

Open Access

Macartan Humphreys*

Reflections on the Ethics of Social Experimentation

Abstract: Social scientists are increasingly engaging in experimental research projects of importance for public policy in developing areas. While this research holds the possibility of producing major social benefits, it may also involve manipulating populations, often without consent, sometimes with potentially adverse effects, and often in settings with obvious power differentials between researcher and subject. Such research is currently conducted with few clear ethical guidelines. In this paper I discuss research ethics as currently understood in this field, highlighting the limitations of standard procedures and the need for the construction of appropriate ethics, focusing on the problems of determining responsibility for interventions and assessing appropriate forms of consent.

Keywords: consent; ethics; field experiments; human subjects.

JEL classification: C90; C93.

DOI 10.1515/jgd-2014-0016

1 Introduction

Social science researchers are increasingly using field experimental methods to try to answer all kinds of questions about political processes and public policies. Unlike traditional “observational” methods, in which you observe the world as it comes to you, the idea right at the heart of the experimental approach is that you learn about the world by seeing how it reacts to interventions. In international development research these interventions can sometimes take the form of researchers from wealthy institutions manipulating citizens from poorer populations to answer questions of little interest to those populations.

These studies raise a host of ethical concerns that social scientists are not well equipped to deal with. US based social science researchers rely on principles

*Corresponding author: **Macartan Humphreys**, Department of Political Science, Columbia University, e-mail: mh2245@columbia.edu

such as respect for persons, justice, and beneficence that have been adopted by health researchers and institutionalized through formal review processes but that do not always do the work asked of them by social scientists.

Consider one example where many of the points of tension come to a head. Say a researcher is contacted by a set of community organizations that want to figure out whether placing street lights in slums will reduce violent crime. In this research the subjects are the criminals: seeking informed consent of the criminals would likely compromise the research and it would likely not be forthcoming anyhow (violation of respect for persons); the criminals will likely bear the costs of the research without benefitting (violation of justice); and there will be disagreement regarding the benefits of the research – if it is effective, the criminals in particular will not value it (producing a difficulty for assessing benevolence). There is no pretense at neutrality in this research since assessing the effectiveness of the lamps is taking sides, but despite the absence of neutrality no implicit contract between researchers and subjects is broken. The special issues here are not just around the subjects however. Here there are also risks that obtain to non-subjects, if for example criminals retaliate against the organizations putting the lamps in place. The organization may be very aware of these risks but be willing to bear them because they erroneously put faith in the ill-founded expectations of researchers from wealthy universities who are themselves motivated in part to publish.

The example raises a lot of issues. It is chosen because despite the many issues raised, the principles that are currently employed provide almost no guidance to deal with the issues raised. It is not however a particularly unusual case and many of the features of the case are shared by other projects including work in spheres such as reduction of violence against women, efforts to introduce democratic institutions in rural communities, job training programs for ex combatants, efforts to alter electoral behaviour of constituents, and efforts to stamp out corruption by politicians [for a discussion of many relevant cases see Baele (2013)]. Unlike classic health and education interventions, these projects routinely deal with interventions that have winners and losers, create risks for some, and are done without the consent of all parties affected by them.

The absence of clear principles to handle these issues leaves individuals and the professions in a difficult situation, at least if they care about the ethical implications of their research designs above and beyond whether they receive formal research approval.

So how should researchers proceed in these cases? At present there are no satisfactory answers. To make progress I discuss three sets of problems raised by research designs like this, which I call the problem of *audience*, the problem of *agency*, and the problem of *consent*.

The *audience question* is about determining what the *professional* ethical issues are. My focus throughout will be on professional ethics rather than more metaphysical questions of what is right or wrong in some objective sense. Thus in Section 2, I highlight a conceptualization of the problem not as a problem of normative ethics – whether any of these designs are right or wrong in any fundamental sense – but as a question of audience. A key purpose of professional ethics is to clarify expectations of members of a profession for relevant groups that are important to their work. For medical ethics the key audience is patients, or particularly subjects: those patients with which medical professionals engage. The current guidelines used by social scientists are inherited from medical ethics, which place a primary focus on human subjects. While subjects perhaps represent the primary audience for medical interventions, this may not be the case for social science interventions for which the key audience can be the general public or policy. This section highlights the need for the construction of an ethics that addresses the preoccupations of social scientists engaging in this type of research. It also highlights the more thorny nature of this problem for interventions in which there are winners and losers, as in the motivating example above.

The *agency problem* is the problem of determining who is *responsible* for manipulations. I discuss this in Section 3, describing 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 the course of research projects. The key idea is that if an intervention is ethical for implementing agencies with respect to the ethical standards of their sphere – which may differ from the ethical standards of researchers – then responsibility may be divided between researchers and implementers, with research ethics standards applied to research components and partner standards applied to manipulations. Put crudely this approach can be considered a way of passing the buck, but in fact the arguments for employing it are much more subtle than that. In a way, the buck-passing interpretation fundamentally misses the point of professional ethics. Even still, this argument is subject to abuse and so this section outlines protections related to agency autonomy and legitimacy which in turn depend on the conceptualization of professional ethics described in Section 2.

The third problem is the critical problem of *consent*. The bulk of this essay focuses on consent and the role it plays in research ethics. Current norms for informed consent are again inherited from medical ethics and reflect answers in the medical community to the first two questions. Yet alternative conceptualizations of consent are possible, and may be more appropriate for social scientists, given the different answers to questions of audience and agency in social science research. I outline a range of these in Section 4.

I close with reflections on implications for practice and for the development of ethical standards that can address the issues raised by experimental research in social science.

2 Problem 1: Audience

What are we worrying about when we worry about whether implementing experiments like that described above is ethical? It often seems as though we are worrying about whether in some fundamental sense these research activities are right or wrong. But framing the question in that way renders it largely unanswerable. The more practical approach of professional ethics is to determine whether one or another action is more or less consistent with the expectations of a relevant “audience” regarding the behaviour of the members of the profession.¹

While the response that ethical action is action that is in line with expectations of a relevant audience is not technically question begging, it does require the existence of some recognized set of norms for a profession. In practice, social scientists largely work within the ethical framework provided by the human subjects protection system.² The system was devised primarily with a view to regulating medical research, but now covers all research involving human subjects, at least for US based researchers or researchers receiving federal funding.

