G., Eichmeyer, S., Glushko, T., Goldsmith-Pinkham, P., Holland, K. Hoppe, E..., and Duflo, E. (2021). Effects of a large-scale social media advertising campaign on holiday travel and COVID-19 infections: a cluster randomized controlled trial. *Nature Medicine*, *27*(9), 1622-1628.

Brook, R. Ware, J., Rogers, W., Keeler, E., Davies, A., Donald, C., Goldberg, G., Lohr, K., Masthay, P., and Newhouse, J. (1983). Does free care improve adults' health? Results from a randomized controlled trial. *New England Journal of Medicine*, *309*(23), 1426-1434.

Brook, R., Ware, J., Rogers, W., Keeler, E., Davies, A., Sherbourne, C., Goldberg, G., Lohr, K., Camp, P. and Newhouse, J., (1984). The effect of coinsurance on the health of adults: results from the RAND Health Insurance Experiment. *RAND Health Insurance Experiment Series.* Santa Monica: The RAND Corporation.

Broome, J. (1990). Fairness. *Proceedings of the Aristotelian Society, 91,* 87-101.

Card, D. and Krueger, A. B. (1994). Minimum wages and employment: a case study of the fast-food industry in New Jersey and Pennsylvania. *American Economic Review*, *84*(4), 772-793.

Cartwright, N. (2007). Are RCTs the Gold Standard? *BioSocieties, 2*(1), 11-20.

Cartwright, N. (2010). What are Randomised Controlled Trials Good For? *Philosophical Studies, 147*(1), 59-70.

Cartwright, N. and Munro, E. (2010). The Limitations of Randomized Controlled Trials in Predicting Effectiveness. *Journal of Evaluation in Clinical Practice, 16*(2), 260-266.

Cartwright, N. and Hardie, J. (2012). *Evidence-Based Policy: A practical guide to doing it better*. New York, NY: Oxford University Press.

Central Intelligence Agency, (2022). Kenya. In *The World Factbook*. Available at: https://www.cia.gov/the-world-factbook/countries/kenya/ (Accessed: 20 July 2022).

Chalmers, T.C. (1978). The Ethics of Randomization as a Decision-Making Technique and the Problems of Informed Consent. In Beauchamp, T. and Walters, L. (eds), *Contemporary Issues in Bioethics* (pp. 426-249). Encino, CA: Dickenson.

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

Christiano, T. (2011). An Instrumental Argument for a Human Right to Democracy. *Philosophy & Public Affairs*, *39*(2), 142-176.

Cohen, G.A. (1997). Where the Action Is: On the Site of Distributive Justice. *Philosophy & Public Affairs, 26*(1), 3-30.

Cohen, J., and Dupas, P. (2010). Free Distribution or Cost-Sharing? Evidence from a Randomized Malaria Prevention Experiment. *The Quarterly Journal of Economics, 125*(1), 1-45.

Cohen, J., and Easterly, W.R. (eds.) (2009a). *What Works in Development: Thinking Big and Thinking Small*. Washington, DC: Brookings Institution Press.

Cohen, J. and Easterly, W.R. (2009b). Introduction: Thinking Big Versus Thinking Small. In J. Cohen and W.R. Easterly (eds.), *What Works in Development: Thinking Big and Thinking Small* (pp. 1-23). Washington, DC: Brookings Institute Press.

Cordelli, C. (2020). *The Privatized State.* Princeton: Princeton University Press.

Crouch, R. and Arras, D. (1998). AZT Trials and Tribulations. *The Hastings Center Report, 28*(6), 26-34.

Deaton, A. (2010). Instruments, Randomization, and Learning about Development. *Journal of Economic Literature, 48*(2), 424–455.

Deaton, A. (2013). *The Great Escape: Health, Wealth, and the Origins of Inequality*. Princeton: Princeton University Press.

Deaton, A. (2020). Introduction: Randomization in the Tropics Revisited, a Theme and Eleven Variations. In F. Bédécarrats, I. Guérin, and F. Roubaud (eds.), *Randomized Control Trials in the Field of Development: A Critical Perspective* (pp. 29-46). New York: Oxford University Press.

