



When We May Not Experiment: On the Ethics of Randomized Controlled Trials in Policy and Development

Bastian Steuwer, Political Science, Ashoka University, bastian.steuwer@ashoka.edu.in

When are governments, non-governmental actors, or researchers permitted to randomly assign citizens to different policies to learn their effects? Randomized policy experiments are increasingly used in development research. Yet randomization involves treating citizens unequally. Some are denied the benefits of a policy which are extended to others. Philosophers and researchers have argued that the concern of unequal treatment can be overcome. They argue that when two policy options are on a par randomization does not treat anyone disadvantageously. In other circumstances randomization is justified as an exercise of fair supererogation. In both cases, randomization is supported by both epistemic and fairness considerations. I challenge this outlook on RCTs and provide a more restrictive perspective on the moral limits of policy experiments. I argue that, even when two policy options are on a par, unique concerns of inequality and, thereby, disadvantageous treatment can arise. I also argue that the scope of fair supererogation is limited by concerns of the costs of experimentation and conditional obligations not to experiment.



ARTICLE

WHEN WE MAY NOT EXPERIMENT: ON THE ETHICS OF RANDOMIZED CONTROLLED TRIALS IN POLICY AND DEVELOPMENT

Bastian Steuwer

When are governments, non-governmental actors, or researchers permitted to randomly assign citizens to different policies to learn their effects? Randomized policy experiments are increasingly used in development research. Yet randomization involves treating citizens unequally. Some are denied the benefits of a policy which are extended to others. Philosophers and researchers have argued that the concern of unequal treatment can be overcome. They argue that when two policy options are on a par randomization does not treat anyone disadvantageously. In other circumstances randomization is justified as an exercise of fair supererogation. In both cases, randomization is supported by both epistemic and fairness considerations. I challenge this outlook on RCTs and provide a more restrictive perspective on the moral limits of policy experiments. I argue that, even when two policy options are on a par, unique concerns of inequality and, thereby, disadvantageous treatment can arise. I also argue that the scope of fair supererogation is limited by concerns of the costs of experimentation and conditional obligations not to experiment.

I. Introduction

In November and December 2018 residents in the poor Nairobi neighborhood of Kayole-Soweto found their water connections shut. Disruptions of water had occurred before. The water system in Kayole-Soweto had been neglected and many residents had outstanding debts with the water company. Unlike before, this time the disconnections were more resolute. The water company argued that they could not continue to provide water while debts accumulated.¹ Some water connections would eventually have to be shut off.

1. Critics of the disruption argued that, even when residents paid, Kayole-Soweto received less water at poorer quality and higher prices than affluent areas of Nairobi. The debts, they argued, accumulated because the costs of upgrades in Kayole-Soweto's neglected water system were pushed onto poor residents rather than fairly shared among all Nairobi residents. Jethron Akallah et al., "When economists shut off your water," *Developing Economics*, December 11, 2023, <https://developingeconomics.org/2023/12/11/when-economists-shut-off-your-water/>.

The households to be disconnected were randomly selected. Unbeknown to the residents, they were part of an experiment by researchers to test whether and how far cutting off water to residents would force property owners (including some landlords not living in the area) to pay outstanding water bills.²

Methodologically, the Nairobi water experiment was nothing out of the ordinary. The study design followed the trend in development economics to rely on randomized controlled trials (RCTs) to test policy options. The intervention randomly divided the population into a treatment and control group. Such randomization is done to reduce selection bias. If the researchers had only observed differences in payment between residents who had their water cut off and those who did not, they would have had a harder time ruling out other causes apart from the intervention. Perhaps the company cut off residents they believed were more likely to pay anyway. Proponents of RCTs see this as a unique epistemic advantage of randomization.³ While the scientific merits and demerits of randomized experiments have been extensively discussed, the Nairobi experiment brought into focus the question of ethical standards guiding RCTs.⁴ I want to focus on a question at the heart of randomizing policy: When are governments, non-governmental actors, or researchers permitted to randomly assign citizens to different policies to learn their effects?

Even if we grant the epistemic upshots, randomizing policy raises moral concerns. After all, randomization involves treating citizens unequally.

2. Aidan Coville et al., “Financing Municipal Water and Sanitation Services in Nairobi’s Informal Settlements,” *The Review of Economics and Statistics* 107, no. 5 (2025): 1215–32, https://doi.org/10.1162/rest_a_01379.

3. See Abhijit V. Banerjee and Esther Duflo, *Poor Economics* (Random House, 2011); Dean Karlan and Jacob Appel, *More than Good Intentions* (Penguin Books, 2011); and Andrew Leigh, *Randomistas* (Yale University Press, 2018). For concerns regarding methodology, see John Worrall, “Why There’s No Cause to Randomize,” *British Journal for the Philosophy of Science* 58, no. 3 (2007): 451–88, <https://doi.org/10.1093/bjps/axm024>; Dani Rodrik, “The New Development Economics: We Shall Experiment, but How Shall We Learn?,” in *What Works in Development? Thinking Big and Thinking Small*, eds. J. Cohen and W. Easterly (Brookings Institution Press, 2009), 24–47, <https://muse.jhu.edu/book/29194>; Angus Deaton and Nancy Cartwright, “Understanding and misunderstanding randomized controlled trials,” *Social Science & Medicine* 210 (2018): 2–21, <https://doi.org/10.1016/j.socscimed.2017.12.005>.

4. Akallah et al., “When economists shut off your water.” The controversy may have caused the re-naming of the paper from the earlier, Orwellian title “‘Stay Connected’ Encouraging Water Service Payments in Nairobi’s Urban Slums” (title of the pre-registration with the *American Economic Association*, 17 November 2018, <https://www.socialscienceregistry.org/trials/3556>). For ethical concerns about the practice of RCTs more broadly, see Edward Asiedu et al., “A call for structured ethics appendices in social science papers,” *Proceedings of the National Academy of Sciences* 118, no. 29 (2021): e2024570118, <https://doi.org/10.1073/pnas.2024570118>; and Reetika Khera, “Some questions of ethics in randomized controlled trials,” *Review of Development Economics* 28, no. 4 (2024): 2072–87, <https://doi.org/10.1111/rode.12996>.

Some are denied the benefits of a policy which are extended to others. Examples include giving some children access to deworming treatments while withholding it to others;⁵ giving some anti-malaria bed nets for free while demanding a price from others;⁶ giving some but not others substantial cash transfers;⁷ and continuing water services to some households but denying it to others.⁸

Does this mean that the entire approach of randomized policy interventions is morally impermissible research? Philosophers and researchers have argued that the concern of unequal treatment can be overcome. Furthermore, as we shall see, their arguments give rather broad license to randomized policy experiments. Many of their arguments rely on the idea that oftentimes randomization is supported by both epistemic and fairness considerations.⁹

While I do not doubt that some randomized policy experiments can be permissible, I wish to challenge the rather permissive outlook on RCTs and provide a more restrictive perspective on the moral limits of policy experiments. I shall also foremostly focus on policy experiments in the context of developing countries, even though much of what I say might also apply *mutatis mutandis* to affluent countries. I begin, in Section II, by outlining a necessary condition on the procedure for randomized policy experiments. In the rest of the paper, I will turn to the question of whether randomizing policy is substantially unjust. Section III lays the groundwork for my argument

5. Edward Miguel and Michael Kremer, "Worms: Identifying Impacts on Education and Health in the Presence of Treatment Externalities," *Econometrica* 72, no. 1 (2004): 159–217, <https://doi.org/10.1111/j.1468-0262.2004.00481.x>.

6. Jessica Cohen and Pascaline Dupas, "Free Distribution or Cost-Sharing? Evidence from a Randomized Malaria Prevention Experiment," *Quarterly Journal of Economics* 125, no. 1 (2010): 1–45, <https://doi.org/10.1162/qjec.2010.125.1.1>.

7. Johannes Haushofer and Jeremy Shapiro, "The Short-Term Impact of Unconditional Cash Transfers to the Poor: Experimental Evidence from Kenya," *Quarterly Journal of Economics* 131, no. 4 (2016): 1973–2042, <https://doi.org/10.1093/qje/qjw025>.

8. Coville et al., "Financing Municipal Water." Incidentally, all of the aforementioned RCTs were conducted in Kenya, a hotspot for the new development economics. See Linda Kinstler, "How poor Kenyans became economists' guinea pigs," *The Economist*, March 1, 2024, <https://www.economist.com/1843/2024/03/01/how-poor-kenyans-became-economists-guinea-pigs>.