The principles embedded in the Belmont report³ and that permeate the work of Institutional Review Boards in the United States self-consciously seek to prescribe a set of common expectations for a community of researchers and their patients and clients. Indeed, sidestepping the question of ethical foundations seems to have been a strategy of the US Commission that produced these reports.⁴ The pragmatic approach adopted by the commission is a strength. As argued by Jonsen (1983), medical ethics, as captured by the documents produced by the

¹ For some scholars, e.g., Binmore (1998), DeScioli and Kurzban (2013), the normative question is also a problem of consistency between actions and expectations, even if it is not always recognized as such.

² At least for US based researchers. Elsewhere authorities have set up more directly focused on social science research. See for example CONEVAL in Mexico.

³ Harms (1978).

⁴ Abram and Wolf (1984) cite Congress as expressing a distaste for establishing a substantive morality; they argue that the “structure Congress gave the commission helped prevent the adoption of any highly abstract ethical theory... agreement on a fundamental moral system was not sought or needed.”

Commission, is “a Concord in Medical Ethics,” a concord “reached by a responsible group drawn from the profession and from the public.”

But this pragmatic approach also limits the pretensions to universality of research ethics in an obvious way. The principles of the Belmont report were developed to address particular problems confronting the medical profession that carry authority because they were developed through a deliberative process that sought to reach consensus in the profession around conventions of behaviour. The result is both elegant in sidestepping the unanswerable questions and messy in its result. The final principles are a mixture of deontological and consequentialist principles, with no overarching principle to refer to to determine what kinds of tradeoffs should be made in cases where interventions that benefit one group harm another. The practical solution is to outsource the problem of making these determinations to the judgments of individuals placed on university institutional review boards. While effective for some purposes, there is *ex ante* no reason to expect that the principles developed provide the appropriate guidelines for social science.⁵

The poor fit stems in part from the fact that medical research differs from social science research in various ways.

Unlike many health scientists, social scientists are commonly working on problems in which:

1. researchers are interested in the behaviour of institutions or groups, whether governmental, private sector, or nongovernmental, and do not require information about individuals (for example if you want to figure out if a government licensing agency processes applications faster from high caste applicants than from low caste applicants)
2. those most likely to be harmed by an intervention are not the subjects (for example when researchers are interested in the behaviour of bureaucrats whose decisions affects citizens, or in the behaviour of pivotal voters, which in turn can affect the outcome of elections)
3. subjects are not potential beneficiaries of the research and may even oppose it (for example for studies of interventions seeking to reduce corruption in which the corrupt bureaucrats are the subjects)
4. consent processes can compromise the research (for example for studies that seek to measure gender or race based discrimination by landlords or employers)

⁵ There are, in addition to the differences in ethical considerations, there are practical concerns. The arrangement whereby social scientists fall under this collective agreement put together by health researchers is in some respects an unhappy one. Levine and Skedsvold (2008) note for example the common complaint that IRBs do not understand the methods used by social scientists and that reviews for non sensitive research are onerous on social scientists.

5. there is disagreement over whether the outcomes are valuable (compare finding a cure for a disease to finding out that patronage politics is an effective electoral strategy); indeed some social scientific interventions are centered on the *distributive* implications of interventions: when different outcomes benefit some and hurt others, the desideratum of benefitting all that are implicated by an intervention is unobtainable
6. there is no expectation of care between the research subjects and the researcher

These features can sometimes make the standard procedures used by Institutional Review Boards for approving social science research irrelevant or unworkable.

The first two differences mean that formal reviews, as currently set up, can ignore the full range of benefits and harms of research or do not cover the research at all. Formal reviews focus on human subjects: living individuals about whom investigators obtain data through intervention or interaction or obtain identifiable private information.

The third and fourth, which again focus on subjects rather than broader populations, can quickly put the principles of justice and respect for persons – two of the core principles elaborated in the Belmont report (upon which standard review processes are based) at odds with research that may seem justifiable on other grounds.

The fifth difference can make the third Belmont principle, beneficence, unworkable, at least in the absence of some formula for comparing the benefits to some against the costs for others (see Baele 2013 on the difficulties of applying beneficence arguments).

The sixth difference means that the stakes are different. If a health researcher fails to provide care for an individual in a control group, this may violate their duty of care and break the public trust in their professions. This may not be true for social scientists however.

Thus, standard considerations inherited from the human subjects protection system can be blind to the salient considerations for social science researchers and their primary audiences. The focus on private data and the protection of subjects may sometimes seem excessive; but the blindness to the risks for non-subjects may be more costly. Specific risks, beyond welfare costs, are that researchers gain a reputation for providing unsound advice to government officials on sensitive issues, encourage the withholding of benefits from the public, interfere with judicial processes, or put vulnerable (non-subject) populations at risk, in order to further research agendas.

Refocussing on the question of audience however can give some guidance here. A preoccupation of medical ethics is the maintenance of relations of trust

between medical professionals and patients. In this sense, patients are a key audience for medical ethics.⁶ Patients can expect care from medical professionals no matter who they are. But the nature of social science questions puts researchers in different relations with subjects, most obviously when interventions are interventions aimed *against* subjects. It seems improbable that social scientists can maintain relations of trust with corrupt politicians, human rights abusers, and perpetrators of violence when the interventions they are examining are designed precisely to confront these groups.

What audiences are most critical for social scientists? Subjects are of course a key audience for social scientists also, not least because for much data collection depends on the trust, generosity, and goodwill of subjects. But two wider audiences are also critical and the fashioning of social science research ethics for field experimentation should focus closely on these. The first are research partners and the second are research consumers.

2.1 Partner Matters

As in the example above, much field experimentation can involve partnerships with local governmental or nongovernmental groups. Partnering in experimental research can be very costly for partners however. And if they do not have a full understanding of the research design, partners can be convinced to do things not in their interests which is a risk when the interests of partners and researchers diverge. One point of divergence 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 useful 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*. At a minimum this requires establishing that partners really understand the limitations and the costs of an experiment.