Deaton, A. and Cartwright, N. (2018). Understanding and Misunderstanding Randomized Controlled Trials. *Social Science & Medicine,* 210, 2-21.

DeMartino, G. (2011). *The Economist's Oath: On the Need for and Content of Professional Economic Ethics*. New York: Oxford University Press.

DeMartino, G. and McCloskey, D. (eds.) (2016). *Oxford Handbook on Professional Economic Ethics*. New York: Oxford University Press.

Devoto, F., Duflo, E., Dupas, P., Parienté, W. and Pons, V. (2012). Happiness on Tap: Piped Water Adoption in Urban Morocco. *American Economic Journal: Economic Policy*, *4*(4), 68-99.

Duflo, E., Glennerster, R., and Kremer, M.R. (2004). Randomized Evaluations of Interventions in Social Service Delivery. *Development Outreach, 6*(1), 26-29.

Easterly, W.R. (2006). *The White Man's Burden: Why the West's Efforts to Aid the Rest Have Done So Much Ill and So Little Good*. New York: Penguin Press.

Easterly, W.R. (2009). Can the West Save Africa? *Journal of Economic Literature*, *47*(2), 373-447.

Easterly, W.R. (2013). *The Tyranny of Experts: Economists, Dictators, and the Forgotten Rights of the Poor*. New York: Basic Books.

Edwards, S., Braunholtz, D., Lilford, R., and Stevens, A. (1999). Ethical Issues in the design and conduct of cluster randomized controlled trials. *BMJ*, *318*, 1407-1409.

Elster, J. (2015). *Explaining Social Behavior: More Nuts and Bolts for the Social Sciences*. Cambridge: Cambridge University Press.

Emmanuel, E. (2012). Global Justice and the 'Standard of Care' Debates.' In J. Millium and E. Emmanuel (eds.), *Global Justice and Bioethics* (pp. 181-212). Oxford: Oxford University Press.

Emmanuel, E. and Miller, F.G. (2001). The Ethics of Placebo-Controlled Trials – A Middle Ground. *The New England Journal of Medicine, 345*(12), 915-919.

Emmanuel, E., Wendler, D., Grady., C. (2000). What Makes Clinical Research Ethical? *JAMA 283*(20), 2701-2711.

Fabre, C. (2006). *Whose Body is it Anyway? Justice and the Integrity of the Person*. Oxford: Oxford University Press.

Fisher, R.A. (1926). The Arrangement of Field Trials. Journal of Ministry of Agriculture of Great Britain, *33*, 503-513.

Fisher, R.A. (1935). *The Design of Experiments*. Edinburgh: Oliver and Boyd.

Freedman, B. (1987). Equipoise and the Ethics of Clinical Research. *The New England Journal of Medicine, 317*(3), 141-145.

Freedman, B. (1990). Placebo-Controlled Trials and the Logic of Clinical Purpose. *IRB: A Review of Human Subjects Research*, *12*(6), 1-6.

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

Friedman, M. (1962). *Capitalism and Freedom*. Chicago: Chicago University Press.

Gates, B. (2013). Bill Gates: My Plan to Fix the World's Biggest Problems. *The Wall Street Journal*, January 25, 2013. Available at: https://www.wsj.com/articles/SB10001424127887323539804578261780648285770 (Accessed: 20 July 2022).

Gelinas, L., Wertheimer, A. and Miller, F.G. (2016). When and Why is Research Without Consent Permissible? *The Hastings Center Report, 46*(2), 35-43.

Gertler, P. (2004). Do conditional cash transfers improve child health? Evidence from PROGRESA′ s control randomized experiment. *American Economic Review*, *94*(2), 336-341.

Gifford, F. (1986). The conflict between randomized clinical trials and the therapeutic obligation. *The Journal of Medicine and Philosophy, 11*(4), 347-366.

Gifford, F. (2000). Freedman's 'Clinical Equipoise' and 'Sliding-Scale All-Dimensions-Considered Equipoise'. *Journal of Medicine and Philosophy, 25*(4), 399-426.

Gifford, F. (2007). Pulling the Plug on Clinical Equipoise: A Critique of Miller and Weijer. *Kennedy Institute of Ethics Journal*, *17*(3), 203-226.

