

ETHICS AND EXPERIMENTS

Problems and Solutions for Social
Scientists and Policy Professionals

Edited by Scott Desposato

 **Routledge**
Taylor & Francis Group
NEW YORK AND LONDON

5

PRISON STATES AND GAMES OF CHICKEN

Jesse Driscoll

You are going to get a 23-year-old Azeri kid mutilated, and afterwards I don't know how you are going to live with yourself.

David Laitin, to the author, August 2008

Authoritarian governance practices represent an important research frontier for comparative politics graduate students. Certain regimes rule through a mix of surveillance, fear, and violence. Many researchers in the subdiscipline agree that there are normative reasons to know more about how these practices function. Experiments are one type of tool in a researcher's arsenal to understand the kinds of behaviors that political institutions incentivize. Scholars will continue to think hard about how to safely and responsibly conduct field experiments in challenging environments because they promise to provide our research community with greater traction on causal impacts.

This chapter is primarily addressed to the next generation of graduate students who are contemplating fieldwork in "hard authoritarian" regimes, and may be thinking about running experiments. The thrust of my argument can be easily summarized: You are basically on your own. If you choose to spend your time in graduate school living in dangerous places—and I believe that there are good reasons to do this—one of the consequences is that you will probably, with time, come to disdain the authority of the bureaucratic entities at your home institution tasked with helping you weigh risks. You will, with time, become *the* area expert on what life is really like in "your" particular poorly governed corner of the planet. No one—not even your dissertation chair or your other advisors—will be better positioned than you to evaluate the risks of whatever experimental interventions you are proposing. With that in mind, this chapter is organized

around a few questions you should ask, with a bit of distance, before you commit yourself. They are not meant to be “gotcha” questions. They do not have easy answers. You should take some time to argue with yourself about them. To convey a sense of how conflicted I am on these questions, and how ambivalent I am about the possibility of a generalized rule-set or regime to incentivize good behavior, this chapter is written in the first person.

A secondary audience for this chapter is the community of mentors (and regulators) tasked with restraining the ambitions of the first audience. Most experimental research interventions have historically sold themselves as positive-sum enterprises: Everyone wins, but a control group may win less. However, at the current moment it is reasonable to think out loud—not just in anonymous referee reports, but also in public meetings—whether this analogy actually holds when our research questions concern raw power in the form of zero-sum distributional politics backed by guns. There may be different stakes associated with accidents and miscalculations, and not just for study participants, requiring different ethical frames. This chapter draws attention to two different kinds of analogies that are sure to invoke moral discomfort for most audiences: *prison states* and *games of chicken*. Research designs that involve cooperation with state authorities in order to build state capacity can be compared to constructing social science-enabled police states, or “prison states” (the catchier moniker). At the other extreme, research designs that eschew cooperation with the state and instead work with NGOs to encourage low-intensity confrontation with state authorities are open to the charge of irresponsibly encouraging “games of chicken” against undemocratic regimes, with any serious costs borne by locals. In their most extreme forms, these paired criticisms suggest that some states are so badly governed that conducting any social scientific research at all should be discouraged. By giving voice to these considerations, I hope to inoculate the discipline against extremist articulations of persuasive reactionary arguments.

5.1 The View From 2015: Confessions From the Experimental Frontier

I am neither an ethicist nor a particularly rabid experimentalist. I spent most of my time in graduate school working on questions related to violence and civil war settlement. In practice, this meant spending a great deal of time living in the former Soviet Union (namely, Georgia, Tajikistan, and Kyrgyzstan), and more recently working closely with the Somali-speaking population of San Diego. The research design of my first book project involved extensive ethnographic research and long-form interviews with former combatants in Georgia and Tajikistan.¹ In 2012, I was the principal investigator for the first representative survey of the city of Mogadishu in 25 years.² These observational research designs required calculated risks with my own safety, the safety of my research team, and my human

subjects. I weighed risks as best I could. I sought advice from trusted friends and mentors and the Institutional Review Board (IRB). I ultimately opted to not publish certain data and to abandon certain research projects when I deemed things were getting too dangerous. I do not think anybody has been hurt as a result of my research, and I am glad for that.

