

Ethical Challenges of Embedded Experimentation
Macartan Humphreys
Columbia University

Introduction

Consider a dilemma. You are collaborating with an organization that is sponsoring ads to inform voters of corrupt practices by politicians in a random sample of constituencies. The campaign is typical of ones run by activist NGOs and no consent is sought among populations as to whether they wish to have the ads placed on billboards in their neighborhoods. You learn that another NGO is planning to run a similar campaign of its own in the same area. Worse (from a research perspective) the other organization would like to target “your” control areas so that they too can make an informed decision on their elected representatives. This would destroy your study, effectively turning it from a study of the effect of political information into a study of the differences in the effects of information interventions as administered by two different NGOs. The organizations ask you whether the new group should work in the control areas (even though it undermines the research) or instead quit altogether (and in doing so, protecting the research but possibly preventing needy populations from having access to important information on their representatives). What should you advise? Should you advise anything?

Consider a tougher dilemma. You are interested in the dynamics of coordination in protest groups. You are contacted by a section of the police that is charged with deploying water cannons to disperse protesters. The police are interested in the effectiveness of water cannons and want to partner with a researcher to advise on how to vary the use of water cannons for some random set of protest events (you could for example propose a design that *reduces* the use of water cannons in a subset of events and examine changes to group organization). As with the first dilemma there is clearly no intention to seek consent from the subjects—in this case the protesters—as to whether they want to be shot at. Should you partner with the police and advise them on the use of water cannons in order to learn about the behavior of non-consenting subjects?

These seem like impossible choices. But choices of this form arise regularly in the context of a mode of “embedded” experimentation that has gained prominence in recent years in which experimental research is appended to independent interventions by governments, politicians, NGOs, or others, sometimes with large humanitarian consequences.

The particular problem here is that *the researcher is taking actions that may have major, direct, and possibly adverse effects on the lives of others*. As discussed below, in these embedded field experiments, these actions are often taken without the consent of subjects; a situation which greatly magnifies the ethical difficulties.

In this essay I discuss the merits and demerits of “embedded” experimentation of this form that is undertaken without subject consent. I compare the approach to one in which researchers create and control interventions that have research as their primary purpose and in which consent may be more easily attained (in the terminology of Harrison and List these two approaches correspond broadly to the “natural field experiment” and the “framed field experiment” approach).¹

¹ Glenn W. Harrison and John A. List, “Field Experiments,” *Journal of Economic Literature* 42, (December 2004):1009-1055.

The issues that arise from the embedded approach are multiple and complex and extend well beyond what is generally covered by Institutional Review Boards (IRBs) that provide a form of ethical approval for research. They are also practical questions and my discussion here should be read as the sometimes pained reflections of a researcher knee deep in these issues rather than as the considered views of a moral philosopher. The view from the trenches is quite grim in that the core questions demand answers but still seem to me largely unanswered and unanswerable.

I try to do five things in this essay.

In section 1 I seek to clarify the advantages and disadvantages of embedded field experiments. While embedded field experiments introduce ethical complexities, there may also be strong ethical reasons for employing them.

In section 2, I argue that common permissive arguments (specifically arguments based on scarcity and ignorance) that embedded field experiments do no harm are often not plausible. Not only do these experiments often involve *risks* of harm (as well as benefits), they may do so in ways that are beyond the purview of institutional review boards. I argue that in such cases researchers need to be able to argue for substantive benefits from the research that can trade off against risk if they are to satisfy a beneficence test. Of course, unless basic knowledge trades off against welfare in this calculation, the beneficence criterion, as usually applied, then places limits on what questions can be addressed using experimental interventions. I highlight that while there are often good reasons to expect important benefits from such research, in general this calculation is difficult because of the lack of a shared value metric.