Glennerster, R. and Powers, S. (2016). Balancing Risk and Benefit: Ethical Tradeoffs in Running Randomized Evaluations. In G. DeMartino and D. McCloskey (eds.), *Oxford Handbook on Professional Economic Ethics* (pp. 367-401). New York: Oxford University Press.

Gneezy, U. and List, J.A. (2013). *The Why Axis: Hidden Motives and the Undiscovered Economics of Everyday Life.* New York: PublicAffairs.

Goodin, R. (1995). Political Ideals and Political Practice. *British Journal of Political Science*, *25*(1), 37-56.

Goodin, R. (2012). The Bioethics of Second-Best. In J. Millium and E. Emmanuel (eds.), *Global Justice and Bioethics* (pp. 153-177). Oxford: Oxford University Press.

Goodrich, J.P. (2021). Do We Own Our Data? The Finders-Keepers Ethics of the Cyber Commons. In P. Aligica, G. Choi, and V. Storr (eds), *Culture, Sociality, and Morality: New Applications of Mainline Political Economy*. Rowan & Littlefield.

Greenberg, D.H. and Shroder, M. (2004). *The Digest of Social Experiments: Third Edition.* Washington, D.C.: The Urban Institute Press.

Guala, F. (2005). *The Methodology of Experimental Economics*. Cambridge: Cambridge University Press.

Gueron, J. and Rolston, H. (2013). *Fighting for Reliable Evidence*. New York: Russell Sage Foundation.

Hacking, I. (1988). Telepathy: Origins of Randomization in Experimental Design. *Isis*, *79*(3), 427-451.

Hardimon, M.O. (1994) Role Obligations. *Journal of Philosophy, 91*(7), 333-363.

Hampton, J. (1997). *Political Philosophy*. Boulder: Westview Press.

Harrison, G. (2011). Randomisation and its Discontents. *Journal of African Economies*, *20*(4), 626-652.

Harrison, G. and List, J.A. (2004). Field Experiments. *Journal of Economic Literature, 42*(24), 1009-1055.

Hausman, D.M. (2011). *Preference, Value, Choice, and Welfare*. Cambridge: Cambridge University Press.

Hawkins, J. (2006). Justice and Placebo Controls. *Social Theory and Practice*, *32*(3), 467-496.

Hawkins, J. (2008a). Research Ethics, Developing Countries, and Exploitation: A Primer. In E. Emmanuel. and J. Hawkins (eds.)*, Exploitation and Developing Countries* (pp. 21-62). Princeton: Princeton University Press.

Hawkins, J. (2008b). Exploitation and Placebo Controls. In E. Emmanuel. and J. Hawkins (eds.)*, Exploitation and Developing Countries* (pp. 246-285). Princeton: Princeton University Press.

Hawkins, J. and Emmanuel, E. (eds.) (2008a). *Exploitation and Developing Countries: The Ethics of Clinical Research*. Princeton: Princeton University Press.

Hawkins, J. and Emmanuel, E. (2008b). Introduction: Why Exploitation? In E. Emmanuel. and J. Hawkins (eds.)*, Exploitation and Developing Countries* (pp. 1-20). Princeton: Princeton University Press.

Heckman, J. J. (1992). Randomization and Social Policy Evaluation. In C. F. Manski and I. Garfinkel (eds.), *Evaluating Welfare and Training Programs* (pp. 201–230). Cambridge, MA: Harvard University Press.

Heckman, J.J. (2020). Epilogue: Randomization and Social Policy Evaluation Revisited. In F. Bédécarrats, I. Guérin, and F. Roubaud (eds.), *Randomized Control Trials in the Field of Development: A Critical Perspective* (pp. 304-330). New York: Oxford University Press, pp. 304-330.

Heckman, J.J. and Smith, J.A. (1995). Assessing the Case for Social Experiments. *Journal of Economic Perspectives*, *9*(2), 85-110.

Hobbes, T. (1994). *Leviathan: with selected variants of the Latin edition of 1688.* Indianapolis: Hackett Publishing.