The subdiscipline of comparative politics is a competitive and psychologically trying subfield of political science. It has historically valued the labor of young scholars who learn difficult languages, travel to uncomfortable places, live for long periods of time far away from loved ones, and eventually bring back data from under-studied parts of the planet. We “walk the walk,” serving as living reminders that if one really wants to know more about the world outside the ivory tower, at some point it becomes necessary to shoulder the burden of going there. Fieldwork is often lonely, as researchers cultivate the self-reliance necessary to engage for months (often years) in distant (often hostile) political environments. Gratification for empirical data collected is often delayed years or decades. But more than the other subfields of political science, comparative politics dangles the promise of getting inside other cultures. We respect and reward efforts to creatively break down the subject-object distinction across language barriers. We tell our students to leave the comfort of the academy, get their hands dirty, and see for themselves how theories interact with the messy details of the real world. The rise of the experimental paradigm suggests that they will be increasingly rewarded for attempts to tinker with the world they find.

Search committees and tenure committees tend to place great value on demonstrations of scholastic aptitude. Scholars who take questions of identification seriously and demonstrate a “go-get-it-done-then-get-it-published” attitude are, as an empirical matter, more likely to rise successfully through the ranks than academic laborers who become enmeshed in moral quagmires and delay publication of work. I am sorry to report that the implication of this may actually be a prisoner’s dilemma for professionally vulnerable untenured researchers. Even controlling for subfield, some researchers are going to be more risk-acceptant and others more risk-averse. And so long as the discipline continues to reward entrepreneurialism and creativity, it may simply be the case that the risk-acceptant young researchers have a competitive advantage in the marketplace of ideas. And perhaps—just perhaps—that is as it should be.

The primary institutional entity tasked with serving as a check on the ambitions of graduate students is the Institutional Review Board (IRB). The IRB derives its moral authority from the claim that it protects the interests of research subjects. The IRB has the authority to evaluate researchers’ ethical intuitions, despite the fact that the researcher will almost always have a comparative advantage in the “ground truth” of the research site.³ I believe that everyone involved in this process has the best intentions. But it is a mistake to pretend that there is interest convergence or incentive compatibility where none actually exists.

Professionally vulnerable young scholars are explicitly told that they will make no progress up the tenure ladder if they do not produce work that lands in top outlets. This makes for a host of “motivated misperceptions” when it comes to assessing risks. If cowboy behaviors are modeled by professionally successful faculty, the lessons will leak to graduate students. The IRBs tasked with limiting university liability, by contrast, are inherently conservative. This can make for an adversarial relationship. Little about checking in with an IRB that only convenes once per month is incentive-compatible for scholars attempting to improvise solutions in rural areas of a developing country.

I currently work at a policy school that credentials terminal MA students. Many are interested in working in the field of international development, and their intuitions seem to be that so long as foreign aid flows from the core to the periphery—in the name of security or charity—the justifications for those aid flows will be articulated in a results-oriented language. Very few of our students matriculate to PhD programs, but our methods sequence—which is state of the art—leaves very little doubt in the minds of our students about what they ought (normatively) to be doing: They should be randomizing what interventions they can, when they can, as best they can. Our best students take these skills and ethical intuitions to very well-funded actors that *do not have IRBs*—investment banks, multinational corporations, advertising conglomerates, Internet start-ups, democracy promotion NGOs, the US military, and the like. The world outside of the ivory tower moves pretty fast. A lot of people in the future are going to get into the business of convincing constituencies in embassies, militaries, and the NGO community to randomize programming. The unexpected results of gradual experimentation will almost certainly teach us all a great deal about the world.

So in summary: I really do believe that there are normative reasons to keep experimenting, and no obvious way to put brakes on the train. Given this constraint, the following two sections are meant to draw attention to ways that the experimental ethos might interact with authoritarian practices in a way that that pose ethical challenges for researchers.