9. The two best philosophical papers in this spirit are Douglas MacKay, "Government Policy Experiments and the Ethics of Randomization," *Philosophy & Public Affairs* 48, no. 4 (2020): 319–52, <https://doi.org/10.1111/papa.12174> and Marcos Picchio, "What is the standard of care in experimental development economics," *Politics, Philosophy & Economics* 23, no. 2 (2024): 205–26, <https://doi.org/10.1177/1470594X231181655>. In other work, MacKay has given cautious support to the experiment shutting water supply in Nairobi. Douglas MacKay, "The Ethics of Public Policy Experiments: Lessons from Clinical Research Ethics," in *The Oxford Handbook of Research Ethics*, eds. A.S. Iltis and D. MacKay (Oxford University Press, 2024), 445–72, <https://doi.org/10.1093/oxfordhb/9780190947750.013.20>. An overview of the field is provided by the very nicely curated bibliography by Emma Cohn and Douglas MacKay available at <https://dmackay.web.unc.edu/ethics-of-field-experiments-a-bibliography/> (July 28, 2025).

and indicates two scenarios in which experimentation may be permissible. Corresponding to these scenarios, I will reconstruct two arguments for randomization which each are taken to justify a sufficient condition for RCTs.¹⁰ Both arguments seek to show how epistemic and fairness considerations are not in tension. I shall show that in each case this is true in only a limited set of cases. I start, in Section IV, by discussing an argument from parity which is invoked in favor of the first sufficient condition. I reply, in Sections V and VI, that the argument from parity overlooks an important form of inequality. In Section VII, I discuss an alternate justification which admits a tension between epistemic and fairness considerations. I point out that this alternative can justify only some RCTs. I then turn to the second scenario of possible randomization and, in Section VIII, discuss one argument which appeals to supererogation and lottery fairness and seeks to license a second sufficient condition. In Section IX, I point out that in many situations in which this condition is meant to apply, we find a conflict between fairness and epistemic considerations. In Section X, I reply to the invocation of supererogation and argue that it fails because it overlooks the presence of conditional obligations. It ultimately turns out that the license to experiment in policymaking is more limited than commonly thought.

II. Procedural Fairness: Transparent and Participatory Authorization

While I want to focus on complaints related to randomizing policies, the Nairobi experiment was criticized on various grounds. For example, the experiment did not obtain informed consent from residents. However, it is relevant that the experiment was conducted by the Nairobi Water Company, a public utility company. As such, the Nairobi Water Company has the authority to set payment policies. It has to adopt payment policies, whether randomized or not, without needing the consent of everyone subject to it.¹¹ But the mere fact that the government—or delegated agents like a public utility company—has the legitimate authority to make policy decisions does not imply that such decisions are justified. The government is still bound by

10. Technically, the conditions are sufficient to justify randomization only if the necessary condition for the procedure on policy experiments is met. They are sufficient in the sense that they are sufficient to justify the substantive policy, although they do not speak to the manner in which it was implemented. For simplicity and readability, I shall simply refer to them as sufficient conditions.

11. On this, see Douglas MacKay and Averi Chakrabarti, "Government Policy Experiments and Informed Consent," *Public Health Ethics* 12, no. 2 (2019): 188–201, <https://doi.org/10.1093/phe/phy015> and Marcos Picchio, "Policy Experiments, Informed Consent, and Democratic Authorization," *Political Philosophy* 2, no. 2 (2025): 349–78, <https://doi.org/10.16995/pp.23389>.

the duties it has towards its citizens. We should ask whether performing or authorizing the experiment entails breaching any duty of the government. There are two possible sources of complaint here.

The first complaint, which I address in this section, is procedural. A plausible ideal of legitimacy requires governments to adopt policies through transparent and participatory means of policymaking. Let us focus first on experiments by governments. One set of experiments are pure government experiments. These experiments are motivated and carried out entirely by a government entity. Insofar as the government has functioning transparent and participatory means of decision-making, such experiments raise no procedural objection. Yet this case is highly unusual in developing countries and almost exclusively found in affluent countries.

More common are embedded experiments in which public authorities work together with (overseas) researchers.¹² The Nairobi water experiment was an example of this kind. Researchers worked together with the publicly owned Nairobi Water Company. The more control is exercised by government officials, the more this type of experiment resembles a government experiment. But in developing countries, more often than not, the motivation and control lie with researchers. Tellingly, a criticism of the Nairobi experiment was titled “When *economists* shut off your water” laying the blame on the researchers behind the experiment rather than the public utility company.¹³ And in the study reporting the Nairobi experiment, the authors (academic researchers) wrote in their abstract that “we” tested the intervention.¹⁴ This raises concerns of democratic control. Researchers riding piggyback on a democratic decision is one thing, but democratically unaccountable researchers making policy is quite another thing. Angus Deaton expresses the worry as follows:

Even in the US, nearly all RCTs on the welfare system are RCTs done *by* better-heeled, better-educated, and paler people *on* lower income, less-educated and darker people. My reading of the literature is that a large majority of American experiments were not done in the interests of poor people . . . That is bad enough, but at least the domestic poor get to vote, and are part of the society in which taxpayers live and welfare operates, so that there is a feedback from them to their benefactors. Not

12. The terminology is due to Jean Drèze, “The perils of embedded experiments,” *Review of Development Economics* 28, no. 4 (2024): 2059–71, <https://doi.org/10.1111/rode.12999>.

13. Jethron Akallah et al., “When economists shut off your water,” emphasis added.

14. Coville et al., “Financing Municipal Water.”

so in economic development, where those being aided have no influence over the donors.¹⁵

Governments ceding control over policies to researchers means ceding control away from voters and the intended beneficiaries. This suggests the following condition for the permissibility of policy experiments:

Democratic Authorization Each policy experiment has to be adopted through participatory and transparent means of policymaking. For experiments by non-government actors, this requires authorization and oversight by government institutions.

In the case of government experiments, the fact that the government agency experimenting is democratically elected is not sufficient to meet *Democratic Authorization*. Electoral regimes can suffer from a democracy deficit in policymaking when decisions are made exclusively in closed door meetings or are excessively delegated. Delegated decision-makers – like a public utility company or a bureaucrat – in particular, should be especially mindful that they are subject to less direct public control. *Democratic Authorization* is admittedly a somewhat vague condition and there are interesting questions about the level of participation and transparency required in each case. I will have to postpone a more careful discussion of these issues for another occasion.¹⁶ In the following, I shall assume that *Democratic Authorization* is met in order to focus on the substantive concerns about randomization.

III. The Duties of Governments

The second possible complaint to government experiments is that they are substantively defective. Governments are required to work for their citizens. This involves achieving a variety of aims including, but not limited to,

15. Angus Deaton, “Randomization in the Tropics Revisited,” in *Randomized Control Trials in the Field of Development: A Critical Perspective*, eds. F. Bédécarrats et al (Oxford University Press, 2020), 43, <https://doi.org/10.1093/oso/9780198865360.003.0002>.

16. A part excuse is that questions of the limits of the authority of bureaucrats and delegated decision-makers in general are neglected in political philosophy. Within democratic theory, there are lengthy debates on democratic representation (democracy and the legislative) as well as the role of judges in a democracy, e.g. judicial review (democracy and the judicative), but much less debate on the role of civil servants, bureaucrats, and the like in a democratic system (democracy and the executive). My diagnosis of a lack of work on the democratic theory of the executive is shared by Joseph Heath, *The Machinery of Government* (Oxford University Press, 2020), 18–23, <https://doi.org/10.1093/oso/9780197509616.001.0001>; and Chiara Cordelli, *The Privatized State* (Princeton University Press, 2020), 85–90, <https://doi.org/10.2307/j.ctv125jsgx>.

improving their well-being, reducing inequality, respecting citizens' rights, ensuring environmental protection, promoting public safety, and so on. Finding the optimal policy will often require a trade-off between several of these target outcomes. Sometimes the trade-offs exist because of factual questions. Experimentation can then help improve policies and identify superior policies. Other times the trade-offs involve moral disagreements. Some of these trade-offs are permissible. After all, citizens can reasonably disagree. But other trade-offs are not. A policy which serves only the interests of a small elite at the expense of most of the population is not an acceptable policy.

Imagine a ranking of policies in terms of their betterness. For example, an inclusive economic policy is better than an economic policy which serves only a small elite. By "better" I mean better in achieving the aims the government ought to achieve. Betterness is a multidimensional judgment which includes an assessment of all the duties a government has with respect to justice, fairness, individual rights, social welfare, and so on.

Now add another consideration: feasibility. Governments may not be able to implement the best policy. This might be because the policy is too expensive. Or the policy requires administrative capacity or know-how which the government does not presently have. It might also be that the government agency lacks the legal authority to implement the policy, for example for reasons relating to the separation of powers or federalism.

If a policy is fully feasible such that the government can apply it to everyone, I call this a fully feasible policy option (FFO). If all policies are feasible and we only deal with FFOs, then governments are required to choose the best policy. Since the policies are ranked in terms of the goals the government ought to achieve, a government may not experiment with an inferior FFO which achieves its goals less well. However, sometimes there might be several policies, none of which is inferior to the others. In such a case, governments must choose one of the maximal FFOs. This scenario leaves room for possible experimentation (Sections IV to VII).

Something is a partly feasible policy option (PFO) if the government could apply a pilot program or, indeed, a randomized experiment in which some citizens are subjected to the policy. Non-government actors like non-government organizations (NGOs) or researchers may also pursue policies. Since they lack the legal authority to universalize policies they can usually only provide PFOs. It is possible that a PFO can be even better than the best FFOs. This opens up a second scenario in which experimentation may be permissible (Sections VIII to X).

IV. Randomization Between Maximal FFOs: The Argument from Parity and the Equipoise Condition

Recall that the core substantive objection to randomization is that it involves treating citizens unequally. Some are subjected to policy A, others are subjected to policy B. What could justify treating citizens unequally? For the moment, let us consider only randomizations between different FFOs. The first argument for randomization exploits the difference between *unequal* treatment and *disadvantageous* treatment. Not all forms of unequal treatment amount to disadvantageous treatment. One possibility is that two options are equally good. Another possibility is that it is neither true that option A is better than B, nor that B is better than A, nor that they are equally good.¹⁷ This might be because there is deep incomparability, incommensurability, or parity in terms of value. But it might also be that the evidence basis available to the policymaker does not allow the policymaker to rank one option above the other. I refer to this as parity, but the same thought could be expressed as incomparability or incommensurability.