In section 3, I articulate an argument—which I call the “spheres of ethics” argument—that researchers sometimes employ as grounds for collaborating in partnerships in which subjects are exposed to risks to an extent not normally admissible in research communities. The basic idea of the spheres of ethics argument is that if the intervention is ethical for implementing agencies with respect to the ethical standards of their sphere—which may differ from those of researchers—and if those agents are ethically autonomous from researchers, then responsibility may be divided between researchers and implementers, with research ethics standards applied to research components and partner standards applied to manipulations. The argument in favor of the approach is simple and strong, though thoroughly utilitarian: if the intervention may be implemented ethically by the implementer, and if the intervention with a research component is at least as good as the intervention without, then the implementation with the research component is ethical also, even if, when undertaken by researchers alone, it violates ethical standards of the research community. In the water cannon example above, the argument would be that if a police force is employing water cannons anyway, it is arguably better to know what their effects are and this might in turn justify a partnership in which the researcher is learning from the behavior of human subjects without their consent. I highlight two prior questions to be addressed if this argument is to be employed—*partner autonomy* and *partner legitimacy*. Without *partner autonomy* (for example if the researcher is also the de facto implementer wearing a different hat), there is a risk that the spheres of ethics argument could simply be used to bypass standards of the research community. Without *partner legitimacy* the spheres of ethics argument could be used to justify the kinds of experimentation that research ethics were specifically intended to prevent. Despite the attractiveness of the argument, I note that the spheres of ethics argument is incomplete insofar as addressing *partner legitimacy* requires a solution to the metaethical problem: that researchers have grounds to deem that actions that are ethical from for the partner’s sphere of action are indeed ethical.

In section 4, I consider other implications of the requirement of beneficence for researchers conducting embedded experiments including implications for relations with partners and the fact that if beneficence is claimed on the basis of the value of learning from the research, this adds an ethical imperative to the professional imperative to produce high quality research.

Finally, adopting the view that ethical principles are constructed by and for communities I use the opportunity here to introduce a set of guidelines that have been recently proposed and endorsed by the Experiments in Governance and Politics (EGAP) network of researchers; while they do not provide answers to the deeper questions they do provide some sort of benchmark for worried researchers.²

In what follows, in order to make progress in thinking about research ethics, I try insofar as I can to sidestep the metaethical problem of why one should take ethics seriously or on what basis one can even begin to claim that one action is more ethical than another. Instead I simply assume that researchers subscribe to the family of principles described in the Belmont report and in particular that they seek to respect the broad (if somewhat crudely defined) principles of beneficence, respect for persons, and justice as described in that report.³ We will see however that for some purposes—in particular the employment of a partner legitimacy test—the attempt to sidestep these issues fails.

Arguments for and Against Embedded Experiments

Let us begin with a description of the kinds of experiments of interest. Embedded experiments are experiments in which researchers form partnerships with other agents who introduce random variation into their projects to allow researchers to learn about the effects of the interventions. Often these experiments are “natural field experiments” in the sense described by Levitt and List:

“Natural field experiments are those experiments completed in cases where the environment is such that the subjects naturally undertake these tasks and where the subjects do not know that they are participants in an experiment. Therefore, they neither know that they are being randomized into treatment nor that their behavior is subsequently scrutinized. Such an exercise is important in that it represents an approach that combines the most attractive elements of the lab and naturally-occurring data: randomization and realism.”⁴

In many political science applications the naturalism arises from the fact that the intervention is implemented by a political actor—a government, an NGO, a development agency. In these cases especially, the term “randomized control trials” can be misleading since often the *research only exists because of the intervention* rather than the other way round. There are now many such field experiments of this form in the area of the political economy of development.⁵ This approach can be contrasted with a “framed field experiment” in which the

² The EGAP guidelines can be found here: <http://e-gap.org/resources/egap-statement-of-principles/>

³ Dan Harms, “The Belmont Report: Ethical Principles and Guidelines for the Protection of Human Subjects of Research”, DHEW Publication No. (OS) 78-0012, 1978.

