

Group Rights in an Era of Large-Scale Field Experiments

Building on McDermott and Hatemi (2020), this reply sketches a group-based ethical framework for the design and implementation of large-scale field experiments by social scientists. It makes the case for including group rights in our risk-benefit calculations and argues that communities, like individuals, have rights to anonymity and confidentiality. The paper also advances six principles to strengthen the ethical guardrails around our field experiments when harmful spillovers are present or suspected.

Word Count: 1,569

McDermott and Hatemi (2020) are right to sound the alarm about the significant, often unrecognized, ethical concerns raised by large-scale field experiments in the social sciences. It is clear that our collective Institutional Review Boards (IRBs) are no longer a guarantee of ethical research — if they ever were — in part because they are focused too narrowly on individual subjects. In a world where we are now randomizing exposure to peacekeepers and (corrupt) police in fragile and post-conflict settings; where we are encouraging high-risk activism in authoritarian contexts; and when our experiments manipulate our beliefs about basic democratic processes, it is time for a new ethical framework to guide our design and deployment of large-scale field experiments. Building on the lead article’s discussion, I first sketch a possible framework here, one built around group rights and our obligations to the communities that host and are affected by our field experiments. I then offer six practical solutions for safeguarding the ethical integrity of our field experiments.

Let me first put my cards on the table. I believe randomized control trials (RCTs) are an important tool in our methodological toolkit; I conduct them myself.¹ But I also believe that field experimenters have a special obligation to ensure we uphold high ethical standards since the downsides of our interventions (even if successful) far surpass other methodologies in sweep and severity. We typically work in teams with far greater resources (and ambitions) than lone researchers conducting interviews, ethnography, or archival work. Our samples routinely include thousands of individuals spanning hundreds of villages. We often collaborate with state agencies and foreign non-governmental organizations, magnifying our access to policy-makers while raising questions about power imbalances relative to study participants. And we usually operate blind to ethical risks over the medium and long term since our research designs emphasize one-time interventions that ignore second-order consequences within and outside our samples.

We must not sidestep these ethical issues. We are not passive observers but active participants in a process of discovery that involves direct intervention into the fabric of communities and nations.² We cannot retreat behind standard platitudes (“they were going to roll the program out anyways”). If we design or assign the intervention, then we are directly implicated in its consequences. RCTs thus represent a great opportunity for scientific advancement, but one that must reflect strong ethical safeguards to protect individuals and groups alike.

¹E.g., Lyall, Zhou and Imai 2020.

²Samii 2020.

1 A Group-Centered Ethical Framework

McDermott and Hatemi present a compelling case for how informed consent and debriefing should guide field experimentation. I propose two additional principles to build out a group-centered ethical framework for large-scale field experiments.³

First, researchers need to incorporate group rights into their risk-benefit calculations when assessing the potential effects of a proposed informational or behavioral intervention. “Groups” here refers to the communities in which individuals within the experiment are embedded *and* adjacent communities that are connected, either spatially or through social networks, to the experiment’s sample. In nearly all settings, researchers must be alive to the possibility that social networks can act as conduits for unintended harm to individuals and groups both inside and outside the defined sample. To date, however, our discussions of spillover largely cast it as a threat to causal inference rather than as a potential multiplier of risk and harm to adjacent communities.⁴

Yet these concerns become apparent when we adopt a group-centered approach to field experiments. Let me provide two simple examples. Imagine, for example, that a particular behavioral intervention like hot-spot policing displaces crime or violence into (unsuspecting) neighboring communities as gangs or rebels respond strategically to state actions. This harm is “off-the-radar” of traditional IRB reviews since these communities are not part of the study population. Or take the example of villages that publicly receive large development programs or cash transfers. In some settings, these interventions can spark jealousy, even violence, among excluded neighboring villages. Here, the intervention has actually increased the risk faced by an experiment’s subjects but not from a direction anticipated by the IRB or measured by standard RCT frameworks.