The policy experiment on anti-malaria bed nets is arguably an illustration of such parity.¹⁸ At the time, there was a debate between distributing bed nets for free and selling them for a nominal price. The argument for free distribution was that free distribution ensures greater coverage. However, proponents of cost-sharing (selling for a nominal price) worried that recipients of free nets are less likely to use them because they attach less value to something they received for free as opposed to something they spent (some) money on.

In such a situation, a policymaker is uncertain which of the two policies is all-things-considered better and unable to rank the options. MacKay dubs this *policy equipoise*.¹⁹ In such policy equipoise, the government could not be faulted with implementing either policy outright. Both FFOs are permissible

17. For ways to spell this out, see Amartya Sen, "Interpersonal Aggregation and Partial Comparability," *Econometrica* 38, no. 3 (1970): 393–409, <https://doi.org/10.2307/1909546>; and Ruth Chang, "The Possibility of Parity," *Ethics* 112, no. 4 (2002): 659–88, <https://doi.org/10.1086/339673>.

18. Cohen and Dupas, "Free Distribution or Cost-Sharing?"

19. Douglas MacKay, "The Ethics of Public Policy RCTs: The Principle of Policy Equipoise," *Bioethics* 32, no. 1 (2018): 59–67, <https://doi.org/10.1111/bioe.12403>; MacKay, "Government Policy Experiments," 327–33; also Marcos Picchio, "What role should equipoise play in experimental development economics?," *Economics & Philosophy* 42, no. 1 (2026): 1–25, <https://doi.org/10.1017/S0266267125000033>. It also plays a role implicitly in Michelle N. Mayer et al., "Objecting to experiments that compare two unobjectionable policies or treatments," *Proceedings of the National Academy of Sciences* 116, no. 22 (2019): 10723–28, <https://doi.org/10.1073/pnas.1820701116>. Lastly, policy equipoise forms part of the reform proposal for social science research by Asiedu et al., "A call for structured ethics appendices."

for the government to implement. Neither justice nor individual rights nor any other consideration tells us that the government must implement a particular policy. The *argument from parity* then goes as follows: In such situations of policy equipoise, since it is permissible for some citizens to be subjected to A and permissible for the other citizens to be subjected to B, it should then also be permissible to randomly mix the policy.²⁰ Researchers testing out support for RCTs implicitly subscribed to that reasoning, writing that “when neither of the two policies is objectionable or perceived as clearly superior, an A/B test [i.e. RCT] comparing them should not be seen as more morally problematic than a unilateral decision to implement either untested policy.”²¹

Consider the position of citizens who are subject to policy A. At the outset, it would have been permissible for the government to apply A to them. They are treated no worse than they could have expected. They are subjected to one of two maximal FFOs. Consider now the position of citizens who are subject to policy B. The same holds for them. Each citizen, from their perspective, is subjected to a policy that would have been permissible outright and so has no grounds to complain.

No fault can be found from the perspective of the government either. The policymaker is not knowingly or foreseeably treating any citizen worse than any other. No one would be subjected to a policy deemed to be inferior—for the policymaker, the options are on a par. From its evidence-relative epistemic standpoint, there is no favoritism or preferential treatment. From this standpoint all options—policy A, policy B, or the mix of both—are on a par in terms of the government’s immediate target outcomes. But the randomly mixed policy has the epistemic advantage that we can learn from the policy. While A is no better than B, it seems that the A/B mix is better than either A or B insofar as it later on enables us to improve policy. The argument from parity suggests the following sufficient condition for randomization²²:

Equipoise Condition If policymakers are, due to epistemic uncertainty, unable to rank different policy options, they may permissibly assign citizens randomly into different groups corresponding to the multiple maximal fully

20. MacKay, “Government Policy Experiments,” 327–33.

21. Michelle N. Mayer et al., “Objecting to experiments that compare two unobjectionable policies or treatments,” 10727. See also Patrick R. Heck et al., “Objecting to experiments even while approving of the policies or treatments they compare,” *Proceedings of the National Academy of Sciences* 117, no. 32 (2020): 18948–50, <https://doi.org/10.1073/pnas.2009030117>.

22. Recall, once more, that this is only a sufficient condition for the substance of the policy, not for the manner of adopting it. I am assuming throughout that *Democratic Authorization* is met and only against the backdrop of this assumption is the proposed *Equipoise Condition* meant to be sufficient to justify the policy experiment.

feasible options (i.e. no one is assigned an option which is worse than an alternative fully feasible option).²³

V. Inequality and the Equipose Condition

I shall now go on to argue that the *Equipose Condition* is false. Take, as a starting point, a concrete and realistic example: Closing schools in response to an infectious disease outbreak like COVID-19.²⁴ We do not know (initially) how effective school closures will be. We know that they are harmful to the education of children and to parents. We know that they will be beneficial to public health, especially the health of older citizens. But we do not know how beneficial they will be to public health and whether this will outweigh the harms to children and parents. It is unknown to everyone which of the two options is all-things-considered best in fulfilling the government aims of providing quality education and safeguarding public health. The policy-maker should treat the options as being on a par. In such an uncertainty, we are in a position of policy equipose. Both the option of keeping schools open is permissible and the option of closing schools is permissible, too. The argument from parity should apply and justify a policy experiment of randomly closing some schools while keeping others open.

However, the argument from parity is fallacious. It does not follow from the fact that either option is permissible that a mix of both options is also permissible. This is because it is permissible *for everyone to be subjected* to A and *for everyone to be subjected* to B. From the perspective of a citizen, it can be one thing to be subjected to A when A is a universal policy and another thing to be subjected to A when some citizens are subject to a different policy. This is salient when the policy mix is introducing a form of inequality between the treatment and control groups.

The school example highlights a possibly troublesome inequality. Students who will continue to go to school are able to reap educational benefits that students whose schools close will not. The children in the closed school group will be disadvantaged relative to their peers. By contrast, there is no such disadvantage in either of the uniform policies. Of course, if schools closed for everyone, then this would be bad for children. But no child would be left behind by the school closure and given less adequate educational pro-

23. Cf. MacKay, "Government Policy Experiments," 329.

24. The example is one among many used in Luchuo Engelbert Bain et al., "One lesson of COVID-19: Conduct more health policy trials," *Proceedings of the National Academy of Sciences* 119, no. 24 (2022): e2119887119, <https://doi.org/10.1073/pnas.2119887119>.

vision by the state. It is only the experiment's policy mix which generates a newfound inequality between the treatment and control group in children's access to education and thereby future opportunities to succeed in life. A uniform school closure would, predictably, also lead to inequalities in educational outcomes. Some children are better equipped to learn at home than others. But the policy mix would lead to greater inequality—high performing students in the open school group would gain even larger advantages over struggling students in the closed school group. This potentially troublesome increase in inequality shows that policy equipoise, as defined, is insufficient to justify a policy experiment. We cannot reason from policy equipoise to permissible randomization.

I shall now go on to explain in greater depth why the inequality in the school experiment can speak against experimentation. What is bad about the inequality between the school children? One might mistakenly suspect that the egalitarian objection in the example appeals to *ex post* inequality. It is well-known that *ex ante* equality can sometimes lead to *ex post* inequality.²⁵ Government experiments can and often do have this implication, too. If one policy turns out to be superior to the other, then citizens subject to the better policy will benefit.²⁶ But my example is different for two related reasons.

First, the inequality in the school closing example will result even if both policies turn out to be all-things-considered equally good. That is, even if the experiment does not resolve our uncertainty about which policy is all-things-considered better, the inequality will result. The conflict between *ex ante* equality and *ex post* inequality is not what is at stake. For, and this is the second point, even from their *ex ante* standpoint the policymakers know that the different policies benefit and harm different people. The intergenerational aspect is quite visible. Keeping schools open is better for children even taking into consideration the risk of catching COVID-19. Keeping schools closed is better for older adults who will benefit from the reduced transmission of the virus. Thus, the decisionmaker knows that the children in the school closing group will be disadvantaged. There is no equipoise about *that* fact.

25. Marc Fleurbaey and Alex Voorhoeve, "Decide as you would with full information! An argument against *ex ante* Pareto," in *Inequalities in Health*, ed. N. Eyal et al. (Oxford University Press, 2013), 113–28, <https://doi.org/10.1093/acprof:oso/9780199931392.003.0009>; Itay Nissan-Rozen, "How to be an Ex-Post Egalitarian and an Ex-Ante Paretian," *Analysis* 77, no. 3 (2017): 550–58, <https://doi.org/10.1093/analys/anx111>.

26. This could be problematic, too, if the foreseen gap between the two policies is expected to be large. However, this is less salient than the concern I am raising for the following reason. If policy equipoise holds, then policymakers will not know which policy is better. If they do not know this, then it is rather unlikely that they know that one of them will be much better than the other.

Instead, the crux of the issue is the heterogeneity in policy impact. The policies are in equipoise because we do not know all the facts to resolve an inter-personal trade-off. We do not know whether educational gains outweigh public health losses or vice versa. This is what makes it impossible for policymakers to rank the different options and is enough for the *Equipoise Condition* to apply. But we know that there is a trade-off here between education and public health. This is quite different from a case in which equipoise exists because we do not know the effectiveness of a policy. The school closure example is not a case in which we simply need to know “what works.”

The presence of a heterogenous policy impact shows that the idea of parity is misleading. For children there is no parity or equality between the two FFOs; they clearly favor schools to be open. Likewise for older adults: there is no parity but a clear preference for schools to be closed. Parity only exists from the perspective of the policymaker who does not know how to resolve this trade-off.