⁴ Steven D. Levitt and John A. List, “Field Experiments in Economics: The Past, the Present, and the Future,” *European Economic Review* 53, no. 1 (2009), 9.

⁵ See for example, Marianne Bertrand, Simeon Djankov, Rema Hanna and Sendhil Mullainathan, “Obtaining a Driver's License in India: An Experimental Approach to Studying Corruption” *The Quarterly Journal of Economics* 122 (4, 2007): 1639- 1676; Macartan Humphreys, William Masters and Martin Sandbu, “The Role of Leaders in Democratic

intervention is established by researchers for the purpose of addressing a research question and done in a way in which participants know that they are part of a research experiment.

In practice, of course, the distinction between these two types of experiment is not always clear. An intervention may be established for non–research reasons, but varied for research reasons; an implementing organization may in practice be dependent on researchers, in which case researchers may be the de facto designers. A framed experiment may be implemented without knowledge of subjects and of course many experiments implemented by third parties may be undertaken with knowledge of participants that they are part of an experiment. In what follows, however, I focus on the particular problems that are manifest when both of the characteristics above are present, that is, on embedded natural field experiments that are implemented by a third party without informed consent.

Arguments for Embedded Natural Field Experiments

There are a number of benefits of the embedded natural field experimental approach relative to the framed experiment.

First there is an *internal validity benefit* to the fact that participants do not know that they are in an experiment – specifically the removal of Hawthorne effects. Second, there may also be an *external validity benefit* from the fact that the researcher did not determine many elements of the design but that these are set at levels determined “by nature”—at least insofar as the natural levels are more likely to be representative of the population of interest. Moreover, the removal of randomization biases arising from individuals refusing to take part in a study allows less problematic assessments of population effects. Third there are *enabling benefits* in that this form of experimentation may not be possible for researchers without partnerships. Partnerships may reduce costs and allow operation at a scale that is not normally feasible for researchers establishing their own interventions. But, as in the watercannon example, partners may be able to implement manipulations that would be illegal for researchers. Fourth, there may also be *epistemological benefits* from the fact that the intervention is not just *like* the class of environments of interest, but that it may in fact *be* an environment of interest; that is one might not just learn about elections in general but also be able to address questions about particular elections of importance.

Finally there is a positive and a permissive ethical reason for employing natural field experiments. The positive ethical argument is a strong one: that interventions—especially those with major consequences—should be informed by the most reliable evidence and not understanding the effects of these interventions is a failing. The permissive ethical argument (which I return to below) is that, unlike framed field experiments, the interventions in question may happen anyway, independent of research, and that interventions that are accompanied by research that allow us to assess their impacts are surely better than interventions that are not.

Arguments Against Embedded Natural Field Experiments

Before noting the ethical concerns of natural field experiments, it is worth noting two other problems. The first is that these experiments may be risky for researchers (in that interventions may collapse or variations introduced in ways that are beyond the control of researchers); the second is that compromises in control

Deliberations: Results from a Field Experiment in São Tomé and Príncipe” *World Politics* 58 (July, 2006): 583-622; Leonard Wantchekon, "Clientelism and Voting Behavior: Evidence from a Field Experiment in Benin," *World Politics* 55 (April 2003): 399-422.

may be considerable (resulting in variations that are either too complex or too modest to produce findings of significance).

The ethical complexities arise from the fact that experiments of this form risk violating all three of the ethical principles described in the Belmont report: *beneficence*, *respect for persons*, and *justice*; and they often do so without informed consent.