One useful practice is to sign a formal Memorandum of Understanding between the researcher and the partner organization at the beginning of a project laying out the roles and responsibilities of both parties. However, even when they exist, these rarely include many of the most important elements that

⁶ As highlighted by Alderman et al. (2013) the distinction between medical and social science researchers may not be so clear for subjects and so subjects that are asked health related or other questions by social scientists may in fact expect that questioning will be followed up by action.

researchers are required to provide to subjects during the informed consent process, such as the potential risks or alternatives to experimentation. These documents could even include discussions of the power of a study to ensure that partners are aware of the probability that their experiment will result in unfavourable findings, even if their program has a positive impact. Having clearer standards for what information should be required before a partner consents to an experiment could facilitate continued positive relationships between researchers and partners.

In addition, concern must be given to how researchers explain technical information to partners. The informed consent process with research subjects defines additional precautions that must be taken to obtain consent from people with limited autonomy. Similarly, there is a burden on researchers to explain the risks and benefits of technical choices to partners in layman's terms. Alderman et al. (2013) highlight the false expectations that subjects can have when they engage with researchers coming from privileged institutions and the responsibilities that this can produce. A similar logic can be in operation for partner organizations. 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.⁷

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 the quality of the research at risk. But the lack of independence cuts the other way also: if staff in partner organizations depend on researchers for access to expertise or funding, this may generate conflicts of interest for them in agreeing to implement some kind of research or other.

One way that independence can be increased is through separation of funding: when researchers are not remunerated for conducting experimental evaluations, they may be freer to report negative results. Another is to clarify from the outset that researchers have the right to the data and the right to publish the results no matter what the findings are. However, even when these measures are taken, there may be psychological or ideological reasons that researchers might still not be fully independent from partners.

⁷ We used this approach in our Congo study. See: <http://cu-csds.org/2011/03/drc-design-instruments-and-mock-report/>.

2.2 Users: Quality of Research Findings

Given the fact that field experiments can impose costs on some groups, including subjects, assessing the beneficence of a study is especially tricky. A part of the consideration of beneficence however 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 produce 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 (Baele 2013).⁸ In practice though researchers sometimes defend research that involves potential risks on the basis of the gains from knowledge there is rarely any kind of systematic accounting for such gains and rarely a treatment of how to assess these gains when there are value disagreements. Moreover researchers, given their interests in the research, are likely the wrong people to try to make this determination. Nevertheless, any claim based on the value of the findings needs to assume that the findings are credible.

The credibility of research depends on many features. I would like to draw attention to one which is the loss in credibility that can arise from weak analytic transparency. *Post hoc* analysis is still the norm in much of political science and economics. Until recently it has been almost impossible to find a registered design of any experiment in the political economy of development (in the first draft of this paper I pointed to one study; there are now close to 200 pre-registered designs housed on the EGAP registry (109), RIDIE (37), and AEA registry (49)). When experiments are not pre-registered there may be concerns that results are selected based on their statistical significance or the substantive claims they make, with serious implications for bias (Gerber and Malhotra 2008; Casey et al. 2012).

As research of this form increases in prominence, there will be a need to develop principles to address these questions of audience. For this, social scientists might

⁸ 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 Lipscomb and Bergh, (1903)].

follow the lead of the National Commission that established the principles for health research and seek not to root assessments of what is or is not ethical research in conflicting moral intuitions or on normative theories that may or may not be broadly shared. Instead in response to the issues raised by field experiments, social scientists could initiate a public process to *decide* what should constitute expected practice in this field in light of the interests of the audiences specific to their research – notably partners, governments, and the general public.⁹

3 Problem 2: Agency

In the example above of an experiment on street-lighting the intervention was initiated and implemented by a local organization and not by the researchers. Is this fact of ethical relevance for researchers taking part in the experiment?

Currently many social science experiments are implemented in this way by political actors of various forms such as a government, an NGO or a development agency. In these cases, and unlike many medical trials, research often only exists because of the intervention rather than the other way round.¹⁰ This approach can be contrasted with a “framed field experiment” in which the 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.¹¹ In practice, of course, the distinction between these two types of

⁹ One could imagine this being done either within or outside the current framework. Levine and Skedsvold (2008) propose further IRB decentralization to allow disciplines to take a lead in making determinations for their fields. This proposal has the advantage of ensuring that determinations are more likely to be made by members of the profession in question, but has the disadvantage perhaps of rendering co-ordination of expectations difficult. Falling within a large scheme has advantages since it helps clarify expectations not just among researchers, but also for the public. If there is a problem, it is not so much that the guidelines inherited from the human subjects protection system do not co-ordinate expectations, but that they do not do so in the most useful ways. The principles of the Experiments in Governance and Politics network (EGAP) provide one example of such collectively formed principles. <http://e-gap.org/resources/egap-statement-of-principles/>

¹⁰ The term “randomized control trials” can be misleading for the simple reason that these are often not trials.

¹¹ In health research the distinction maps partially to the distinction between research, quality improvement, and programme evaluation. As discussed for example by Tolleson-Rinehart (2008), however, these distinctions are fraught since they depend on the notion that it is possible to generate systematic evidence for reliable relations but that the knowledge generated not be generalizable.

experiment is often not clear,¹² even still it raises an important point of principle: can things be arranged such that the ethical responsibility for experiments can be shared with partners?

Assume heroically that there is agreement among researchers about appropriate standards of research. Say now, still more heroically, that there are other standards of behaviour 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 international treaty obligations. One might think of these ethical principles in different spheres as stemming from a single theory of ethics, or simply as the possibly incompatible principles adopted by different communities. In either case, these different standards may specify different behaviours for different actors. Thus for example by the ethical principles of research, a researcher interviewing a *genocidaire* in Rwanda should seek fully informed consent prior to questioning and stop questioning when asked by the subject or if they sense discomfort on the part of the subject. However, a government interrogator might not, but still act ethically according to the principles adopted by governments by eschewing other behaviour, such as torture. In this example, the ethical constraints on the researcher seem more demanding. There may be more intractable 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.