Second, communities deserve the same protections as individual subjects. Nearly all of our current data practices revolve around the assurance of confidentiality and anonymity to participants in our field experiments. Few, however, extend these same protections to the communities where the experiments are conducted or to adjacent populations. At a minimum, scholars should consider redacting place names, administrative districts, and

³The new APSA guidelines “Principles of Ethics for Human Subjects Research and Guidance” marks a welcome effort to build a common framework. It is centered around individual subjects, however, not larger groups and communities. For important discussions of ethical fieldwork, see Fujii 2012; Campbell 2017; Cronin-Furman and Lake 2018.

⁴See, for example, Glennerster and Takavarasha 2013.

other metadata from public datasets to shield the identity of these communities. Most importantly, geographic coordinates (latitude/longitude) should be removed from public datasets to prevent identification.⁵

Without these procedures, researchers may inadvertently increase the risk of harm to participating and adjacent communities. For example, behavioral interventions designed to reduce corruption, foster support for local politicians, or improve perceptions of government performance and legitimacy might fail to produce desired results. In turn, these failures might expose communities to various forms of retribution by government authorities. The knowledge that one’s village is a “control” while neighboring villages profit from development assistance might also increase alienation and frustration against the government, external donors, or nearby recipients, sparking violence. Even behavioral interventions that lift incomes and boost support for a government can have downside risks if these successes raise the odds that insurgents or criminals will target these populations.⁶

While groups have a right to public erasure of their involvement in a field experiment, honoring this obligation poses logistical and other challenges. Collaboration with a state agency might prove problematic, for example, if anonymity cannot be guaranteed and if scholars suspect negative downstream consequences for participating communities. Yet measures are available that obscure, if not sever entirely, the link between an experiment and its host location(s). Large-scale interventions can forgo public branding that draws (hostile) attention to experimental sites, for example. Experiments could also be conducted “off-site” in safe (neutral) locations to shield communities from negative externalities or reprisals. While safeguards will be context-specific, the notion that groups have rights to anonymity will force scholars to consider how best to maintain this obligation from the planning stages to well after the experiment has concluded.

2 Applied Ethics

What practical steps could field experimenters adopt to safeguard group rights? I offer six suggestions below.

⁵These data could be made available privately for replication or follow-on studies.

⁶Crost, Felter and Johnston 2014. Treatment status, in other words, is guiding insurgent/criminal targeting decisions.

Incorporate group rights into net risk-benefit calculations. Well before research designs are crafted, researchers should explicitly theorize the proposed intervention’s direct and indirect effects on both the intended sample and adjacent populations. Central to this task will be articulating *who* these populations are, possible spillover channels, and the potential benefits and harm that might result for (non-)participating communities. Assessing the likelihood of these risks and benefits, along with the net benefits, for both individuals and all relevant communities is crucial. These considerations should inform how the sample is constructed, the intervention’s actual design, and the adoption of safeguards to minimize harmful spillover.

Informed consent. Researchers should not only seek informed consent from individual subjects in their field experiments but also their host communities (or their representatives).⁷ Informed consent should also be obtained among adjacent communities that might be indirectly affected through spillover. Exceptions to obtaining informed consent from some or all of these populations should be rare, explicitly acknowledged in resulting publications, and must not offload risk to unsuspecting adjacent communities.

Preregistration. Scholars should preregister their ethical concerns about the field experiment, including potential spillover and general equilibrium effects, as part of their preregistration exercise.⁸ Detailed plans for the measurement and mitigation of potential harms could easily be included in standard preregistration plans (such as DeclareDesign).⁹ In addition, researchers should articulate an explicit mechanism for shutting down an experiment (a “kill-switch”) in case of adverse externalities. Pre-specified metrics, including upticks in crime or violence among experimental or adjacent populations, could be used to monitor for harmful effects from an intervention in near-real time.

Robust piloting. Before launching the full experiment, researchers should pilot both the proposed intervention and survey instrument in the intended sample and among adjacent communities where spillover is suspected. Desk reviews of past interventions can be helpful here, but spillover is often contextually specific, and so devoting substantial resources to a

⁷See Wood 2006.