In a way, the lack of consent is the crux of the matter. A benefit of the principle of consent is that it instantiates respect for persons, but it also results in a *sharing* of responsibility with subjects, which is especially important when the benefits of an intervention is in doubt. But for many natural field experiments informed consent may be very incomplete. Informed consent is routinely sought for measurement purposes, for example when survey data is collected. It is sometimes sought at least implicitly for interventions, although individual subjects may often not be consulted on whether for example they are to be exposed to particular ads or whether a school is to be built in their town. But consent is often not sought for participation in the experiment per se, for example subjects are often not informed that they were randomly assigned to receive (or not receive) a treatment for research purposes, nor, often, is there a debriefing afterwards. The common absence of consent makes the question of beneficence especially difficult for researchers because the responsibility for determining beneficence cannot be shared with subjects.

Evidently, if informed consent is possible in the context of a natural field experiment, it should be employed as this would mitigate many of the concerns raised here. However there are at least three arguments for why consent might not be sought in natural field experiments. The first is that because the intervention is naturally occurring the need for consent is obviated. For example, if in the normal course of events a politician airs a variety of ads on the radio that a voter might listen to, then a systematic altering of which ads are aired when operates within the sphere of activities to which the subject has implicitly consented. The second is that because the intervention is naturally occurring an attempt to gain consent would be especially damaging. In the last example it is precisely because listening to the ad is a routine event that preceding the ad with an announcement that the ad is being aired to understand such and such an effect will have particularly adverse consequences. A third, more difficult reason, is that the withholding of consent may not be within the rights of the subjects. Consider for example a case where a police force seeks to understand the effects of patrols on reducing crime; the force could argue that the consent of possible criminals (the subjects in this case) is not required for the force to decide where to place police. This third argument is the most challenging here since it highlights the fact that consent is not even notionally required by all actors for all interventions, even if it is generally always required of researchers for subjects.

Establishing Beneficence in Embedded Natural Field Experiments

Researchers are often advised that manipulations should “do no harm” (note, not no *net* harm, but no harm at all). Two arguments are commonly used for why randomized assignment does no harm. The *scarcity argument* is that randomized assignment takes place in the context of scarcity and so randomization only affects *which* individuals of a set of equally deserving individuals are allocated benefits not *how many*. Moreover in such settings a random assignment is an ex ante equitable way of assigning scarce resources. The *ignorance argument* is that ex ante it is often not known whether a treatment is beneficial or not, indeed establishing this may be the purpose of the research in the first place.

But neither of these arguments holds up nearly so well in practice as in theory. The scarcity argument runs into problems when there are low marginal costs – such as interventions that provide information of various forms. It also comes under stress when goods are divisible. For example a cash allocation of \$100 per person may be optimal from a beneficence perspective, but to generate stronger effects, a \$200 allocation to half as many beneficiaries may be optimal from a design perspective. Factorial designs in which some subjects receive multiple benefits while others receive none seem also to give the lie to the scarcity argument. Moreover the idea that all individuals are equally needy is something of a fiction in many actual settings. Finally the introduction of randomization may itself increase scarcity. For example, if an allocation of benefits as determined by a randomization scheme is more expensive to deliver than that from a purposive scheme.

The ignorance argument is also often hard to defend. While it is certainly true that we in general cannot be *certain* whether a given treatment is beneficial or not, we nevertheless generally have prior beliefs, and in the context of many natural field experiments, the prior beliefs, at least on the side of implementing organizations, are often very strong.⁶

In practice then we generally cannot claim that experiments do no harm, even if this is sometimes the case. This does not mean however that they do not do more good than harm; or conversely, it may still be that although they may do harm, not doing them may do more harm than good. The problem here is that while the “do no harm principle,” being a negative injunction, is compelling at first blush, it is exceptionally restrictive. Almost all research carries some imaginable risk of harm; we are always in a world of tradeoffs, measuring risks against possible benefits.⁷ *Thus researchers may often need to make a more direct case for possible benefits and not just absence of harm.*

This brings us to the first question for researchers:

Q 1 The beneficence test: *Is there value in answering the research question and can that value trade off against harm?*

We can usefully distinguish between three loci of costs and benefits. The *participation* costs and benefits for subjects taking part, the broader *process* costs and benefits associated with the implementation of the research, including the immediate social effects of the intervention; and the *outcome* costs and benefits resulting from the findings of the research.