5.2 Four Questions

I think that a decentralized regulatory regime—one in which everyone asks themselves difficult questions as they course-correct in the field—is superior to the feasible alternatives. It is not clear that a centralized regime to gauge the ethics of research is either feasible or desirable. When one multiplies the number of research institutions by the myriad possible combinations of field sites, dependent variables, and experimental methodologies, and complicates the picture further by acknowledging the vast diversity of leverage-able personal networks and potential second- and third-tier harms, we are suddenly in a space that is so vast that it hardly seems useful to generalize about best practices. Abkhazia is not Somaliland.

The best practices of doing research on elections in Afghanistan may not travel even as far as Tajikistan, let alone Kenya or Russia. India's military is not Mexico's municipal police force. Internet monitoring works differently in Turkmenistan than in China. Democracy promotion NGOs work very differently in Burma than their counterparts do in Bahrain. The basic norm governing the scope and scale of interventions at the research frontier certainly seems to be: “Let the area specialists, who can gauge risks best, figure out what they think they can get away with and try to get published afterwards.” The magnification of the possible harms emerging from field experiments would not serve the interests of the discipline. Most experimental interventions tend to be very small, and cannot possibly do much harm. Large interventions are almost always randomizations of practices that happen all the time, or are representative of the kinds of things that were probably going to happen anyway. No one has any idea how to assess the probability of a black swan-style “nuclear” failure.⁴ The people who are potentially threatened in a “black swan” scenario, where everything goes as wrong as it could possibly go, could easily be people who are not subjects.⁵

For graduate students reading this, let me make the point as explicitly as possible: It is *your* job to assess these risks and stop the project if these risks are serious—not the job of the IRB. The IRB framework is not always going to be able to help you sort through these questions (and, as a community, we should not try to force them to do any more work than they already are). Conducting scholarly work in authoritarian environments on politically sensitive topics while staying safe requires keeping one's eyes open and responding flexibly to highly local and contextualized variables. To their credit, many individuals who work at IRBs will readily admit that they do not have a clue about those kinds of details. So: You are basically on your own. Be really honest with yourself. You have a long memory.

Will Your Findings Legitimize Authoritarian Behaviors or Perpetuate Bad Government?

I have spent much of the last five years working on a series of projects that assess welfare outcomes in Somalia. The stakes of this research project are high, and we have found many willing allies. But these allies usually wanted something from us: They wanted evidence that what the Somali government was doing was working. And when they noticed that the evidence did not seem to support this theory, but that we were continuing the research anyway, their partnership became more and more costly.⁶ I mention this particular story only to draw attention to a familiar fact: Doing work in authoritarian or badly governed societies requires taking advantage of the idealism of people who see themselves engaged in public goods provision or charity. And when we “sell” projects to these local partners, we are often tempted to pretend as if we share their assessment of an ideal outcome. (Otherwise why would we be there?)

If we are truly honest with ourselves, one of the of the important reasons that field experiments have gotten so popular is born from a logic of pure pragmatism: They are a technology that facilitates a positive-sum exchange relationship between academics (who desire original datasets) and practitioners (who want to be able to report to constituencies that they are *positive* that the programs are working). This exchange relationship can be easily overlooked if one focuses strictly on the philosophy of science arguments supporting the experimental tradition in social inquiry.⁷ This pragmatic rationale also has a mercenary variant. One begins by simply noticing, as a matter of fact, that the World Bank is probably not going to stop making loans or discontinue commissioning papers. Development economists arrive to this ongoing conversation armed with the confidence that their methods produce the only kinds of answers that count. What this means, in practice, is that even if influential political science departments coordinated to put a halt to experiments because of shared ethical considerations, *the experiments would not stop and might not even slow down*. And the truth is that clever identification strategies are a more valued currency in the discipline today than in the past. Articles in the *American Political Science Review* and the *American Economic Review* do not look as different as they did two decades ago. There are real, measurable rents to being the kind of political scientist who can pinch-hit as an economist. I could be wrong, but I doubt this observation will have become obsolete any time soon.

There are many outcomes that are interesting to social scientists but that cannot be experimentally studied without cooperation from state organs (e.g., taxation; economic redistribution; the orderly production of justice; efficient counterinsurgency; the collection of social intelligence, using technologies like the census; and the functioning of the education, prison, or pension systems). In a well-functioning society, it is difficult to imagine even an observational study of most institutions without state permission. Sometimes states are principals, scholars are agents.