⁸Lyall 2019.

⁹Blair et al. 2019.

pilot large enough to detect indirect effects is valuable, especially if little is known about spillover channels.

Debriefing. Researchers should make a concerted effort to debrief subjects in the experimental and adjacent populations. For the latter group, community-wide notification, perhaps through public meetings shortly after the program concludes, of the field experiment and its purpose would represent a good-faith effort to mitigate harm arising from the experimental intervention. Such efforts are especially important where informational interventions sought to alter or reinforce attitudes toward out-groups or where behavioral interventions created negative spillovers that affected coping strategies (e.g., with crime) or intergroup relations. Public hearings also represent a low-cost opportunity to collect additional data on the spatial reach and temporal persistence of our interventions, helping refine our theories while improving public policies.

Post-experiment harm reduction. In conjunction with debriefing, researchers should preregister their plans for mitigating any potential second-order consequences that might arise for communities that participated in the experiment or that were adjacent to it. To avoid “informational harm,”¹⁰ scholars should redact community-identifying information, including latitude and longitude coordinates, from public datasets. Scholars, in conjunction with implementing partners, can also devise plans to monitor relevant communities for post-experiment harm. Above all, we need to consider what we owe these communities for their participation, even if unknowingly, in our experiments and how long these obligations extend to them.

3 Conclusion

This ethical framework is meant to guide, not straitjacket, field experiments. Not every eventuality can be foreseen; mistakes made in good faith will continue. But the current status quo of large-scale field experiments unleavened by ethical considerations is likely untenable. The framework proposed here reduces unintended harm by forcing us to consider the net benefits to both individual and groups, including adjacent communities often neglected in current empirical strategies. This comes at a cost; our field experiments will

¹⁰Raymond 2019.

likely be slower and more expensive if conducted within this framework. But the advantages are clear, both in terms of our ethical conduct and better social science, including new insights into spillover mechanisms, the spatial and temporal persistence of our treatments, and greater policy relevance. In short, a focus on group rights helps ensure our research is both cutting-edge and ethically responsible to the communities that make our advances possible.

References

- Blair, Graeme, Jasper Cooper, Alexander Coppock and Macartan Humphreys. 2019. “Declaring and Diagnosing Research Designs.” *American Political Science Review* 113(3):838–859.
- Campbell, Susanna. 2017. “Ethics of Research in Conflict Environments.” *Journal of Global Security Studies* 2(1):89–101.
- Cronin-Furman, Kate and Milli Lake. 2018. “Ethics Abroad: Fieldwork in Fragile and Violent Contexts.” *PS: Political Science & Politics* 51(3):607–614.
- Crost, Benjamin, Joseph Felter and Patrick Johnston. 2014. “Aid Under Fire: Development Projects and Civil Conflict.” *American Economic Review* 104(6):1833–1856.
- Fujii, Lee Ann. 2012. “Research Ethics 101: Dilemmas and Responsibilities.” *PS: Political Science & Politics* 45(4):717–723.
- Glennerster, Rachel and Kudzai Takavarasha. 2013. *Running Randomized Evaluations: A Practical Guide*. Princeton, NJ: Princeton University Press.
- Lyall, Jason. 2019. “Preregister Your Ethical Redlines.” *Unpublished Manuscript* .
- Lyall, Jason, Yang-Yang Zhou and Kosuke Imai. 2020. “Can Economic Assistance Shape Combatant Support in Wartime? Experimental Evidence from Afghanistan.” *American Political Science Review* 114(1):126–143.
- Raymond, Nathaniel. 2019. “Safeguards for human studies can’t cope with big data.” *Nature* 568:277.
- Samii, Cyrus. 2020. “Reasons for policy experimentation that have nothing to do with bias.” *World Development* 127:1–2.
- Wood, Elisabeth. 2006. “The Ethical Challenges of Field Research in Conflict Zones.” *Qualitative Sociology* 29(3):373–386.