In practice Institutional Review Boards often seek to weigh the outcome benefits against the participation costs (the formula given in the Belmont report is to “maximize possible benefits and minimize possible harms”). However embedded natural field experiments involve process costs to subjects that are very distinct from the participation costs that IRBs focus on. To see the difference consider the situation we noted above.

⁶ In some variants ignorance arguments focus on the uncertainty of the researcher, in others on the uncertainty of the research community. In some statements the requirement is that there should be some uncertainty about which treatment is best—a condition which is generically satisfied in social science settings; in other statements the condition is that there should be exact indifference—a condition that is generically not satisfied if researchers have informative priors. See Freedman, B. “Equipose and the ethics of clinical research.” *The New England Journal of Medicine*, 317 (July 1987):141-145.

⁷ Cedric M. Smith, “Origin and Uses of *Primum Non Nocere*--Above All, Do No Harm!” *Journal of Clinical Pharmacology* 45 (April 2005), 371-377.

In the absence of a partnership with a researcher, an NGO plans to provide cash allocations of \$100 to 1000 needy families. With the partnership in place the NGO decides instead to randomly select 500 families to receive \$200 each. Thus, *as a result of the study design, 500 families will go without allocations*. Nevertheless, *given the study's design*, participation comes at no risk to households.

So while IRB boards will be concerned about risks and benefits of participation as well as the outcome benefits of the research relative to a situation with no intervention, a researcher concerned about ethics may be concerned about the total effect of the partnership relative to what would happen in the absence of the partnership. This is of concern because process costs from natural field experiments may be considerable; in principle variations introduced on the advice of researchers may determine who gets aid, which parties get elected to power, which officials get incarcerated, or who lives and who dies.

The rough utilitarian injunction adopted in the Belmont report to maximize net benefits might be attempted using all three types of costs and benefits. In cases in which the benefits and the costs are measured in the same units and for the same individuals (for example jobs created and destroyed in a single population) this may be easy, but more generally (for example when different populations suffer benefits and costs, as may be the case in the water cannon example) the calculation runs into difficulties: we do not have agreed criteria for assessing the value of those effects, and little reason to believe that agreement would be possible. The case for research that creates concrete risks in order to advance basic knowledge but without demonstrable benefits, seems especially difficult (another argument for why researchers should not motivate their projects by the need to simply fill a gap in the literature). While I attempted to sidestep the metaethical problem by adopting the principles of the Belmont report, metaethical problems resurface once we actually try to apply them.

In short, while Question 1 may have an easy answer, and the quality of learning from an experimental design may make the question easier to answer, in some cases researchers face an impossible decision calculus: a cost benefit calculation of some form seems to be needed but there is no agreed metric for making the calculation. Of course individuals may have their own answers to these questions to guide their decisions but as a profession we do not.

Spheres of Ethics and the Clarification of Responsibilities

In deploying embedded natural field experiments, researchers may take actions that have major consequences over outcomes that researchers have no business determining, *qua* researchers, and for which there may not be clear lines of accountability. But perhaps there are benefits if in fact in these cases researchers are *not* responsible. Can things be arranged such that the ethical responsibility for embedded experiments can be shared with partners?

Above I assumed heroically that there is basic agreement among researchers about appropriate standards of research. Say now, still more heroically that there are other standards of behavior for other actors in other spheres that are also generally accepted. For NGOs for example we might think of the INGO Accountability Charter; for governments we might think of the international treaty obligations. One might think of these ethical principles in different spheres as stemming from a single theory of ethics, or as simply the possibly incompatible principles adopted by different communities. In either case, these different standards may specify different behaviors for different actors. Thus for example a researcher interviewing a *genocidaire* in Rwanda should inform the prisoner of the purpose of the questioning and stop questioning when asked by the subject; a government interrogator could act ethically and ignore such principles, even if other behavior,

such as torture, is eschewed. Here the ethical constraints on the researcher seem stronger; but there may be stronger incompatibilities if constraints are not nested. For example a researcher may think it unethical to give over information about a subject suspected of criminal activities while a government official may think it unethical not to.