J-PAL (2022). About Us. Available at: https://www.povertyactionlab.org/about-j-pal (Accessed: 14 June 2022).

Kahneman, D. and Tversky, A (1979). Prospect Theory: An Analysis of Decision Under Risk. *Econometrica*, *47*(2), 263-292.

Kahneman, D. and Smith, V. (2002). Foundations of Behavioral and Experimental Economics: Daniel Kahneman and Vernon Smith. *Nobel Prize Documents, 1*(7). Available at: https://www.nobelprize.org/uploads/2018/06/advanced-economicsciences2002-1.pdf (Accessed: 5 March 2023).

Kamm, F.M. (2000). Does Distance Matter Morally to the Duty to Rescue? *Law and Philosophy*, *19*(6), 655-681.

Kant, I. (1991). *Kant: Political Writings*. Cambridge: Cambridge University Press.

Karlan, D. and Appel, J. (2011). *More than Good Intentions: Improving the Ways the World's Poor Borrow, Save, Farm, Learn, and Stay Healthy*. New York: Penguin.

Kershaw, D. (1972). A Negative-Income-Tax Experiment. *Scientific American, 227*(4), 19-25.

Kornhauser, L. A., and Sager, L. G. (1988). Just lotteries. *Social science information*, *27*(4), 483-516.

Kremer, M.R. and Holla, A. (2009). Pricing and Access: Lessons from Randomized Evaluations in Education and Health. In J. Cohen and W. Easterly (eds.), *What Works in Development: Thinking Big and Thinking Small* (pp. 91-119). Washington, DC: Brookings Institute Press.

Kukla, R. (2007). Resituating the Principle of Equipoise: Justice and Access to Care in Non-Ideal Conditions. *Kennedy Institute of Ethics Journal, 17*(3), 171-202.

LaLonde, R. J. (1986). Evaluating the econometric evaluations of training programs with experimental data. *The American Economic Review*, 604-620.

Leamer, E. E. (1983). Let's take the con out of econometrics. *The American Economic Review*, *73*(1), 31-43.

Levine, R.J. (1988). *Ethics and Regulation of Clinical Research*. New Haven: Yale University Press.

Levine, R.J. (1998). The 'Best Proven Therapeutic Method' Standard in Clinical Trials in Technologically Developed Countries. *The Hasting Center Report, 20*(1), 5-9.

Levine, R.J. (1999). The Need to Revise the Declaration of Helsinki. *The New England Journal of Medicine, 341*(7), 531-534.

Levitt, S.D. and List, J.A. (2009). Field Experiments in Economics: The Past, the Present, and the Future. *European Economic Review, 53*(1), 1-18.

Levy, S. (2007). *Progress against poverty: Sustaining Mexico's Progresa-Oportunidades Program*. Washington, D.C.: Brookings Institution Press

Lipsey, R.G. and Lancaster, K. (1956). The General Theory of Second Best. *The Review of Economic Studies*, *24*(1), 11-32.

Locke, J. (1980). *Second Treatise on Government*. Indianapolis: Hackett Publishing.

London, A.J. (2000). The Ambiguity and the Exigency: Clarifying 'Standard of Care' Arguments in International Research. *Journal of Medicine and Philosophy, 25*(4), 379-397.

London, A.J. (2001). Equipoise and International Human-Subjects Research, *Bioethics, 15*(4), 312-332.

London, A. J. (2006). The Moral Foundations of Equipoise and Its Role in International Research. *The American Journal of Bioethics, 6*(4), 48-51.

London, A.J. (2020). Equipoise: Integrating social value and equal respect in research with humans. In D. MacKay and A. Iltis (eds.), *The Oxford Handbook of Research Ethics*. Oxford: Oxford University Press.

Longino, H.E. (1990). *Science as Social Knowledge*. Princeton: Princeton University Press.

Lurie, P. and Wolfe, S.M. (1997). Unethical Trials of Interventions to Reduce Perinatal Transmission of the Human Immunodeficiency Virus in Developing Countries. *The New England Journal of Medicine, 337*(12), 853-855.

MacKay, D. (2014). Standard of Care, Professional Obligations, and Distributive Justice. *Bioethics*, *28*(7), 352-359.