We also know that there will be inequality between old and young whichever option is chosen. This intergenerational inequality which is inherent in both the closing and opening school policy is permissible. *Ex hypothesi*, we assumed that given our state of knowledge it would be reasonable for a policymaker to weigh this harm as justifiable given what we know about the benefits to public health. It would have been, for example, permissible to close all schools despite the adverse outcomes on children’s education.

What my example brings out however is a *different* inequality which is not present in either of the uniform policies. I am referring to the inequality between children in the treatment (school closure) and control (open schools) arm. The randomized policy mix uniquely brings about this inequality which is not present in any other policy option. This introduced inequality is better than other sources of inequality insofar as it is random and does not compound already existing injustices. But the fact that the inequality is random does not alleviate all concerns. A venerable tradition of egalitarianism is precisely concerned with people being worse off than others through no choice or fault of their own. Insofar as this inequality between treatment and control widens overall educational inequalities, it speaks against the policy mix.

Furthermore, the added inequality due to randomized policy is concerning insofar as it impacts a positional good.²⁷ For positional goods it matters

27. Indeed, my argument can be endorsed even by people who reject the intrinsic value of equality. Equality in positional goods can be justified as an implication of a general prioritarian concern. Harry Brighouse and Adam Swift, “Equality, Priority, and Positional Goods,” *Ethics* 116, no. 3 (2006): 471–97, <https://doi.org/10.1086/500524>. This does not undercut, but rather broadens my point that policymakers have good grounds to be concerned with the inequality of educational opportunities.

how much you have in comparison to others. There surely are non-positional benefits to education. Being better educated brings advantages in how to lead one's life, it brings enlightenment, and so on. But being better educated also brings the advantage to outcompete others in the labor market and for socially desirable positions. Such foreseeable inequality speaks against the randomized policy and in favor of either of the uniform policies which do not impact positional goods.

I should emphasize that the inequality in question pertains to the overall inequality between children in educational advantage. What matters is how the inequality between treatment and control impacts overall patterns of inequality.²⁸ Even without the experiment there are likely some inequalities in educational advantage. Even in highly egalitarian societies, children vary in their ability to learn through online classes. It is imaginable that the resulting inequality of the experiment would be no worse than the inequality that results from closing schools. If so, then there is no egalitarian objection here. But since the inequality introduced by the experiment is random, oftentimes it will not counterbalance other sources of inequality. In the school example, suppose that children's ability to learn from online classes falls into three groups. One third is able to learn at 20 percent of the pace, another third learns at 10 percent of the pace, a last third does not learn at all. A policy mix would mean that half of the children learn normally while a sixth each will learn at 20, 10 or 0 percent. In this version of the example, there is a substantial inequality in educational advantage which results from the experiment and thus an egalitarian reason not to experiment.

VI. Equipose, Parity, and Heterogeneous Policies

One lesson from the school closure example is that policy equipose is too coarse grained. The *Equipose Condition* applies when neither of the two FFOs is deemed all-things-considered better or worse than the other. But the school example shows that there can be distributional considerations within each of the options. My discussion highlights that policy equipose and the argument from parity work quite differently in two different situations. The argument works quite well in situations in which we simply want to find out which policy is more efficient. Suppose we are trialing whether health advice is more effective when delivered by fully qualified nurses or by medical students training to become doctors. Let us suppose that we not only do not know which option works better but also lack knowledge about how the different mechanisms will impact different people. In this case, the

28. I thank a reviewer for *Free & Equal* for pressing me to clarify the point in this paragraph.

government in fact treats no one worse than anyone else and subjects everyone to a policy which is also maximal *for them*. But this is not the case when there are trade-offs between the interests of different sets of people. Politics and policies are typically about such conflicts.

The argument from parity is, thus, more limited than it initially appears and applies only in rather special circumstances. We should reject the *Equipoise Condition* as it stands and insist that its scope is narrower. For the *Equipoise Condition* to apply we need a further condition—namely that there is no *ex ante* inequality between (subsets of the) treatment and (subsets of the) control arm. The lack of such inequality is a necessary condition for the sufficient condition of the *Equipoise Condition* to apply.

Part of the reason why such egalitarian reasons are overlooked stems from the reliance on medical research ethics frameworks. Borrowing the terminology of “equipoise” from medical research gives rise to the confused impression that inter-personal trade-offs are not at stake. In medical research it is much more common that if we do not know which of two medicines is better, we simply need to know what works best (at least for our target population). If both medicines are equally promising, then we give each patient an equally good prospect. But social science is not like medicine. In cases of policy-relevant social science, oftentimes we do not know which policy is all-things-considered better because we do not know how to resolve inter-personal trade-offs. This suggests that framing the debate on the ethics of randomizing policy in terms of medical research ethics has its drawbacks.

Another drawback is the following. Other critics of RCTs have suggested that instead of policy equipoise, we should demand clinical equipoise, mirroring medical research ethics.²⁹ This turns out to have rather restrictive implications. Many development RCTs involve treatments which are known to be effective. There is no doubt that deworming is medically advisable and will have *some* effect on school performance. Yet my argument does not

29. See e.g. Stephen T. Ziliak and Edward R. Teather-Posadas, “The Unprincipled Randomization Principle in Economics and Medicine,” in *The Oxford Handbook on Professional Economic Ethics*, eds. G. DeMartino and D.N. McCloskey (Oxford University Press, 2016), 423–52, <https://doi.org/10.1093/oxfordhb/9780199766635.013.44>. Some slightly more hedged support is given by Stéphane J. Baele, “The ethics of the New Development Economics: is the Experimental Approach to Development Economics morally wrong?,” *Journal of Philosophical Economics* 7, no. 1 (2013): 2–42, <https://doi.org/10.46298/jpe.10653>; and Michel Abramowicz and Ariane Szafarz, “Ethics of RCTs: Should Economists Care About Equipoise,” in *Randomized Control Trials in the Field of Development*, ed. Bédécarrats et al, 280–92, <https://doi.org/10.1093/oso/9780198865360.003.0012>. For the canonical formulation of clinical equipoise, see Benjamin Freedman, “Equipoise and the Ethics of Clinical Research,” *New England Journal of Medicine* 317, no. 3 (1987): 141–5 <https://doi.org/10.1056/nejm198707163170304>.

presuppose anything as restrictive as clinical equipoise. Clinical equipoise rules out, for example, studying pilot programs before universal rollout.³⁰ The problem is that this opposition of clinical equipoise and policy equipoise conceals important issues. I do not doubt that policy equipoise can sometimes justify an experiment. Instead, my argument points to the fact that policy equipoise only justifies experiments in rather special cases when policy impacts are homogenous. When policy impacts are heterogenous, by contrast, assignment into a different treatment group is often not just unequal but also disadvantageous treatment. In such cases we need to weigh the costs of inequality against gains from the experiment.

Even though my argument is novel, it connects with two criticisms made against RCTs. A unifying thread of both these arguments and my own is that we should not reduce questions of development to simple questions of “what works.” The first connection is to arguments that criticize RCTs as cohering best with utilitarian and distribution-insensitive theories of justice which focus on cost-effectiveness or efficiency alone.³¹ The reason for this is that the unique methodological advantages of RCTs pertain exclusively to *average* treatment effects. RCTs cannot provide us, by themselves, with any information about the distribution of outcomes. Any attempt to extrapolate such data requires us to use tools familiar from standard regression analysis (e.g. controlling for potential causes) which randomization was supposed to avoid. Thus, ultimately the strong epistemic advantages of RCTs are limited to situations in which we only care about average treatment effect—how well something works on average—and ignore any distributional considerations.

Second, the RCT movement has been criticized for depoliticizing development questions by portraying them as simple questions of “what works.”³² RCTs provide us with evidence for the effectiveness of measurable interventions. But they do not address larger political questions around the development process. This makes them technocratic tools which ignore political processes, including questions of democracy and freedom. My concern dovetails this criticism. The argument from parity works only if we ignore conflicts of interests, something which gives rise to politics in the first place.³³

30. Rather than the framework of equipoise, I think this is more usefully analyzed as a comparison between an FFO and PFO and thus falls within the discussion of Sections VIII to X.

31. Donal Khosrowi, “Trade-Offs Between Epistemic and Moral Values in Evidence-Based Policy,” *Economics & Philosophy* 35, no. 1 (2019): 49–78, <https://doi.org/10.1017/S0266267118000159>.

32. E.g. Jean Drèze, *Sense and Solidarity* (Permanent Black, 2017), 3–20, <https://doi.org/10.1093/oso/9780198833468.001.0001>; and Deaton, “Randomization in the Tropics Revisited.”

33. Indeed, conflicts of interests are one of the two Rawlsian circumstances of justice. John Rawls, *A Theory of Justice* (Oxford University Press, 1999, rev. edn.), 109–12, <https://doi.org/10.4159/9780674042582>

VII. Trading Off Costs to Fairness and Epistemic Gains

Thus far, I have argued that whenever we compare two FFOs and we are not encountering an unusually simple case of “what works,” we need to carefully examine any possible inequalities resulting from the experiment. We cannot justify such experiments on the basis of the *Equipoise Condition* or the argument from parity. The resulting inequality shows that the unequal treatment inherent in the experiment is not just unequal but also disadvantageous treatment. Without any countervailing moral consideration, the experiment would not be permissible. The obvious and most promising contender to justify a policy experiment despite such inequality would be to invoke the benefits which result from the experiment.³⁴

The decision whether or not to experiment can be seen as a trade-off between epistemic and fairness (or inequality) considerations. There is an egalitarian reason speaking against experimentation, but it is possible that this reason is outweighed by the knowledge gains from the experiment. Knowledge gained may, after all, help in improving policy in the future.³⁵ I certainly admit that some experiments can be justified in such a way. In my school example, if even a brief experiment of, say two weeks, could yield significant knowledge gains, then the experiment could still be justified.