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 behaviour of partners in ways that may limit their effectiveness. 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

¹² An intervention may be established for non-research reasons, but still shaped by research considerations in various ways; 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. See for example, Bertrand et al. (2007), Wantchekon (2003).

implementers; it could for example prevent 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. This is what I call above the “spheres of ethics” argument. 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.

Quite distinct rationales for this approach are that partner organizations may be better placed to make decisions in the relevant areas and may be more effectively held to account if things go wrong. In addition partners may be seen by others as having legitimacy to take actions which might (correctly) be seen as meddling by outsiders (see Baele (2013) on the “Foreign Intervention problem”).

As a practical matter researchers can do this in an underhand way by advising on interventions *qua* consultants and then returning to analyse 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 for 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.

An argument against the spheres of ethics approach is that it is simply passing the buck and not engaging with the ethical issues at all. But this response misses the point of professional ethics; professional ethics is not about what outcomes should obtain in the world but about who should do what. Allocating responsibility to partners is no more buck-passing than calling on police to intervene in a threatening situation rather than relying on self-help.

The sphere of ethics approach is consistent with ideas in medical research for assessing non-validated practice. On this issue the Belmont report notes: “Research and practice may be carried on together when research is designed to evaluate the safety and efficacy of a therapy. This need not cause any confusion regarding whether or not the activity requires review; the general rule is that if there is any element of research in an activity, that activity should undergo review

for the protection of human subjects.” In terms of the standards to be applied in such a review, however, Levine (1988) notes: “the ethical norms and procedures that apply to non-validated practice are complex. Use of a modality that has been classified as non-validated practice is justified according to the norms of practice. However, the research designed to develop information about the safety and efficacy of the practice is conducted according to the norms of research.”

Levine’s interpretation of the division of labour appears consistent with the spheres of ethics approach. But the approach raises at least two critical difficulties. The first is a problem of implementer *autonomy*. In practice implementers may not be so autonomous from the researchers, in which case the spheres of ethics 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 normative 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 benefiting from weak ethical standards of partners. Indeed without an answer to this question, any collection of people could engage in any action which they claim to be ethical with respect to their “sphere;” a version of this argument could for example serve as grounds for doctors participating in medical experimentation in partnership with the Nazi government.

In line with the principle of socially constructed professional ethics, described in Section 2, a solution might be the formal recognition by the professions of classes of legitimate partners for various spheres – such as all governments, or all governments satisfying some particular criteria. The incompleteness of the spheres of ethics argument then adds urgency to the need for an answer to the problem of audience.

4 Problem 3: Consent

Medical ethics places considerable focus on the principle of informed consent, and indeed consent can in principle allay the twin concerns of audience and agency discussed in Sections 2 and 3: If the relevant audience provides consent then the expectations of the audience are arguably met and there is also a clearer allocation of responsibility for action. Both of these arguments confront difficulties however. Moreover different conceptualizations of audience and agency have different implications for consent.

The US National Commission motivated the principle of consent as follows:

Respect for persons requires that subjects, to the degree that they are capable, be given the opportunity to choose what shall or shall not happen to them... there is widespread agreement that the consent process can be analyzed as containing three elements: information, comprehension and voluntariness.

In promoting the concept of consent, the commission also sought to produce definitional clarity around it. Whereas the terms can mean many things in different settings, as described by Levine (1988), “the Commission [...] abandoned the use of the word “consent,” except in situations in which an individual can provide “legally effective consent” on his or her own behalf.”¹³

In practice however in many social experiments, consent is very imperfect. imperfect 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 even if consent for exposure to a treatment is sought, individual level consent may not be 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.¹⁴

To assess how great a problem this is, it is useful to consider the rationales for informed consent that inspired medical professionals and other rationales that may be relevant for social scientists.

4.1 The Argument from Respect of Persons

The argument provided for informed consent in the Belmont report and related documents is the principle of “respect for persons.” Manipulating subjects

¹³ Critical elements of consent are that subjects understand treatments they will receive and have freedom to agree or not to those treatments. In some accounts consent requires providing subjects with very extensive detail about what will happen to them. According to the Nuremberg code for example, a subject should have “sufficient knowledge and comprehension of the elements of the subject matter involved as to enable him to make an understanding and enlightened decision. This latter element requires that before the acceptance of an affirmative decision by the experimental subject there should be made known to him the nature, duration and purpose of the experiment; the method and means by which it is to be conducted; all inconveniences and hazards reasonably to be expected; and the effects upon his health or person which may possibly come from his participation in the experiment.”

¹⁴ Moreover it is in general not known, even if consent is sought, to what extent informed consent is achieved. Even in non-research settings for major or invasive operations, the quality of informed consent is likely very poor. Cassileth et al. (1980).

without their consent diminishes their autonomy and instantiates a lack of respect. Consent, conversely, can serve two functions.

The first is *diagnostic*: that consent can provide a test of whether people are in fact being used “merely as ends.”¹⁵ Critically, this *diagnostic* function of consent can in principle be achieved without actual consent; though actual consent eliminates the need for guesswork.

The second is *effective*: that consent may enhance autonomy (or conversely, forgoing consent reduces autonomy). Thus the Belmont report advises the importance of maximizing the autonomy of subjects: “Respect for persons requires that subjects, to the degree that they are capable, be given the opportunity to choose what shall or shall not happen to them.” There are multiple aspects of autonomy that may be affected by engagement with an experiment, with somewhat different implications for what is required of consent. I distinguish here between three: *participation autonomy*, *behaviour autonomy*, and *product autonomy*.¹⁶

The first, *participation autonomy*, relates to the decision of whether or not to be involved with the research. The absence of choice reduces subject autonomy at least with respect to the decision to take part. *Behavioural autonomy* may be compromised due to lack of consent because of information deficits (see example below) resulting in subjects making decisions that they would not otherwise make, given the options available to them. Behavioural autonomy can also be compromised if individuals’ choice sets are constrained because of the manipulation. Third, as a subject’s actions yield a research product, a lack of consent means that the subject loses control over how their labour is to be used, or a loss of *product autonomy*.¹⁷ To illustrate: say an intervention broadcasts information

¹⁵ Note I emphasize a *diagnostic* interpretation rather than a *definitional* interpretation. As noted by Pallikkathayil (2010), there is a risk of inducing circularity in the definition of being a means to an end, at least if we seek to impute consent on the basis of what a person would rationally consent to. Pallikkathayil invokes the example of determining whether shoving someone out of the way is using them as a means to an end. Answering the question requires determining if the person would rationally consent to this action which they should only do if it is not treating them merely as a means to an end.