MacKay, D. (2015). Standard of Care, Institutional Obligations, and Distributive Justice. *Bioethics, 29*(4), 262-273.

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

MacKay, D. (2020). Government Policy Experiments and the Ethics of Randomization. *Philosophy & Public Affairs, 48*(4), 319-352.

MacKay, D. and Chakrabarti, A. (2019). Government Policy Experiments and Informed Consent. *Public Health Ethics, 12*(2), 188-201.

MacKay, D. and Cohn, E., (2023). Public Policy Experiment Without Equipoise: When is Randomization Fair?" *Ethics & Human Research*, *45*(1), 15-28.

MacKay, D. and Iltis, A. (eds.) (2020). *The Oxford Handbook of Research Ethics*. Oxford: Oxford University Press.

MacKay, D. and Rulli, T. (2017). The Duty to Rescue and Investigators' Obligations. *Kennedy Institute of Ethics Journal, 27*(1), 71-105.

Marquis, D. (1983). Leaving Therapy to Chance. *The Hastings Center Report, 13*(4), 40-47.

McRae, A., Weijer, C., Binik, A., Grimghsaw, J., Boruch, R., Brehaut, J., Donner, A., Eccles, M., Saginur, R., White, A., Taljaard., M. (2011). When is informed consent required in cluster randomized trials in health research? *Trials*, *12*(202), 1-13.

Miguel, E. and Kremer, M.R. (2004). Worms: Identifying Impacts on Education and Health in the Presence of Treatment Externalities. *Econometrica*, *72*(1), 159-217.

Miller, D. (2011). Taking Up the Slack? Responsibility and Justice in Situations of Partial Compliance. In Z. Stemplowska and C. Knight (eds.), *Responsibility and Distributive Justice* (pp. 230-245). Oxford: Oxford University Press.

Miller, D.L. (2020). The Nature and Limits of the Duty of Rescue. *Journal of Moral Philosophy, 17*(3), 320-341.

Miller, F.G. (2012). Clinical Equipoise and Risk-Benefit Assessment. *Clinical Trials*, *9*, 621-627.

Miller, F.G. and Brody, H. (2002). What Makes Placebo-Controlled Trials Unethical?' *The American Journal of Bioethics, 2*(2), 3-9.

Miller, F.G. and Brody, H. (2003). A Critique of Clinical Equipoise. *The Hastings Center Report, 33*(3), 19-28.

Miller, F.G. and Brody, H. (2007). Clinical Equipoise and the Incoherence of Research Ethics. *Journal of Medicine and Philosophy, 32*(2), 151-165.

Miller, F.G. and Joffe, S. (2011). Equipoise and the Dilemma of Randomized Clinical Trials. *The New England Journal of Medicine*, *364*(5), 476-480.

Miller, P.B. and Weijer, C. (2003). Rehabilitating Equipoise. *Kennedy Institute of Ethics Journal*, *13*(2), 93-118.

Mills, C.W. (2005). "Ideal Theory" as Ideology. *Hypatia*, *20*(3), 165-183.

Mills, C.W. (2008). Racial Liberalism. *PMLA*, *123*(5), 1380-1397.

Mills, C.W. (2009). Rawls on Race/Race in Rawls. *The Southern Journal of Philosophy*, 47(Supplement), 161-184.

Millum, J. and Bromwich, D. (2020). Respect for Persons. In A. Iltis and D. MacKay (eds.), *The Oxford Handbook in Research Ethics.* Oxford: Oxford University Press.

Millum, J. and Emmanuel, E. (eds.) (2012). *Global Justice and Bioethics*. Oxford: Oxford University Press.

Morduch, J. (2020). The Disruptive Power of RCTs. In F. Bédécarrats, I. Guérin, and F. Roubaud (eds.), *Randomized Control Trials in the Field of Development: A Critical Perspective* (pp. 108-125). New York: Oxford University Press.

Murphy, L.B. (1998). Institutions and the Demands of Justice. *Philosophy & Public Affairs,* 27(4), 251-291.

Murphy, L.B. (2000). *Moral Demands in Nonideal Theory*. Oxford: Oxford University Press.