For the spheres of ethics argument, the question then is whose ethical principles to follow when there are collaborations? One possibility is to adhere to the most stringent principle of the partners. Thus researchers working in partnerships with governments may expect governments to follow principles of research ethics when engaging with subjects. In some situations, discussed below, this may be a fruitful approach. But as a general principle it suffers from two flaws. The first is that in making these requirements the researcher is altering the behavior of partners in ways that may limit their effectiveness; this runs counter to the goal of reducing the extent of manipulation. The second is that, as noted above, the constraints may be non-nested: the ethical position for a government may be to prosecute a criminal; but the researcher wants to minimize harm to subjects. In practice this might rule out appending research components to interventions that would have happened without the researcher and that are ethical from the perspective of implementers; it could for example stymie the use of experimental approaches to study a large range of government strategies without any gain, and possibly some loss, to affected populations.

An alternative approach is to divide responsibilities: to make implementers responsible for implementation and researchers responsible for the research. The principle of allocating responsibility of implementation to partners may then be justified on the grounds that in the absence of researchers, partners would be implementing (or, more weakly, that they could implement) such interventions anyhow and are capable of bearing ethical responsibility for the interventions outside of the research context.

As a practical matter researchers can do this in an underhand way by advising on interventions *qua* consultants and then returning to analyze data *qua* researchers; or by setting up an NGO to implement an intervention *qua* activist and then return for the data *qua* researcher. But this approach risks creating a backdoor to simply avoiding researcher responsibilities altogether.

Instead, by appealing to spheres of ethics, researchers collaborating with autonomous partners can do something like this in a transparent way by formally dividing responsibility. Although researchers play a role in the design of interventions it may still be possible to draw a line between responsibility for design and responsibility for implementation. Here, responsibility is understood not in the causal sense of who contributed to the intervention, but formally as who shoulders moral and legal responsibility for the intervention. Researchers hoping to employ such an argument need to be able to answer Question 2:

Q 2 Is there clarity over who is ethically responsible for the intervention?

The first of five principles endorsed by the EGAP network addresses this question:

Principle 1: [...] In cases in which researchers are engaged alongside practitioners, an agreement should state which party, if either, has primary responsibility for the intervention. Researchers should disclose the role that they play in the design of interventions implemented by practitioners or third parties.

There are two critical difficulties with the spheres of ethics approach however. The first is the *autonomy* concern: that in practice implementers may not be so autonomous from the researchers, in which case the argument may simply serve as a cover for avoiding researcher responsibilities. The second is deeper: the argument is *incomplete* insofar as it depends on an unanswered metaethical question: it requires that the researcher have grounds to deem actions that are ethical from the partner's perspective are indeed ethical—perhaps in terms of content or on the grounds of the process used by partners to construct them. This is the *partner legitimacy* concern. A researcher adopting a spheres of ethics argument may reasonably be challenged for endorsing or benefitting from weak ethical standards of partners. Indeed a version of this argument could otherwise serve as ammunition for doctors participating in medical experimentation in partnership with the Nazi government.

Given the incompleteness, researchers may still seek to use design to ensure beneficence even if responsibility for the intervention is borne by a partner. Design choices have implications for beneficence. For example in cases where selection effects may not be strong but there is clear variation in need or merit, a regression discontinuity design may be better than a fully randomized design. In some cases variation on the upside rather than on the downside of treatments may improve participant beneficence (for example is the water cannon case, variation in the use of cannons can be introduced by *reducing* the use of water cannons rather than increasing it; the former would be less harmful for subjects, but perhaps more harmful for third parties).