¹⁶ In principle it is possible to have one of these types of autonomy without the other. I may not give consent to a treatment but nevertheless feel that the treatment has no bearing on my actions; in this case the loss in autonomy stems only from the lack of consultation. Alternatively a subject may feel adequately consulted but nevertheless feel that a treatment infringes on agency; and this may be the case even if the subject provides continuing consent. Or a subject might suffer participation and behaviour losses in autonomy but, when if consent is sought post factor, product autonomy may be restored.

¹⁷ *Product autonomy* seems at once the most important and the least important of these. On the one hand the complicity in producing a research product is what appears in some settings to invoke the need for consent. Consider an experiment in which organizations vary which political

about political performance on the radio in order to assess how the information alters voting behaviour by the politician's constituents. Done without consent, the listeners had no option but to take part in the study (participation autonomy), their subsequent actions are affected by the treatment and might have been different had they known the information was provided for research purposes (behavioural autonomy) and they will have no say in the publication of knowledge that is derived from their actions (product autonomy).

A problem with this formulation is that consent, or even notional consent, is not clearly either a necessary or sufficient condition for respect for persons. That is, unless respect for persons is defined in terms of consent (rather than, for example, a concern with the welfare or capabilities of others), the diagnostic function of consent as described above faces difficulties. There is a logical disconnect between consent and respect since determining respect requires information about the disposition of the *researcher* but consent provides information on the disposition of the subject. Consent might not be a necessary condition for establishing respect for persons since it is possible that the subject would never consent to an action that is nevertheless taken by a researcher with a view to enhancing their welfare or their capabilities. And of course, subjects may consent to actions not in their interests and not consent to other actions that are, or they may unknowingly take actions that limit their autonomy. The specific markers sometimes invoked to indicate that respect for persons is violated, such as the use of deceit or force, also suffer difficulties since one can construct instances in which a deceived person can recognize that deceit was necessary to achieve a good in question.¹⁸ In addition, consent might not be sufficient since it is possible that a subject consents to an action that is not being done *because* it is in their interest, but nevertheless has their welfare as a byproduct.

ads are aired in different locations. If this is done without any engagement with researchers but just as part of political campaigning then under current norms no consent issues arise. If this is done as part of research the consent requirement arises because the individuals are being *used* to make a product but not being consulted along the way. Thus lack of product autonomy appears to be at the heart of the matter. On the other hand there is generally no expectation that consent is required for using "found" information from behaviour observed in a public setting. Although people's histories are being turned into products consent is generally not called for. One key difference is that in the second case the product is in fact a *byproduct*.

18 In practice in everyday interactions we expect deceit in specified fora. I tell a car salesman that I am only willing to offer US\$5000 and she agrees to the sale. But she would not have consented to the sale at that price had she known that I would have been willing to go as high as US\$6,000. Kant (1788) would have disapproved of my lie to the car salesman (I am abusing the institution of language) but my bargaining strategy is consistent with a position in which I have respect for the salesman, but as a strategic actor.

Consider again the three types of autonomy that are threatened by an incomplete consent process. Loss in *participation autonomy* does not necessarily imply that individuals are treated simply as ends. Holding a surprise birthday for a friend deliberately compromises participation autonomy in order to provide a benefit for the friend – one that they might consent to if only the consent did not destroy the surprise.¹⁹ In some situations, where providing consent may put individuals at risk, not seeking consent may even increase participation autonomy by providing the choice to participate *de facto* or not even if risks make formal consent impossible. Even in the absence of consent however it is possible that participation in an experiment enhances *behaviour autonomy* either by expanding information or by expanding choice sets. *Product autonomy* can be restored by *ex post* consent, for example allowing a subject to determine whether they want data collected from them to be used in an analysis. Thus consent, as currently required, does not seem to be necessary or sufficient for the work asked of it.

4.2 Other rationales for Consent

Beyond respect for persons there are at least three other rationales, one nonethical and two ethical, for seeking consent for research with human subjects (not counting the desire to abide by current norms):

1. **Legal protection from charges of abuse:** A nonethical reason for seeking consent is to protect researchers from civil or criminal charges of abuse. For medical trials, the need for protection is obvious since actions as simple as providing an injection involve physical injury, which would under normal circumstance have criminal implications.²⁰ Consent clarifies that the action is non-criminal in nature (although this depends on the action – consent to be killed does not generally protect the killer). The rationale for *documenting* consent is primarily legal. As noted by Levine, HEW regulations “require that if there are risks associated with research then ‘legally effective informed consent will be obtained... The purpose of documenting consent

¹⁹ The birthday surprise example is also of interest because the action could possibly also fail the diagnostic text even though it treats someone as an end in themselves. You might know that under no circumstance would your friend consent to having people go to the trouble of organizing a party for them, but still be very pleased when it happens.

²⁰ The recognition of the importance of legal protection has a long history. When the Neisser scandal broke in 1899, government inspection of the case focused on consent processes from an ethical perspective; in addition however, according to Vollmann and Winau (1996), the minister responsible, “also sought legal advice on the Neisser case. Lawyers stated that conducting non-therapeutic research on a subject without consent fulfilled the criteria for causing physical injury in criminal law.”

on a consent form is [...] to protect the investigator and the institution against legal liability” (Levine 1979).

2. **Information aggregation/subject filtering:** Consent may also provide researchers with information regarding the relative costs or benefits of an intervention. If a researcher discovers that an individual is unwilling to take part in a study, this provides information on the perceived benefits of the study. In such cases there are double grounds not to proceed, not just because it compromises autonomy but also because it violates beneficence. As discussed below however, this goal of information aggregation may be met at a population level by seeking consent from a subset of potential subjects.
3. **Maintaining the reputation of the academy:** A third rationale for consent is that consent preserves the reputation of the academy. It clarifies to the public the nature of relations between researchers and populations, that this relation is based on respect, and that populations should not expect that their trust in researchers will be abused or that they will be put at risk without consent. Though clearly of pragmatic benefit to the academy this argument is ethical insofar as it reflects a standard of behaviour that is expected of a particular group. Note that this argument, more than any of the others, provides a rationale for ethical standards specific to researcher-subject relations that maintain higher standards than is expected of general interactions.