But suppose that the experiment would have to continue for many months. Suppose we also know that the policy mix will create gaps between the students which later cannot be effectively remedied and that are much worse than any gaps arising from closed schools. Perhaps students generally learn poorly in online classes and by giving some students access to in-person education, they gain massive advantages. Further, suppose that the knowledge gain is modest. It is limited to COVID-19, and we expect to be able to fine-tune our policies only for a short time until mass vaccinations will have changed the situation dramatically. The results will be peculiar to COVID-19 and not be generalizable to any other health emergency. In this version of the school closure example, the benefits from the experiment could not outweigh the harms of inequality. The inequality would be disproportionate to any knowledge gains and experimentation would be impermissible. This

34. MacKay briefly hints at this possibility, too. MacKay, “Government Policy Experiments,” 352. Notably, though, this possibility only occupies a single paragraph of his article which otherwise seeks to justify experimentation on the grounds that fairness is not compromised.

35. For this, it is necessary that the knowledge created is morally valuable. Not everything that is scientifically interesting is morally beneficial. Recall the complaint by Deaton cited in Section II that many RCTs are not performed in the interests of the poor. See also Ankur Sarin, “Indecent Proposals in Economics,” *The India Forum*, October 29, 2019, <https://www.theindiaforum.in/article/indecent-proposals-economics>.

suggests that while benefits from the experiment can justify a policy experiment, they can only do so if the resulting inequalities are not disproportionate to the benefits from the experiment.³⁶

One upshot of this alternate justification for randomization is that it is not possible to fully divorce ethical questions about RCTs from methodological questions. Fairness and epistemic reasons push in opposite directions. If skeptics are right that oftentimes other methods work (almost) as well as randomization and that randomization faces problems of generalizing conclusions, then this would undercut not only the scientific merits of RCTs.³⁷ It would also influence the balance of reasons determining the moral permissibility of randomizing policies whenever the reasoning behind the *Equipose Condition* fails to hold.

VIII. Randomization between a PFO and FFO: Supererogation, Fairness, and the Fair Trial Condition

My discussion of the *Equipose Condition* has focused on situations in which policymakers face epistemic limitations in identifying the best policy. However, many policy experiments take place outside of such ignorance. Consider the deworming experiment. There is little doubt that deworming children is a good idea and that it will have beneficial effects on children's health and education. What could justify providing deworming treatment to some but not to others? Furthermore, what could justify doing so randomly rather than in a targeted manner?

If deworming all children was financially viable, then these questions would have no good answer. If deworming was feasible—what I called an FFO—then governments would be required to implement the best FFO. But sometimes good policies are only partly feasible—they are PFOs. Insofar as deworming is only partly feasible, experimenting with deworming then en-

36. If the inequalities resulting from the experiment are no worse than the inequalities that would have occurred otherwise—recall the last paragraph of Section V—then the inequalities are always proportionate.

37. This relates to what is discussed as the external validity problem of RCTs, see e.g. Deaton and Cartwright, "Understanding and misunderstanding"; Deaton, "Randomization in the Tropics Revisited," esp. pp. 34–37; Nancy Cartwright and Jeremy Hardie, *Evidence-Based Policy: A Practical Guide to Doing It Better* (Oxford University Press, 2012), <https://doi.org/10.1093/acprof:osobl/9780199841608.001.0001>; Rodrik, "The New Development Economics." Some suggestions for how to use RCTs better (and differently from much established practice) can be found in Cartwright and Hardie, *Evidence-Based Policy* as well as Angus Deaton, "Instruments, Randomization, and Learning about Development," *Journal of Economic Literature* 48, no. 2 (2010): 42–455, <https://doi.org/10.1257/jel.48.2.424>.

tails an experiment between the best (or maximal) among all FFOs and an even better PFO. Such experimentation can serve valuable epistemic ends. We might find out more information about how cost-effective the PFO really is. This can help us determine whether the policy should be rolled out universally. If the policy turns out more cost-effective than previously thought, then this would support universalization.

While some observers have demanded a principle of clinical equipoise which would rule out all such experimentation between an FFO and an even better PFO, defenders of such experiments have justified them on grounds other than equipoise.³⁸ In particular, researchers and philosophers have provided the following two-step defense of such experiments which I call the *argument from fair supererogation*.³⁹ First, governments, in providing the best policy among all the FFOs, are already providing everything that justice and individual rights demand. The government is not required to do more and indeed it cannot feasibly further its ends any better for everyone. What is feasible for governments is to spend a bit more on trying out something new. But it would not be unjust or violate anyone's rights if it did not, rather it would go beyond the call of justice. The situation is even clearer for NGOs and researchers. The treatment in NGO-run policy experiments is a benefit to the recipients. But it is also the case that the NGOs are not required to intervene. Nothing entitles the research subjects to the benefit. If researchers or NGOs decide to benefit some, then they act in supererogation. Insofar as the researchers already go beyond the call of duty, it seems that they cannot be faulted for not going even further.

The fact that agents act in supererogation explains why they can withhold benefits from some. But it does not tell us why it is permissible to do so in a random fashion. This is where the second step comes in. How should we decide who will be subject to the status quo (the best among all FFOs) and who will be subject to the new policy (the even better PFO)? In this scenario, the status quo is everything that citizens are entitled to. Neither justice nor individual rights nor a concern for their welfare grants an entitlement for any particular citizen to the new policy. No one can demand preferential treat-

38. For the critics see footnote 29. Defenders like MacKay make it clear that this argument is separate from any ideas of (policy) equipoise. See MacKay, "Government Policy Experiments," 333-48.

39. See MacKay, "Government Policy Experiments," 333-48; Till Bärninghausen et al., "HIV Treatment-as-Prevention Research at a Crossroads," *PLoS Medicine* 11, no. 6 (2014): e1001654, <https://doi.org/10.1371/journal.pmed.1001654>; Abhijit V. Banerjee and Esther Duflo, "The Experimental Approach to Development Economics," in *Field Experiments and their Critics*, ed. D.L. Teele (Yale University Press, 2014), 101-2, <https://doi.org/10.12987/9780300199307-006>.

ment and say that they should be first in line. The situation can be seen as one in which there is a scarce good—access to the new policy—and all citizens are equally deserving recipients with equally strong claims of justice to it.

The government needs to find a fair way to select who will receive the scarce good. It should not resort to any kind of favoritism, partiality, or discrimination. Even though the government is under more stringent duties to treat all citizens as equals, researchers, too, should avoid at the very least discrimination or undue favoritism. The fairest thing to do would be to run a lottery. This is so for three reasons. First, a lottery prevents bias, favoritism, partiality, or discrimination. It functions as a prophylactic against improper allocations.⁴⁰ Second, lotteries not only prevent impermissible modes of allocation, but they also publicly express a commitment to impartial allocation.⁴¹ Lastly, lotteries also give every citizen an equal chance of receiving the good. Insofar as it is impossible to give everyone access to the new policy, each citizen at least gets the next best thing—a chance of receiving it.⁴²

While not everyone accepts all three reasons, almost everyone accepts at least the prophylactic reason. This broad consensus is what supports random allocation when the better PFO counts as a scarce resource. All three approaches imply that randomization becomes a demand of equality. The best way to treat citizens as equals when distributing the scarce policy is to randomize. The beauty of this argument is that the knowledge gain is incidental to what is anyway the fairest option.

The argument from fair supererogation even works for those who detract from this near consensus. Those, like Tim Henning, who doubt the arguments for lottery fairness object to the claim that lotteries are *required* for fairness.⁴³ But even Henning accepts that lotteries

40. This is also called the sanitizing function of lotteries. George Sher, "What Makes a Lottery Fair?," *Noûs* 14, no. 2 (1980): 203–16, <https://doi.org/10.2307/2214861>; Lewis Kornhauser and Lawrence Sager, "Just Lotteries," *Social Science Information* 27, no. 4 (1988): 483–516, <https://doi.org/10.1177/053901888027004001>; Peter Stone, *The Luck of the Draw* (Oxford University Press, 2011), <https://doi.org/10.1093/acprof:oso/9780199756100.001.0001>.

41. David Wasserman, "Let them Eat Chances: Probability and Distributive Justice," *Economics & Philosophy* 12, no. 1 (1996): 29–49, <https://doi.org/10.1017/S0266267100003709>.

42. John Broome, "Selecting People Randomly," *Ethics* 95, no. 1 (1984): 38–55, <https://doi.org/10.1086/292596>; John Broome, "Fairness," *Proceedings of the Aristotelian Society* 91, no. 1 (1991): 87–102, <https://doi.org/10.1093/aristotelian/91.1.87>; Kai Spiekermann, "Good Reasons for Losers: Lottery Justification and Social Risk," *Economics & Philosophy* 38, no. 1 (2022): 108–31, <https://doi.org/10.1017/S0266267121000043>; Michael Otsuka, "Equal Chances versus Equal Outcomes: When Are Lotteries Fair and Justified?," *Political Philosophy* 1, no. 2 (2024): 381–98, <https://doi.org/10.16995/pp.16917>.