Nagel, T. (2005). The Problem of Global Justice. *Philosophy & Public Affairs, 33*(2), 113-147.

National Commission for the Protection of Human Subjects of Biomedical and Behavioral Research, (1978). *The Belmont Report: ethical principles and guidelines for the protection of human subjects of research.* Washington, DC: US Department of Health, Education, and Welfare.

Newhouse, J.P. (1993). *Free for all?: lessons from the RAND health insurance experiment*. Cambridge, MA: Harvard University Press.

Newhouse, J.P. (2004). Consumer-Directed Health Plans and The RAND Health Insurance Experiment. *Health Affairs*, 23(9), 107-113.

Newhouse, J.P. and Normand, S.L. (2017). Health Policy Trials. *The New England Journal of Medicine*, *376*(22), 2160-2167.

Nozick, R. (1974). *Anarchy, State, and Utopia.* New York: Basic Books.

Ogden, T. (ed.) (2017). *Experimental Conversations: Perspectives on Randomized Trials in Development Economics*. Cambridge, MA: MIT Press.

Ogden, T. (2020). RCTs in Development Economics, Their Critics and Their Evolution. In F. Bédécarrats, I. Guérin, and F. Roubaud (eds.), *Randomized Control Trials in the Field of Development: A Critical Perspective* (pp. 126-151). New York: Oxford University Press. New York: Oxford University Press.

Parker, S.W. and Teruel, G.M. (2005). Randomization and Social Program Evaluation: The Case of Progresa. *The Annals of the American Academy of Political and Social Science, 59*(1), 199-219.

Peirce, C.L. and Jastrow, J. (1885). On Small Differences in Sensation. *Memoirs of the National Academy of Sciences*, *3*, 73-83.

Petticrew, M., McKee, M., Lock, K., Green, J. and Phillips, G., (2013). In Search of Social Equipoise. *BMJ, 347*.

Pettit, P. (1997). *Republicanism: A Theory of Freedom and Government.* Oxford: Oxford University Press.

Picchio, M. (2023). Being Good in a World of Need, Larry S. Temkin. Oxford University Press, 2022, 432 pages. *Economics & Philosophy,* 1-6. doi:10.1017/S0266267123000111

Pierson, L .and Millum, J. (2018). Health Research Priority Setting: The duties of Individual Funders. *The American Journal of Bioethics*, *18*(11), 6-17.

Pogge, T. (2002). *World Poverty and Human Rights: Cosmopolitan Responsibilities and Reforms.* London: Polity Press.

Posner, R. (2006). Charitable Foundations—Posner's Comment. Available at: https://www.becker-posner-blog.com/2006/12/charitable-foundations--posners-comment.html (Accessed: 14 June 2022).

Pritchett, L. (2020). Randomizing Development: Method or Madness? In F. Bédécarrats, I. Guérin, and F. Roubaud (eds.), *Randomized Control Trials in the Field of Development: A Critical Perspective* (pp. 79-107). New York: Oxford University Press. New York: Oxford University Press.

Ravallion, M. (2014). On the Role of Aid in *The Great Escape*. *Review of Income and Wealth, 60*(4), 967-984.

Ravallion, M. (2020). Should the Randomistas (Continue to) Rule? In F. Bédécarrats, I. Guérin, and F. Roubaud (eds.), *Randomized Control Trials in the Field of Development: A Critical Perspective* (pp. 49-78). New York: Oxford University Press. New York: Oxford University Press.

Rawls, J. (1971). *A Theory of Justice*. Cambridge, MA: Harvard University Press.

Rawls, J. (1999). *The Law of Peoples with "The Idea of Public Reason Revisited."* Cambridge, MA: Harvard University Press.

Rawls, J. (2005). *Political Liberalism: Expanded Edition.* New York: Columbia University Press.

Raz, J. (1986). The Morality of Freedom. Oxford: Oxford University Press.

Reich, R. (2019). *Just Giving: Why Philanthropy is Failing Democracy and How it Can Do Better.* Princeton: Princeton University Press.

Reiss, J. (2013). *Philosophy of Economics: A Contemporary Introduction*. New York: Routledge.