In the context of naturally occurring field experiments, there are also arguments for why consent might *not* be sought.

One is that because the intervention is naturally occurring, an attempt to gain consent would be intrusive for subjects and especially damaging for research. Consider for example an experiment that focuses on the effects of billboard ads. In this experiment it is precisely because seeing government ads is a routine event that preceding (if that is possible) viewing of the ad with an announcement that the ad is being posted to understand such and such an effect will have particularly adverse consequences. Preceding the ad with a disclaimer may moreover falsely suggest to subjects that some unusual participation or measurement is taking place, even if a purpose of the disclaimer is to deny it.

A second, 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, and indeed is undesirable, for the force to decide where to place

police. This argument is the most challenging 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. In this example the police can argue that the subject has no rights over whether or how the intervention is administered (participation autonomy). One might counter that even if that is correct, the subject may still have rights regarding whether his responses to the interventions can be used for research purposes (product autonomy). However, one might in turn counter that even these concerns might be discounted if the actions are public information.

In Section 2, I noted that maintaining the trust of subjects is of paramount concern to medical researchers. This provides a basis for insisting on informed consent by subjects. As argued in Section 2, for social scientists, the confidence of the general public and of policy makers in particular are also critical. Moreover the welfare of non-subjects may be of critical importance. These considerations have two implications: first that depending on the treatment of the problem of audience, the form of consent needed may differ from the current standard; second that depending on the population affected, the focus on subjects as the locus of consent may not be appropriate: the informed consent of practitioner partners and affected third parties may be just, or perhaps more, critical.

4.3 Varieties of Consent

Given the multiple desiderata associated with consent we may expect that variations of the informed consent process might succeed in meeting some or other of these.

For example, if what is valued is *participation autonomy*, then this seems to require actual *ex ante* consent. The loss in autonomy consists of the absence of choice to be subjected to a treatment. The demands of *product autonomy*, unlike participation or behaviour autonomy, can be met with *ex post* consent. The demands of the *diagnostic test* can in principle be met by notional consent, and so on.

With this in mind, Table 1, considers how eight approaches to the consent process fare on different desiderata.²¹

1. ***Ex ante informed consent***: *Ex ante informed consent* fares well on autonomy principles as well as on legal protection of researchers (if documented)

²¹ In addition one can imagine various combinations of these strategies: e.g. delegated superset consent or superset package consent.

Table 1: Consent Strategies.

Strategy		1	2	3	4	5	6	7	8
		Informed consent	Implied consent	Proxy (delegated) consent	Superset consent	Package consent	Deferred (ex post) consent	Inferred (surrogate) consent	Spheres of ethics: (compartmentalized consent)
Illustration/description:	Subject provided intelligible information and asked for consent before implementation	Patient holds arm out to doctor to provide injection	[Statistical] Subject asked to appoint someone who they trust to receive information about experiment A and provide consent on their behalf [Authoritative] Perhaps: representative used that is not specifically appointed by subject for this purpose	Subject asked in advance of experiment A if they would be willing to take part in each of experiments A, B, C	Subject asked in advance of experiment A if they would be willing to be assigned to take part in one of a set of experiments	Subject asked in advance of experiment A if they would be willing to be assigned to take part in fact if data may be used	Subject unwittingly takes part and asked after the fact if data may be used	Sample of nonsubjects that are "like" the subject are asked if they would be willing to take part in experiment A. Inferences on hypothetical consent of subject are drawn	A researcher partners with a practitioner to implement a study. The practitioner is responsible for the intervention and the researcher for measurement
Consideration	++	+	++	++	++	+	-	+	-
Application of diagnostic test	++	+	+	++	++	+	-	-	-
Individual participation autonomy	++	+	+	++	++	+	-	-	-
Individual behaviour autonomy	++	+	+	++	++	+	++	-	-
Individual product autonomy	++	+	+	++	++	+	++	++	-
Researcher gains knowledge about likely felt costs of intervention	++	+	+	++	++	+	++	++	-
Researcher gain prior knowledge about particular risks facing individual	++	+	+	++	++	+	-	-	-
Legal protection of researchers (assuming documentation)	++	+	+	++	++	++	-	-	+
Reputation of discipline	?	?	?	++	++	++	-	+	+
Beneficence (towards subjects)	?	?	?	?	?	?	?	?	?
Avoids Hawthorne and related biases	?	?	+	+	+	+	++	++	++
Low cost	?	+	-	-	-	-	?	+	+

Source: Author.

and the reputation of the discipline. As argued above however it is not a necessary or sufficient condition for respect for persons, in addition it may impose costs on subjects, weaken the quality of some kinds of research, and be costly to achieve.

2. **Implied consent:** An alternative is **implied consent** which arises when there are grounds to think that consent is given even if consent is not formally given or elicited. Implied consent might include cases in which voluntary participation is itself considered evidence of consent to be in a study. Implied consent can reduce costs to subjects and researchers but may leave researchers in a legally weaker position and may put their reputation more into question.
3. **Proxy (delegated) consent:** Both *ex ante* consent and implied consent suppose that subjects are informed of the purpose of the experiment *ex ante*. In some settings, this can threaten the validity of the research. An approach to maintain a form of participation autonomy but keep subjects blind to treatment is to ask subjects to delegate someone who will be given full information and determine on their behalf whether to give consent.²² Insofar as the subject sees the delegate as their agent in the matter, proxy consent inherits the benefits of *ex ante* informed consent, but with reduced risks to the research. A weaker alternative – the “**authoritative**” approach – is to seek consent from a proxy that is not specifically delegated for the purpose by a subject. In some settings for example the consent of community leaders is sought for interventions that take place at a community level; this procedure invokes the principles of proxy consent but assumes that individuals that are delegated for one purpose inherit the authority to be delegates for the consent process. Baele (2013) for example recommends this form of consent.
4. **Superset (Blanket) Consent:** Another way to protect research integrity while preserving subject autonomy is to seek what might be called “Superset consent.” Say a researcher identifies set X of possible experiments, including the experiment of interest. The researcher then asks subject to identify set $C \cap X$ of interventions for which the subject is willing to take part.²³ Given this procedure, if set C includes the experiment of interest, a researcher can

²² The *Institutional Review Board Guidebook* argues that “the ‘proxy consent’ of someone other than the subject is not the same as the subject’s own consent, although it may be an acceptable substitute when a subject is unable to give informed consent.” Indeed in general the commission refers to proxy consent as permission so as not to confuse it with consent.