The question for researchers then is:

Q 3 Have variations that reduce risks and costs to subjects been examined?

The question is especially salient at a time when the recognition given to randomized experimentation in the discipline may provide professional incentives for researchers to employ it beyond what is merited by the problem at hand.

Beneficence Beyond Human Subjects

Embedded field experiments raise a set of ethical questions around partner relations and research results that factor into beneficence calculations but that are not present in other approaches and are not covered by standard Institutional Review Board considerations. Broadly these process concerns stem from extending the Belmont principles for subjects to partners and to users of research findings.

Partner Matters

Engaging in field experimentation can be very costly for partners. And if they do not have a full understanding of the research design, partners can sometimes be convinced to do things they should not. On various points of design, partners and researchers may have divergent interests. One of these is with respect to statistical power. For a partner, an underpowered study can mean costly investments that result in ambiguous findings. Underpowered studies are in general a problem for researchers too with the difference that they can still be beneficial if their findings can be incorporated into metaanalyses. Researchers may also be more willing to accept underpowered studies if they are less risk averse than partners and if they discount the costs of the interventions. Thus to account for global beneficence researchers need to establish some form of informed consent with partners and address the question.

Q 4 Do your partners really understand the limitations and the costs of an experiment?

Sharing (and explaining) statistical power calculations is one way of ensuring understanding. Another is to generate “mock” tables of results in advance so that partners can see exactly what is being tested and how those tests will be interpreted.⁸

A second concern relates to the researchers’ independence from partners. The concern is simple, that in the social sciences, as in medical sciences, partnering induces pressures on researchers to produce results that make the partner happy. These concerns relate to the credibility of results, a problem I return to below. The problems are especially obvious when researchers receive remuneration but they apply more generally and may put at risk the quality of the research.

Q 5 Can you demonstrate independence of the research from the implementation?

The third and fifth principles endorsed by the EGAP group propose guidelines for ensuring and demonstrating independence.

Principle 3: Rights to Review and Publish Findings. In collaborations between researchers and practitioners it should be agreed in advance, and not contingent upon findings, what findings and data can be used for publication. In cases in which such agreement is not made in advance, and unconditional on findings, this fact should be noted in publications.

Principle 5: Remuneration. Researchers should normally not receive remuneration from project implementers whose projects they are studying. In cases in which researchers receive remuneration from such agencies, this fact should be disclosed in footnotes to publications.

Users: Quality of Research Findings

Finally, part of the consideration of beneficence involves an assessment of the quality of the work and the lessons that can be drawn from it. If an argument in favor of a research design is that the lessons from the research produces positive effects, for example by providing answers to normatively important questions, then an assessment of beneficence requires an expectation that the design is capable of generating credible results.⁹

There are clearly many aspects to the quality of research but here I would like to point to one area in which basic standards are not at present being met. The question for researchers is:

Q 6 Are you really testing what you say you are testing?

⁸ We used this approach in our Congo study. See: <http://cu-csds.org/2011/03/drc-design-instruments-and-mock-report/>.

⁹ This line of reasoning is contestable, although it appears important to claim beneficence. Arguably researchers should not be in the business of trying to estimate the outcome costs and benefits of the impact of their work beyond the participation and process costs and benefits. Thus for example the injunction to go where the truth leads scorns such weighing of costs and benefits, on the optimistic presumption that the truth is in league with the good. (For a classic articulation see Thomas Jefferson on the University of Virginia: "This institution will be based on the illimitable freedom of the human mind. For here we are not afraid to follow truth wherever it may lead, nor to tolerate any error so long as reason is left free to combat it." Cited in Andrew Lipscomb and Albert E. Bergh, eds. *The Writings of Thomas Jefferson*. Washington, D.C.: Thomas Jefferson Memorial Association of the United States, 1903-04. 20 vols.