43. Tim Henning, "From Choice to Chance? Saving People, Fairness, and Lotteries," *Philosophical Review* 124, no. 2 (2015): 169–206, <https://doi.org/10.1215/00318108-2842176>.

are permissible. If some have to be selected, a lottery is one (of several, he would add) permissible options. And if a lottery—random allocation—has other, for example epistemic, advantages, then this would speak for lotteries on grounds other than fairness. Whether one believes that random selection increases fairness or that it does not diminish fairness, both views concur that randomization is *a* fair method of selecting citizens for the new policy.

This completes the two-step defense. The first step establishes that any decision to provide a policy which is better than the status quo but not universally feasible is a form of supererogation—going beyond the call of duty or justice. The second step establishes that random allocation is a fair method to decide who gets access to this improved policy. The argument from fair supererogation supports the following condition:

Fair Trial Condition Policymakers may permissibly randomly allocate citizens between an FFO and a PFO if (1) no citizen is given anything less than the best or maximal FFO, and (2) no citizen has a greater entitlement under justice to the better PFO.

The *Fair Trial Condition* and the argument which underpins it are popular among researchers. Edward Miguel, one of the authors of the deworming study, embraced its reasoning when justifying the study. Miguel points first to the fact that NGOs do not have the resources to provide deworming to all Kenyans—invoking ideas of supererogation—before pointing out that random allocation is the fairest selection procedure in such scarcity—invoking ideas of lottery fairness.⁴⁴ The reasoning behind the *Fair Trial Condition* also gives broad license to experiment. Marcus Picchio adopts its reasoning to argue that researchers can almost always take the status quo as a control in their experiments.⁴⁵

In the following, I shall argue against the *Fair Trial Condition* and the broad permission it gives to experiment. The condition suffers from two key defects. First, it overlooks that even though the policy as such may be beneficial, transition costs of switching policies can speak against randomization (Section IX). Second, the argument from fair supererogation ignores the significance of conditional obligations to help optimally while randomization often helps suboptimally (Section X).

44. Cited in Linda Kinstler, “How poor Kenyans became economists’ guinea pigs.”

45. Picchio, “What is the standard of care.”

IX. Transition Costs and Maximizing Knowledge versus Minimizing Harm

The argument from fair supererogation is less sweeping than it initially appears. Again, the culprit is that it seeks to extend its reasoning to policies which will eventually be all-things-considered better. This can be best seen in cases in which experimentation is done during a policy transition. Such policy transitions often come with transition costs. If we accept that there are transition costs, then we already accept that the new policy is no unadulterated manna from heaven. We assumed that the new policy is an improvement once it is up and running. But this does not entail that being part of its teething problems is a benefit. Insofar as teething problems are of concern, a more sensible reply would be to roll out the policy in a way that minimizes teething problems. Randomizing does not necessarily achieve this.

Take a real-life example. India's public distribution system (PDS) for subsidized food grains (and some vegetables) suffers from inefficiency losses. The government decided to tackle this problem through biometric verification.⁴⁶ Rolling out this technology has teething problems, especially in rural areas with less infrastructure. One option is to minimize teething problems by starting in urban areas before moving on to rural areas which predictably will have greater difficulty in adapting. But this rules out randomization. Even if the policy transition is permissible, a randomized policy transition often incurs additional and preventable costs.⁴⁷ Any assessment of randomizing a policy shift would need to take into consideration added transition costs. We cannot leap, as the *Fair Trial Condition* does, from the fact that there is a policy transition to a superior policy to the conclusion that randomizing that transition is superior.

There is a possible reconciliation of the aim to minimize transition costs and the aim to learn through randomization. One suggestion is to limit randomization within a "risk group" of likely transition costs. In the PDS example, this would mean starting among urban areas with a randomized group of urban areas first. Only then one moves to suburban areas, again randomizing which suburbs shift first, and so on.

46. Let's suppose that this is indeed superior to the status quo ante; for skepticism, see Reetika Khera, "Impact of Aadhaar on Welfare Programmes," *Economic & Political Weekly* 52, no. 50 (2017): 61–70, <https://www.jstor.org/stable/45132600>.

47. For an analysis of such costs in the PDS, see Jean Drèze et al, "Aadhaar and Food Security in Jharkhand: Pain without Gain?," *Economic & Political Weekly* 52, no. 50 (2017): 50–59, <https://www.jstor.org/stable/45132599>.

This reconciliation at best reduces the problem but does not eliminate it. There remains a trade-off between minimizing transition costs and maximizing learning. Minimizing transition costs pulls towards the direction of ordering the rollout by predicted impact. The more subdivisions, the better. Maximizing the effects of learning from randomization pulls towards large cohorts that are randomized. Sample sizes matter for RCTs. This is for two reasons. First, randomization does *not ensure* that all possible causal factors are balanced between the treatment and control group. Rather, as sample sizes tend to infinity the difference between the observed treatment effect and the actual treatment effect becomes arbitrarily small. The smaller the sample size, therefore, the more likely that relevant causal factors are unbalanced, i.e. disproportionately distributed across control and treatment groups. In medical research, this is called random confounding. Thus, epistemically it would be better to have larger sample sizes and fewer subdivisions of the sample to reduce random confounding.⁴⁸

A second reason are statistical outliers.⁴⁹ Consider microfinance programs. Suppose that the success of microfinance programs depends largely on a small number of people with great entrepreneurial skill but without the ability to finance their business.⁵⁰ If so, then the evaluation of microfinance programs largely depends on few outliers. In a randomized evaluation, everything then turns on how these few business geniuses are allocated into control and treatment groups. The smaller the sample size, the larger this problem.

More generally, randomization within selected sub-groups goes against the spirit of randomization. Recall that RCTs are supposed to alleviate concerns of selection effects. But the stratification into risk groups is precisely such a form of a selection effect. Suppose the shift to biometric authentication was not done by random assignment but rather by self-selection of villages. An observational study could analyze the effects. In both cases there is a selection effect—in the RCT it is the policymaker who selects, in the observational study it is the village itself.

In all of these cases, one can still use randomized procedures to try to learn. But identifying randomly confounded allocations and dealing with statisti-

48. For discussion of this dynamic, see Deaton and Cartwright, “Understanding and misunderstanding,” 5–6.

49. Again, the discussion of Deaton and Cartwright is very helpful in “Understanding and misunderstanding,” 8–9.

50. Deaton and Cartwright, “Understanding and misunderstanding,” 8. The possibility is also discussed in a review paper on microfinance RCTs. Abhijit Banerjee et al., “Six Randomized Evaluations of Microcredit: Introduction and Further Steps,” *American Economic Journal* 7, no. 1 (2015): 1–21, <https://doi.org/10.1257/app.20140287>.

cal outliers as well as selection effects all require background assumptions which proponents of RCTs seek to avoid. Researchers would, for example, need to identify possible causal effects to test whether they are randomly confounded, or they would need to identify whether statistical outliers are more likely data errors or causally relevant. This requires the same kind of “soft evidence” for which randomistas criticize observational studies. RCTs are then in the same boat as their rivals. If one wants to maintain the strict methodological approach of the randomistas, then epistemic and fairness considerations pull in different directions.

The argument from fair supererogation, too, is thus more limited than it appears. Consequently, the *Fair Trial Condition* has a narrower scope than stated. As in the case of the *Equipose Condition*, we need a further necessary condition for the *Fair Trial Condition* to be a sufficient condition for randomization. We need the absence of transition costs or other costs to the study population. If, on the other hand, costs are present, then we once again face a trade-off between fairness and epistemic considerations. The situation then mirrors the one discussed in Section VII. Sometimes this trade-off will license experimentation. But other times it will not. Suppose that the experiment will reveal only modest policy-relevant knowledge. Once the transition is made, the knowledge about how best to transition this policy is moot and while the results can be used to inform other policy transitions, its generalizability is limited. At the same time, the experiment causes grave transition costs, for example large-scale disruptions in food support for people in extreme poverty.⁵¹ In such a case the transition costs would be disproportionate to the benefits generated from the experiment. As with the case of inequality, the transition costs and other costs cannot be disproportionate to the benefits due to knowledge gains that can be reasonably expected.

X. Conditional Obligations to Benefit Effectively

I shall now go on to argue that the *Fair Trial Condition* fails to hold even in cases in which there are no transition costs and the intervention truly benefits recipients. Let me recap the first part of the reasoning behind the argument from fair supererogation. The *Fair Trial Condition* stipulates that no citizen has a greater entitlement under *justice* to the better PFO. This renders the *Fair Trial Condition* compatible with many actual RCTs. For example, in cash transfer programs some members of the target population will

51. Drèze et al., “Aadhaar and Food Security in Jharkhand” suggest that significant disruptions indeed happened in the PDS reform.

be poorer and others will be relatively less poor. Proponents of RCTs insist that random, rather than targeted, allocation is, nevertheless, permissible in such scenarios. MacKay justifies such programs on the grounds that insofar as cash transfers for all are too expensive, they are a discretionary good and no one has a claim of *justice* to them even if they have differing needs. In discussing the opening example of Nairobi's water system, he similarly suggests that insofar as the Nairobi Water Company cannot provide water to Kayole-Soweto without payment, no one has a moral claim to a continued water connection.⁵² When it comes to experiments run by non-government entities such as researchers or NGOs, defenders of RCTs have similarly invoked supererogation—researchers already go beyond the call of duty or justice.⁵³ Researchers or NGOs are not obligated to help anyone at all. What objection can there be if they help in the way they prefer?