²³ A practical complication arises here in that subjects can make inferences on x given that they learn $x \in C \cap X$; in the extreme case if $x = C \cap X$ the subject knows the nature of the treatment for sure; moreover their final beliefs depend not just on their prior belief about x but their prior

conclude that consent has been given for the experiment of interest even though the experiment has not been specified as being the one of interest. In practice, abstract descriptions may suffice to generate consent for large classes of experiments (for example a subject may consent to any experiment that seeks to answer some question in some class for which there is no more than minimal harm); greater coarsening of this form implies less specific information (see Easton on waived consent).²⁴

5. **Package consent:** An alternative to superset consent is a process in which subjects are asked whether they are willing to take part in an experiment that will involve some intervention in set X , including the intervention of interest. If the subject agrees, then consent for the intervention is assumed. This differs from superset consent insofar as it is possible that a subject would be willing to accept the package but not accept the individual component if offered that component alone. For example if X contained experiment A in which I could expect to win \$1000 and experiment B in which I expect to lose US\$10 I might consent to set X , but only in the hope that I will be assigned to experiment A . To enhance informedness, the subject may be provided with the probabilities associated with the implementation of each possible experiment. Critically, this approach may be inconsistent with a desire to have *continuous* consent – in the sense of consent not just at study outset but in the course of the study also. In a sense under this design a deal is struck between researcher and subject and the subject is expected to follow through on their side of the deal; this limitation runs counter to common practice but is not inconsistent with respect for persons.

beliefs about all the interventions that might be implemented. In principle it is even possible that willingness to provide consent might depend on certainty about the true treatment – for example I might benefit more from an experiment if I know what the experiment is about (for example perhaps I could then game it, or simply steel myself to the emotional turmoil the intervention might produce), in which case the set C itself depends on prior beliefs. In principle these problems could be reduced if proceeding depended on characteristics of C other than $x \in C$, for example that C contain at least, or exactly, n elements.

24 It is possible to generate mechanisms for systematizing superset consent. For example one could create an exchange where prospective subjects provide *ex ante* consent to participate in various classes of experiment; they might elect experiments that address topics they care about, or that are accompanied with varying degrees of risk, and they might rule out participation in certain types of trial also. In addition, prospective subjects could specify the maximum number of experiments that they would like to take part in per year and perhaps a fee for participation. Such a mechanism, though creating obvious risks of weak external validity that need to be assessed, might allow for the implementation of research like that conducted by Milgram and others with participants that are blind to treatment but supportive of the research goals and willing to take on risks.

6. **Deferred (retrospective, *ex post*) consent:** When consent is not sought before the fact, it is common to provide a debriefing after the fact. In some cases this might be important to avoid harm. In the Milgrom experiments debriefing could help remove guilt, if subjects find out that they did not in fact torture the confederates. But beyond debriefing it is possible to seek *consent* after the fact (Fost and Robertson 1980). For some purposes this is too late: it does not restore participation or behaviour autonomy,²⁵ but it does provide product autonomy and it does satisfy the diagnostic test. In some situations, retrospective consent might impose costs on subjects however and generate a sense of lost autonomy.
7. **Inferred (surrogate) consent:**²⁶ Consent is inferred (sometimes, “presumed”) if there are empirical grounds to expect that consent would be given were it elicited. As described above, the diagnostic test for respect for person is not that consent has been obtained but that it would not be refused if sought. This question is partly answerable. A number of different approaches might be used. For example one might describe an experiment to a random subset of subjects and ask them if they would be happy to take part in this experiment, or if they would be happy to take part in this experiment, *even if their consent were not sought*. One could also combine this with *ex post* consent by implementing the experiment with a subset of actors and then ask them if they are happy that they took part, even though they were not told the purpose; or alternatively if, knowing what they know now, would they have been willing to give their consent to take part *ex ante*? Inferences may then be made to the willingness of the larger population to provide consent. This might be called the **statistical** approach.²⁷ Again a weaker, authoritative alternative may be invoked by seeking consent from a third person that does not have legitimacy to speak on behalf of the subject but who is believed to have insight into the subject’s disposition.

²⁵ As noted by Gray (2001): “Deferred consent allows researchers to breach the subject’s autonomy by making the enrolment decision themselves, then asking for retrospective approval. If the patient ... objects to the study, it is possible to withdraw from the study but not to ‘undo’ the exposure to the study intervention. Deferred consent cannot be seen as a substitute for full informed consent.”

²⁶ See Love and Fost (1997).

²⁷ A natural question arises regarding what to do if, as is likely the case, <100% of sampled subjects provide consent. One approach would be not to proceed if there are any objections. Another is to use a rule of thumb. For a related problem of determining how much information to provide for informed consent a 95% satisfaction rate has been proposed. See Veatch (2007) for more on this issue.

The final approach marked in column 8 Table 1, is the spheres of ethics approach, described in Section 3.

Thus although currently researchers use a very narrow operationalization of the principle of consent the broader menu of possibilities is quite large. As of now, researchers could test and develop these in settings in which consent is not routinely sought. Though most of these fall short of fully informed consent, many meet the principles of respect for persons more effectively than consent as sometimes practiced. Looking forward, collective answers to the question of audience and agency can help determine which type of consent is optimal when.

5 Conclusion

I have described the primary problem of assessing the ethical implications of social experiments as a problem of audience. Medical ethics have been developed in large part to regulate relations between medical researchers and patients. Social scientists have adopted the framework created for medical researchers but their audiences are different: at least in the area of experimental research on public policy, relations with policy makers, practitioner organizations, and the general public can be just as important as the relationship with research subjects. Moreover the interests of these different groups often diverge, making the problem of constructing ethics more obviously political.