Reiss, J. (2019). Expertise, Agreement, and the Nature of Social Scientific Facts or: Against Epistocracy. *Social Epistemology*, *33*(2), 183-192.

Risse, M. (2005). What We Owe to the Global Poor. *The Journal of Ethics*, *9*(1/2) 81-117.

Rodrik, D. (2009). The New Development Economics: We Shall Experiment, but How Shall we Learn? In J. Cohen and W.R. Easterly (eds.), *What Works in Development: Thinking Big and Thinking Small* (pp. 24-47). Washington, DC: Brookings Institute Press.

Rodrik, D. (2015). *Economics Rules: The Rights and Wrongs of the Dismal Science*. New York: W.W. Norton & Company.

Rothman, K.J. and Michels, K.B. (1994). The Continuing Unethical Use of Placebo Controls. *The New England Journal of Medicine 331*(6), 394-398.

Rousseau, J.J. (1987). *On the Social Contract.* Indianapolis: Hackett Publishing.

Rubin, D. B. (1974). Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of educational Psychology*, *66*(5), 688.

Rulli, T. and Millum, J. (2016). Rescuing the Duty to Rescue. *Journal of Medical Ethics*, *42*, 260-264.

Sabin, J., Mazor, K., Meterko, V., Goff, S., and Platt, R. (2008). Comparing Drug Effectiveness at Health Plans: The Ethics of Cluster Randomized Trials. *The Hastings Center Report*, 35(5), 39-48.

Sachs, J.D. (2005). *The End of Poverty: Economic Possibilities for our Time*. New York: Penguin Press.

Samuelson, P.A. (1954). The pure theory of public expenditure. *The Review of Economics and Statistics*, *36*(4), 387-389.

Schafer, A. (1982). The Ethics of the Randomized Clinical Trial. *The New England Journal of Medicine,* 307(12), 719-724.

Schmidtz, D. (1991). *The Limits of Government: An Essay on The Public Goods Argument*. Boulder: Westview Press.

Schultz, T.P. (2004). School Subsidies for the Poor: Evaluating the Mexican Progresa Poverty Program. *Journal of Development Economics, 74*(1), 199-250.

Sen, A.K. (1999). *Development as Freedom*. New York: Anchor Books.

Sen, A.K. (2009). *The Idea of Justice.* Cambridge, MA: Harvard University Press.

Shaw, L.W. and Chalmers, T.C. (1970). Ethics in Cooperative Clinical Trials. *Annals of the New York Academy of Sciences*, *169*(2), 487-495.

Siegel, A. (2008). Kantian Ethics, Exploitation, and Multinational Clinical Trials. In E. Emmanuel. and J. Hawkins (eds.)*, Exploitation and Developing Countries* (pp. 175-205). Princeton: Princeton University Press.

Sim, J. and Dawson, A. (2012). Informed Consent and Cluster Randomized Trials. *American Journal of Public Health*, *102*, 480-485.

Simmons, A. (1999). Justification and Legitimacy. *Ethics,* 104(4), 739-771.

Singer, P. (1972). Famine, Affluence, and Morality. *Philosophy & Public Affairs,* 1(3), 229-243.

Smith, A. (1993). *An Inquiry into the Nature and Causes of the Wealth of Nations* (abridged). Indianapolis: Hackett Publishing.

Smith, V. (1962). An Experimental Study of Competitive Market Behavior. *Journal of Political Economy, 70*(2), 111-137.

Sreenivasan, S. (2007). Health and Justice in Our Non-Ideal World. *Politics, Philosophy and Economics, 6*(2), 218-236.

Stiglitz, J.E. (1989). On The Economic Role of the State. In A. Heertje (ed.), *The Economic Role of the State* (pp. 9-88). Oxford: Basil Blackwell.

Stiglitz, J.E. (2003). Globalization and its Discontents. New York: W.W. Norton & Company.

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

Sun, D., Gao, W., Hu., H., Zhou, S. (2022) Why 90% of Clinical Drug Development Fails and How to Improve It. *Acta Pharmaceutica Sinica B*. DOI: https:// doi.org/10.1016/j.apsb.2022.02.002

Svensson, J. (2005). Eight Questions about Corruption. *Journal of Economic Perspectives, 19*(3), 19-42.