The problem with the argument from fair supererogation is that it implicitly assumes a particular view of supererogation which we should reject. According to this view of supererogation, once we enter the realm of the supererogatory, there is freedom to choose and hardly any moral constraints.⁵⁴ But this view of supererogation is false. The decision to provide help may trigger more stringent conditional duties. Here is an illustration. A building is on fire. No one is required to run into the burning building to save any of the people inside the building. The risks to oneself are so high that it is permissible not to help. Any help would be supererogatory. Still, you decide to run into the burning building. You grab a first child. There is a second child present whom you can also save, but you decline to do so. In this case, it would not

52. MacKay, "Government Policy Experiments," 335; MacKay, "The Ethics of Public Policy Experiments."

53. See e.g. Picchio, "What is the standard of care," 212–14. See also the quotations by practitioners cited in Kinstler, "How poor Kenyans became economists' guinea pigs."

54. Most formulations of this view on supererogation omit the possibility of wrongful discriminatory supererogation. But even if we add this constraint, the view retains a broad freedom of choice among supererogatory options. Tina Rulli calls this the free space view of moral options. Daniel Muñoz calls this the 1D or one-dimensional solution to the paradox of supererogation. Christian Barry and Seth Lazar dub this as a disconnection between weighing up costs and benefits (i) until the call of duty and (ii) beyond the call of duty. Tina Rulli, "Conditional Obligations," *Social Theory and Practice* 46, no. 2 (2020): 373, <https://doi.org/10.5840/soctheorpract20204189>; Daniel Muñoz, "Three paradoxes of supererogation," *Noûs* 55, no. 3 (2021): 701, <https://doi.org/10.1111/nous.12326>; and Christian Barry and Seth Lazar, "Supererogation and Optimisation," *Australasian Journal of Philosophy* 102, no. 1 (2024): 26–29, <https://doi.org/10.1080/00048402.2022.2074066>. Notably, all these authors reject—as I do—this view. Barry and Lazar even discuss claims as opposed to discretionary goods as a possible source of support for this view in ways reminiscent of MacKay invoking discretionary goods to justify the *Fair Trial Condition*. Compare Barry and Lazar "Supererogation and Optimisation," 27 with MacKay, "Government Policy Experiments," 335.

be any justification for your action that you could have permissibly saved no one. Given that you ran into the building, you acquired the obligation to help both.⁵⁵

One way to explain what is going on in the burning building case is the interplay of requiring and permitting reasons.⁵⁶ Requiring reasons are considerations counting in favor of morally requiring a given act. Permitting reasons are considerations counting in favor of making a given act permissible. One important permitting reason is personal cost to an agent. In the burning building example, there is sufficient permitting reason to allow you to refuse to enter the burning building simply because this will be risky and costly to yourself. But insofar as you decide to enter the building, you cannot appeal to the cost of entering in order to justify saving one rather than two people. There is no permitting reason to save one rather than two and there are strong requiring reasons to save a person's life.

There are two possible conditional obligations that researchers may have. One possibility is that if we give to some, we ought to give to all. The difficulty in many RCTs is that they budget the target intervention only for some and providing it for everyone would incur extra financial costs. This raises difficult issues of what makes policy options "feasible." For this reason, I want to focus on another possible conditional obligation. Even if we take the financial cost as fixed, one can redistribute benefits and thereby help more effectively. If we only give to some, we ought to give to the neediest first. A version of my burning building story supports this approach. Suppose that the children are surrounded by fires. The children are safe from the fire, but one cannot reach them without crossing through the fires. The children have injuries of different severity. Waiting for the fires to burn down would mean that help reaches the children too late and consequently they will suffer long-term harms. You are not required to run through a fire to bring the children to safety and ensure that their injuries are attended to. But if you

55. See also Derek Parfit, "Future Generations: Further Problems," *Philosophy & Public Affairs* 11, no. 2 (1982): 131, <https://www.jstor.org/stable/2264925>; Shelly Kagan, *The Limits of Morality* (Clarendon Press, 1989), 16–17, <https://doi.org/10.1093/0198239165.001.0001>; Theron Pummer, "Whether and Where to Give," *Philosophy & Public Affairs* 44, no. 1 (2016): 83–86, <https://doi.org/10.1111/papa.12065>; Theron Pummer, *The Rules of Rescue* (Oxford University Press, 2023), 62–64, <https://doi.org/10.1093/oso/9780190884147.001.0001>; Joe Horton, "The All or Nothing Problem," *Journal of Philosophy* 114, no. 2 (2017): 94, <https://doi.org/10.5840/jphil201711427>; Jeff McMahan, "Doing Good and Doing the Best," in *The Ethics of Giving*, ed. P. Woodruff (Oxford University Press, 2018), 81–84, <https://doi.org/10.1093/oso/9780190648879.003.0004>; Rulli, "Conditional Obligations," 365–67; Muñoz, "Three paradoxes of supererogation," 699; and Barry and Lazar, "Supererogation and Optimisation," 21–22.

56. This adapts the framework by Pummer, *The Rules of Rescue*, esp. chs. 1 and 3.

do, then you are conditionally obligated to save the children whose needs are greatest—if, as in this case, the neediest will benefit at least as much as anyone else. You could not randomly choose among them.

This is a good reply to the argument from fair supererogation in its original form. If we provide cash transfers, for example, we know that people are of differing need and can prioritize accordingly. Another way to redistribute benefits is to share them more equally. If we have to ration water supply, we could spread out water more thinly for everyone to ensure that everyone has access to at least some water.

While the argument from fair supererogation seeks to justify randomization without relying on epistemic gains, a defender of RCTs could seek to rebut the concern of conditional obligations by relying more heavily on the knowledge gains from RCTs. Suppose that the RCT creates substantial knowledge gains which will enable more people to be helped in the future. If so, then unlike in the burning building example, random allocation is only suboptimal relative to the immediate study population, but the experiment will produce overall more good than bad. This makes it different from the burning building case in which one action is unambiguously better (and thus has more requiring reason) than the other. Hence, there cannot be a conditional obligation in this case. Or so the argument goes.

In response, it first bears repeating that this argument admits one of my central points that fairness and epistemic considerations pull in opposite directions and that the ethical evaluation of RCTs cannot be divorced from the methodological evaluation. Second, I do not believe that it suffices to show that running the experiment would be overall axiologically best. This is because we can be conditionally obligated to perform actions which are axiologically suboptimal.⁵⁷ To see why consider first that non-consequentialists believe that we can be unconditionally obligated to do what is (axiologically) suboptimal. For example, you are obligated to keep a promise to help a friend even when volunteering at your local charity would be axiologically better. Suppose now there is a snowstorm such that it is permissible for you to stay at home. If you decide to go out nevertheless, then you are conditionally obligated to honor your promise and help your friend. You now lack

57. This is often overlooked because the focus in the debate on conditional obligations is on the interplay between personal prerogatives and moral reasons to benefit others. In this set-up which, for example animates the “all-or-nothing problem” that has dominated debates on conditional obligations, the putative conditional obligation is to do what is unambiguously best axiologically. See, for example, the set-up of Barry and Lazar, “Supererogation and Optimisation”; Pummer, “Whether and Where to Give”; or Horton, “The All or Nothing Problem.”

a permitting reason not to help her, and you have most requiring reason to honor your promise.

The same dynamic can occur due to human needs. Most non-consequentialists believe that sometimes a person's need is weighty enough that we should fulfill that person's need rather than do what is impersonally best.⁵⁸ If you have a permitting reason that allows you to do neither, but you decide to do something anyway, then once more you would be conditionally obligated to fulfill that person's need. Here is a thought experiment to illustrate. You are taking a stroll by a lonely beach. The tides are rising, and the sea is treacherous. You suddenly see a child drowning in the sea desperately calling out for your help. You also see a painting on a rock which will soon be submerged in the sea. You correctly estimate that if you brought the painting back to shore, you could (and would) sell it and donate the proceeds to effective charities which will enable them to benefit many people equivalent to an aggregate estimate of 1.5 statistical lives. You also correctly estimate that there is no time to save both the child and the painting. What may you do? Since it is dangerous for you to enter, it is permissible for you to stay safely on shore. Yet in an act of bravery, you jump into the sea. Would it now be permissible for you to swim past the child to rescue the painting instead?

I do not believe this to be case. If you jump into the sea, you are conditionally obligated to save the child. Once you jump into the water and swim out far enough, any personal cost to you is a sunk cost (if you drown, it is *literally* a sunk cost) and you lack a permitting reason not to save the child. In choosing between the painting and the child, it appears to me that the need of the child is more pressing than the benefits from selling the painting.