The question seems obvious but in fact post hoc analysis is still the norm in much of political science and economics. It is almost impossible to find a registered design of any experiment in the political economy of development (for an exception see the work of Casey and colleagues on Sierra Leone¹⁰). This raises the serious concern that results are selected based on their significance, with serious implications for bias.¹¹ It is obvious but worth stating that *research designs that create risks cannot claim beneficence on the basis of their potential findings when those findings are not credible.*

Two EGAP principles seek to address these concerns:

Principle 2: Transparency: To maintain transparency and limit bias in reporting, researchers should seek to register research designs, hypotheses and tests in advance of data collection and analysis. In presentation of findings, researchers should distinguish between analyses that were planned ex ante and those that were conceptualized ex post.

Principle 4: Publication of Data: In collaborations between researchers and practitioners, researchers and practitioners should agree in advance that data used for analysis will be made publicly available (subject to masking of identifiable information) for replication purposes within a specified time period after data collection.

Conclusion

Are we now in a better position to face the dilemmas I introduced at the beginning of this essay? I think both dilemmas remain hard but that the considerations here help focus on the key issues. The difficulty in both dilemmas is that design decisions have impacts on the lives of populations but the exposure of populations to different treatments is done without their consent. In both cases however there may be grounds for employing a spheres of ethics argument to justify researcher participation. This refocuses attention on the issues of partner autonomy—which may be especially important in the first case—and partner legitimacy—which may be especially important in the second.

For the political information problem, the first option, in which the research is essentially abandoned may seem ethically the most defensible option. It seems most faithful to the injunction to do no harm. This is the option that I and my colleagues have advised when confronted with problems like this in the past. But I am not sure that we were always right to do so: sacrificing the research in this way may not be ethically the best option. It effectively assumes that the research is of no ethical import, which, if true, puts the ethical justifiability of the original research plan in question. The second option brings direct costs to populations of a form not covered by normal human subjects considerations. Done without their consent however the principles of beneficence and justice normally suggest that this option would require that the lessons from the research plausibly produce effects that are commensurable with these costs; a hard calculation. The comparison of options then requires a calculation for which researchers are unlikely to have defensible metrics.

¹⁰ Casey, Katherine, Rachel Glennerster and Edward Miguel. “Reshaping Institutions: Evidence on External Aid and Local Collective Action.” NBER Working Paper, 17092 (2011).

¹¹ Gerber, Alan and Neil Malhotra. “Do Statistical Reporting Standards Affect What Is Published? Publication Bias in Two Leading Political Science Journals.” *Quarterly Journal of Political Science*, 3 (3, 2008) 313–326.

The argument presented here suggests that handling this problem hinges on autonomy. If in practice the intervention exists to support the research then the experiment should be treated as a framed field experiment and the usual requirements for consent sought by IRB boards should be applied. If however the grounds for not having consent by subjects is that the intervention is being implemented by an autonomous third party, that can bear ethical responsibility for the decision, then the researcher can employ a spheres of ethics argument to justify stopping the intervention of the second group. Autonomy is key here: under this argument, the choice of strategy is not for the researcher to make but for the implementing partner subject to their adhering to ethical criteria of their sphere. The responsibility of the researcher is a different one: to provide conditions for the implementer's decision to be made in an informed way and for the ensuing research to be credible.

The second dilemma turns out to involve many similar considerations. In this case the research could not legally be implemented as a framed experiment, or without a partnership with government. The problem here however is not that the proposed modification reduces benefits to subjects—indeed under some designs it may reduce risks—but that even with these reductions, the collaboration involves learning from harmful manipulations to subjects that are undertaken without their consent. It is unlikely that the researcher can make a simple argument for beneficence for subjects in this case; even if there are benefits, these may accrue entirely to government and not to the subjects. The researcher might however employ a spheres of ethics argument in this case; but, assuming autonomy of government, engaging in the partnership shifts focus to the prior, and unanswered, question of the basis for the legitimacy of the police to decide whether to take actions of this form in the first place.