Temkin, L.S. (2022). *Being Good in a World of Need*. Oxford: Oxford University Press.

Temple, R. and Ellenberg, S. (2000). Placebo-Controlled Trials and Active-Control Trials in the Evaluation of New Treatments. *Annals of Internal Medicine, 133*(6), 455-463.

Thaler, R. (1980). Toward a Positive Theory of Consumer Choice. *Journal of Economic Behavior & Organization*, 1(1), 39-60.

The Committee for the Prize in Economic Sciences in Memory of Alfred Nobel (2019). Understanding Development and Poverty Alleviation. Available at: https://www.nobelprize.org/uploads/2019/10/advanced-economicsciencesprize2019.pdf (Accessed: 15 March 2021).

The Committee for the Prize in Economic Sciences in Memory of Alfred Nobel (2021). Answering Casual Questions Using Observational Data. Available at: https://www.nobelprize.org/uploads/2021/10/advanced-economicsciencesprize2021.pdf (Accessed: 5 March 2023).

U.S. Department of Health and Human Services, National Institutes of Health, and Office for Human Research Protections (2018). The Common Rule. Title 45 (Public Welfare), Code of Federal Regulations, Part 46.

Varian, H.R. (1992). *Microeconomic Analysis: Third Edition*. New York: W.W. Norton & Company.

Varmus, H. and Satcher, D. (1997). Ethical Complexities of Conducting Research in Developing Countries. *The New England Journal of Medicine, 337*(14), 1003-1005.

Valentini, L. (2012). Ideal vs. Non-Ideal Theory: A Conceptual Map. *Philosophy Compass,* 7(9), 654-664.

Valentini, L. (2021). The Natural Duty of Justice in Non-Ideal Circumstances: On the Moral Demands of Institution Building and Reform. *European Journal of Political Theory, 20*(1), 45-66.

Veatch, R.M. (2002). Indifference of Subjects: An Alternative to Equipoise in Randomized Clinical Trials. *Social Philosophy and Policy*, *19*(2), 295-323.

Veatch, R.M. (2007). The Irrelevance of Equipoise. *The Journal of Medicine and Philosophy*, *32*(2), 167-183.

Vrousalis, N. (2013). Exploitation, vulnerability, and social domination. *Philosophy & Public Affairs*, *41*(2), 131-157.

Vrousalis, N. (2020). *Exploitation as Domination.* Oxford: Oxford University Press.

Weijer, C. (2000). The Ethical Analysis of Risk. *Journal of Law, Medicine & Ethics*, *28*, 344-361.

Weijer, C. and Miller P. (2004). When are Research Risks Reasonable in Relation to Anticipated Benefits? *Nature Medicine, 10*(6), 570-573.

Wenar, L (2021). Rights. In Zalta, E. (ed.), *The Stanford Encyclopedia of Philosophy* (Spring 2021 Edition). URL = <https://plato.stanford.edu/archives/spr2021/entries/rights/>.

Wertheimer, A. (1996). *Exploitation*. Princeton: Princeton University Press.

Wertheimer, A. (2010). *Rethinking the Ethics of Clinical Research: Widening the Lens*. New York: Oxford University Press.

Widerquist, K. (2005). A Failure to Communicate: What (if anything) Can We Learn from the Negative Income Tax Experiments? *The Journal of Socio-economics*, *34*, 49-81.

Wolff, R.P. (1970). *In Defense of Anarchism*. New York: HarperCollins

World Medical Association (1975). *WMA Declaration of Helsinki – Ethical Principles for Medical Research Involving Human Subjects.* Ferney-Voltaire, France: World Medical Association.

World Medical Association (2013). *WMA Declaration of Helsinki – Ethical Principles for Medical Research Involving Human Subjects*. Ferney-Voltaire, France: World Medical Association.

Ziliak, S. and Teather-Posadas, E. (2016). The Unprincipled Randomization Principle in Economics and Medicine. In G. DeMartino and D. McCloskey (eds.), *Oxford Handbook on Professional Economic Ethics* (pp. 423-452). New York: Oxford University Press.