These considerations suggest two conclusions.

First, rather than seeking some fundamental answer to ethical dilemmas or seeking to address the practical problems facing social scientists using the tools generated for another discipline, there is a need for a social process of *construction* of ethical principles that address the preoccupations of social scientists in this field, especially in settings in which there are power imbalances between lead researchers and research partners and in which there are value disagreements regarding what constitutes beneficent outcomes. Such a process will be inherently political. Just as social scientific interventions are more likely to have distributive implications – generating costs for some and benefits for others – so ethical principles of engagement, if there is to be engagement at all, may require the principled taking of sides, that is, the choice of an audience. The importance of constructing an appropriate ethics for this field is of some urgency since there is no reason to expect that all researchers working in this domain will independently converge on consistent standards for experimental research in grey areas.

Second, depending on answers to the problem of audience, it may turn out that answers to the questions of agency (Section 3) and consent (Section 4) will be

different for social scientists than for medical researchers. I have sketched some possible answers to the questions of agency and consent that diverge somewhat from standard practice. Currently when researchers engage in studies that generate risks, they defend the research on the basis of its social value. But they do so often as interested researchers and without equipment to weigh benefits in the presence of value disagreements. Greater efforts to share the responsibility of research, whether through more carefully crafted relations of agency with developing country actors or more diligent focus on consent may reduce these pressures on value assessments and may also reduce risks to both populations and the professions.

Acknowledgments: Warm thanks to the WIDER research group on Experimental and Non-Experimental Methods in the Study of Government Performance. Earlier version presented at UCSD conference on ethics and experiments in comparative politics. My thanks to Jasper Cooper and Lauren Young for very generous comments on this manuscript. This paper draws on previous work titled “Ethical Challenges of Embedded Experimentation.”

References

- Abram, M. B. and S. M. Wolf (1984) “Public Involvement in Medical Ethics. A Model for Government Action,” *The New England Journal of Medicine*, 310(10):627–632.
- Alderman, Harold, Jishnu Das, and Vijayendra Rao (2013) *Conducting Ethical Economic Research: Complications from the Field*. World Bank Policy Research Working Paper No. 6446.
- Baele, S. J. (2013) “The Ethics of New Development Economics: is the Experimental Approach to Development Economics morally wrong?,” *Journal of Philosophical Economics*, 7(1):2–42.
- Bertrand, M., S. Djankov, R. Hanna, and S. Mullainathan (2007) “Obtaining a Driver’s License in India: An Experimental Approach to Studying Corruption,” *The Quarterly Journal of Economics*, 122(4):1639–1676.
- Binmore, K. G. (1998) *Game Theory and the Social Contract: Just Playing*. Vol. 2. Cambridge: MIT Press.
- Casey, K. R. Glennerster, and E. Miguel (2012) “Reshaping Institutions: Evidence on Aid Impacts Using a Preanalysis Plan,” *The Quarterly Journal of Economics* 127(4):1755–1812.
- Cassileth, B. R., R. V. Zupkis, K. Sutton-Smith, and V. March (1980) “Informed Consent – Why are its Goals Imperfectly Realized?” *The New England Journal of Medicine*, 302(16):896–900.
- DeScioli, P. and R. Kurzban (2013) “A Solution to the Mysteries of Morality,” *Psychological Bulletin*, 139(2):477.
- Fost, N. and J. A. Robertson (1980) “Deferring Consent with Incompetent patients in an Intensive Care Unit,” *IRB*, 2(7):5.

- Gerber, A. and N. Malhotra (2008) "Do Statistical Reporting Standards Affect What Is Published? Publication Bias in Two Leading Political Science Journals," *Quarterly Journal of Political Science*, 3(3):313–326.
- Gray, J. D. (2001) "The Problem of Consent in Emergency Medicine Research," *Canadian Journal of Emergency Medicine*, 3(3):213–218.
- Harms, D. (1978) "The Belmont Report: Ethical Principles and Guidelines for the Protection of Human Subjects of Research," DHEW Publication No. (OS) 78-0012.
- Jonsen, A. R. (1983) "A Concord in Medical Ethics," *Annals of Internal Medicine*, 99(2):261–264.
- Kant, I. (1956). *Critique of Practical Reason*, translated by Lewis White Beck. Indianapolis, Ind.: Bobbs-Merrill.
- Levine, R. J. (1979). "Clarifying the Concepts of Research Ethics," *Hastings Center Report*, 9(3):21–26.
- Levine, R. J. (1988). *Ethics and Regulation of Clinical Research*. Yale University Press.
- Levine, F. J. and P. R. Skedsvold (2008). "Where the Rubber Meets the Road: Aligning IRBs and Research Practice," *PS: Political Science and Politics*, 41(3):501–505.
- Lipscomb A. and A.E. Bergh, eds. (1903) *The Writings of Thomas Jefferson*. Washington, DC: Thomas Jefferson Memorial Association of the United States, 1903-04. 20 vols.
- Love, R. R. and N. C. Fost (1997) "Ethical and Regulatory Challenges in a Randomized Control Trial of Adjuvant Treatment for Breast Cancer in Vietnam," *Journal of Investigative Medicine*, 45:423–431.
- Pallikkathayil, J. (2010) "Deriving Morality from Politics: Rethinking the Formula of Humanity," *Ethics*, 121(1):116–147.
- Tolleson-Rinehart, S. (2008) "A Collision of Noble Goals: Protecting Human Subjects, Improving Health Care, and a Research Agenda for Political Science," *PS: Political Science and Politics*, 41(3):507–511.
- Veatch, R. (2007) "Implied, Presumed and Waive Consent: the Relative Moral Wrongs of Under and Over-informing," *The American Journal of Bioethics*, 7(12):39–41.
- Vollmann, J. and R. Winau (1996) "Informed Consent in Human Experimentation Before the Nuremberg Code." *British Medical Journal*, 313(7070):1445.
- Wantchekon, L. (2003) "Clientelism and Voting Behavior: Evidence from a Field Experiment in Benin," *World Politics*, 55:399–422.