One reason for this is that the benefit of selling the painting, just as the benefit of enlightened policies, is dispersed across many individuals and uncertain. By contrast, the benefit to the person here and now does not raise any such issues of aggregation and uncertainty. A second reason is the following. In this situation of direct encounter with the child, the direct confrontation of their needs provides you with a strong agent-relative reason to

58. This is implied by a non-consequentialist opposition to full aggregation according to which sometimes fulfilling a person's claim based on their need takes precedence over aggregate welfare. It is also a rather plausible extension of Raz's interest theory of rights according to which individuals have rights when their interest (such as their need) is weighty enough to hold others under a duty. Joseph Raz, *The Morality of Freedom* (Clarendon Press, 1986), ch. 7, <https://doi.org/10.1093/0198248075.003.0007>. If we add the modest claim that such rights are not always overridden when they conflict with the aggregate good, then we should sometimes fulfill a person's need rather than the aggregate good.

help *that* child which can outweigh your impersonal duty to aid.⁵⁹ It is hard to explain why exactly this direct encounter with the needs of someone else grounds this agent-relative duty. But the idea of encountering and recognizing the “other” and their unique significance seems to play an important role at the bedrock of morality.⁶⁰

This ties in with a critique of effective altruism made by Amia Srinivasan. Srinivasan takes issue with the suggestion that helping a person one encounters is “arbitrarily” favoring them. She suggests that while it is arbitrary whom we encounter, once we encounter people our relationship to them is far from being arbitrary.⁶¹ Although Srinivasan does not link her point to conditional obligations, she implicitly invokes conditional obligations. Conditional on helping and encountering particular individuals, we acquire special reasons to help *them*. In a similar vein, social scientists engaging in policy experiments sometimes encounter special reasons to help their study subjects in particular simply because they provide cash transfers, deworming treatments, or the like.⁶²

Whether these reasons translate into a conditional obligation will depend on the case. The reasons to help those we encounter are stronger the greater the needs of the individuals. Likewise countervailing reasons are stronger the greater the opportunity costs. The reasons in favor of a conditional obligation can be defeated in some cases if there are sufficiently good out-

59. Larry Temkin agrees with me in his thoughtful discussion of an unconditional obligation to help one child at the expense of losing the ability to do more good via charitable donations. Larry Temkin, *Being Good in a World of Need* (Oxford University Press, 2022), 67–81, <https://doi.org/10.1093/oso/9780192849977.001.0001>.

60. For different attempts to spell this out, see Temkin, *Being Good in a World of Need*, 77–78; Kieran Setiya, “Other People,” in *Rethinking the Value of Humanity*, eds. S. Buss and N. Theunissen (Oxford University Press, 2023), 325–26, <https://doi.org/10.1093/oso/9780197539361.003.0014>; and Sergio Tenenbaum, “Can’t Kant Count? Innumerate Views on Saving the Many over Saving the Few,” in *Oxford Studies in Normative Ethics*, Vol. 13, ed. M. Timmons (Oxford University Press, 2023), 218–20, <https://doi.org/10.1093/oso/9780198895909.003.0011>. Setiya and Temkin invokes Levinas’s idea of recognizing the “other.” Tenenbaum sees this as reflecting Kantian ideas about human dignity having no price.

61. Amia Srinivasan, “Stop the Robot Apocalypse,” *London Review of Books*, September 24, 2015, <https://www.lrb.co.uk/the-paper/v37/n18/amia-srinivasan/stop-the-robot-apocalypse>.

62. A similar claim is sometimes made with respect to clinical researchers. For a skeptical discussion, see Alan Wertheimer, *Rethinking the Ethics of Clinical Research* (Oxford University Press, 2010), ch. 6, <https://doi.org/10.1093/acprof:oso/9780199743513.003.0006>. For a sympathetic discussion, see Rulli, “Conditional Obligations,” 383–88. Several putative explanations for this phenomenon focus on the clinical context, yet my argument above shows that it generalizes rather naturally outside clinical contexts. Thus, the best argument for such conditional obligations does not rest on the role obligations of physicians, *pace* Picchio, “What is the standard of care,” 212; and, *pace* Henry S. Richardson, “Moral Entanglements: Ad Hoc Intimacies and Ancillary Duties of Care,” *Journal of Moral Philosophy* 9, no. 3 (2012): 376–409, <https://doi.org/10.1163/174552412X628922>, such obligations (of ancillary care) are not deeply tied to privacy waivers.

comes to be expected from the experiment. But in other cases, especially if other means of learning are not much worse, researchers are conditionally obligated to help effectively. In these cases, a decision to provide a benefit triggers a conditional obligation to provide this benefit efficiently by targeting the neediest as opposed to allocating it randomly. For example, further experiments on deworming may generate some additional knowledge. But we already have substantial evidence on deworming and could also generate more knowledge through careful observational studies. The gains of the experiment are, plausibly, not sufficient to undermine the reasons for a conditional obligation to help the neediest. If so, then an NGO providing deworming treatments would have to prioritize children based on medical need. Despite what the argument from fair supererogation suggests, the fact that the NGO was not obligated to provide anyone with the benefit does not by itself license it to deliver the benefit in a random fashion. The argument from fair supererogation and the *Fair Trial Condition* only apply if there is no conditional obligation to help effectively. And, as I suggested, it is not sufficient to argue that performing the experiment will be overall, all-things-considered, best.⁶³ This is another necessary condition limiting the scope of the sufficient *Fair Trial Condition*. It also suggests something stronger: If it is the case that researchers are conditionally obligated to help the neediest, then they may not randomly allocate their help—they may not experiment.

XI. Conclusion

Randomized policy experiments are increasingly used in development research. It sometimes appears as if they are the only game in town. But such policy experiments often involve giving some citizens a benefit which is withheld from others. The arguments from parity and from fair supererogation seek to justify this unequal treatment. They suggest that either unequal treatment does not entail disadvantageous treatment (parity) or that citizens are not entitled to the benefit (supererogation) and that such benefits are given fairly (fairness). What these arguments have in common is that they seek to reconcile epistemic considerations with considerations of fairness. The crux of my arguments has been that this reconciliation often fails. Oftentimes epistemic and fairness considerations pull in opposite directions.

63. A reviewer for *Free & Equal* reminds me that this argument is compatible with the idea that governments or NGOs can randomize among the neediest. This mirrors the suggestion discussed in Section IX. For the same reasons this would be a methodological cost to RCTs which is why current practice in development research does not attempt to prioritize the neediest first.

As my argument in Sections V and VI showed, even in conditions of policy equipoise policymakers have to consider any possible inequalities between treatment and control arm. Such inequalities can also arise between sub-sets of the treatment and control arm. Policymakers need to take a deeper look at the policies and their (likely) distributional profiles. They have to move past merely thinking in terms of “what works.” Only if such inequalities are non-existent, the *Equipoise Condition* holds. As a working proposal, they could be guided by the following revised condition.

Revised Equipoise Condition If policymakers are, due to epistemic uncertainty, unable to rank different policy options, they may permissibly assign citizens randomly into different groups corresponding to the multiple maximal fully feasible options (i.e. no one is assigned an option which is worse than an alternative fully feasible option) *as long as the experiment does not introduce any inequality between (subsets of the) treatment and (subsets of the) control arm which increases overall inequality.*

The other case of experiments concerns comparisons of FFOs with even better PFOs. An example of such a comparison is the Kayole-Soweto water experiment. Here I pointed to two flaws. First, the fact that the PFO is overall better does not mean that it is better for everyone. Second, the fact that the PFO is not unconditionally mandatory does not preclude the PFO from being conditionally mandatory. Again, as a working definition researchers would do better to consider the following revised condition.

Revised Fair Trial Condition Policymakers may permissibly randomly allocate citizens between an FFO and a PFO if (1) no citizen is given anything less than the best or maximal FFO *and no citizen is subjected to transition costs or other costs arising from the experiment*, and (2) no citizen has a greater entitlement under justice to the better PFO *nor does any citizen have a conditional entitlement to the better PFO.*

Both revised conditions are much narrower in scope than the condition they replace. Even if the Kayole-Soweto experiment could be justified on grounds of the original *Fair Trial Condition*, the *Revised Fair Trial Condition* would be harder to meet. It is plausible that even if shortages are unavoidable, residents do have a conditional entitlement to have the water shared more fairly among them—whether by distributing it more thinly or by prioritization according to need.

For this reason, researchers will often need to resort to a different route for justifying their experiments. They can invoke benefits which result from

knowledge gains. Such benefits could outweigh some costs to research participants. But this acknowledges a tension between epistemic and fairness concerns. Furthermore, as I argued, the fairness costs (whether in terms of inequalities or transition costs) cannot be disproportionate to the benefits we can reap from the epistemic gains of the experiment. This means that we cannot divorce methodological disputes about RCTs as cleanly from their moral permissibility as we initially believed.

There is more to be said on the proportionality of costs relative to epistemic gains, and resolving these trade-offs is likely going to be complex.⁶⁴ However, what already emerges is that only some experiments can be justified on these grounds. In other cases, the costs are disproportionate to the gains of the experiment. Ultimately, the picture of permissible randomization that we end up with licenses fewer experiments than the picture we started out with.

Acknowledgements Versions of this paper were presented at the Center for Population-Level Bioethics at Rutgers University, the Analytic Political Philosophy Workshop, and the Centre for the Experimental-Philosophical Study of Discrimination at Aarhus University. I benefited from interactions with the audiences and their comments, concerns, and criticisms. I would also like to thank Andreas Bengtson, Mark Budolfson, Nir Eyal, Jean Drèze, Dan Hausman, Reetika Khera, Doug MacKay, Monica Magalhães, Tarun Menon, Søren Flinch Midtgaard, Mike Otsuka, Yasashvi Paarakh, Vikki Møller Lyngby Pedersen, Marcos Picchio, Thulasi K. Raj, Danny Weltman, Dan Wikler, Bridget Williams, and Miklós Zala for their input and suggestions. A special thanks goes to Nir Eyal who was very helpful and supportive during the early stages of this project. Three anonymous reviewers for *Free & Equal* as well as the Editor-in-Chief, Anna Stilz, have been very generous and patient with their feedback and guidance.

Competing Interests The author has no competing interests to declare

64. This can be unsatisfying for anyone seeking clear guidance. But then why should we believe moral theory to be simple? If our investigation reveals that a moral phenomenon is more complex, then moral theory should reflect that complexity. See Judith Jarvis Thomson, *Rights, Restitution, and Risk* (Harvard University Press, 1986), 255–57.