And none of this is a problem, really, until we notice that many states function quite a bit like prisons. One can imagine a spectrum. On one end you would find places where citizens basically elect their government and slowly shape its institutions. On the other end you would find places where the governing authority is analogous in most ways to a prison warden.⁸ It is important to notice that at both ends of the spectrum, the state not only experiments but also observes—often archiving vast quantities of citizen data without their permission. The rapid proliferation of computer and smartphone technology is expanding the state's ability to do these things. As social scientists, many of us badly want these data. Such desire will lead some of us into talking ourselves into helping states—even the bad ones—to collect and analyze more data. Now: If a student hopes to someday be a principal investigator working on political violence, part of her graduate school training ought to include a serious probing of her own threshold for dealing with the agents of state security bureaucracies. My threshold is high. I invested

many hours lobbying to randomize the placement of police cameras in the city of Newark. I spent two weeks with no security living in the unrecognized state of Abkhazia, attempting to secure political permissions from elites in the unrecognized government to conduct the first representative household survey since the Soviet era. Both projects would have involved substantial cooperation with “the state.” Both would have involved populations that could not easily opt out of the study.⁹ We have ethical intuitions that study populations that cannot easily “opt out” of the study—especially prisoners—are somehow deserving of different protections than regular subject pools. It is difficult for a researcher to extract herself from the background-level of coercion in a prison by blithely invoking the language of informed consent, once the ghosts of the Zimbardo experimental subjects have been summoned forth.¹⁰

As social scientists, there are reasons to be wary of the top-down, “eyes of the sovereign” perspective. The entire conversation about randomizing scarce public goods is implicated by this line of thought. It is vulgar, somehow, for the foreign observer to assume a “right to treat.” But when matters of life and death are at stake, and the state is badly governed, it is not always clear that there is a local domestic moral authority to legitimize the enterprise. In some cases it is possible to argue that any findings that could serve to legitimize certain very bad states, or make these very bad states more efficient and effective at controlling their populations (“in the name of order, in the name of development, in the name of making the trains run on time . . .”) is morally compromising. When the PI doubts her own ability to publish or publicize non- or negative findings on the intervention, political compromises are being made. Every researcher is responsible for finding her own threshold of comfort for self-censorship, of course. But we would do well to notice that the hard-won ethical intuitions that IRB professionals have developed about the experimental paradigm are calibrated to the harms that can arise for subjects that reside on college campuses.

I do not have strong ethical intuitions on whether there are some states that are just *so badly governed* that we should not study them at all. The truth is there are many states with ruthless security services, whose leadership is desperate for foreign aid and hungry to validate certain outcomes for impatient donors. What “informed consent” is supposed to mean when the study is backed by a state of this kind is by no means clear.¹¹ In the third round of telephone call-backs to respondents residing in Somalia—individuals who had been read lengthy Somali-language consent scripts multiple times—at the end of a survey we asked respondents who they thought we were. Some of the responses were very funny and self-aware (“What do you think I am, stupid? You’re researchers from a university in California! You made me sit through a three-minute script explaining that to me!”). Some were equally self-aware, but not at all funny (e.g., “You’re working for the US military—you just don’t realize it.” “You’re going to probably sell this to a NGO who is trying to decide whether or not to send aid.”) An intuition

that the researcher is in league with the prison warden is likely a source of bias in much of the data generated by surveys and experiments in certain places.

How Will You Publish If There's No Effect—Or, Worse, If Your Results Are "Signed Wrong"?

If the experiment is well designed and implemented with integrity, non-findings ought to be publishable. Our science is young and our theories have low predictive power. But the reality is that our discipline has a poor track record with the publication of non-findings. And today there are a lot more scholars planning RCTs than there are slots in top field journals. This means well-identified papers that are not *quite* meshing with where the gears of interest are at the time of submission will be rejected at top outlets and ultimately land in journals with low impact factors. Sometimes this is going to occur because of completely insincere ethical hand-wringing, and sometimes because of completely unjustifiable parochial interests ("... But what would it *mean* if our interventions were having *no effect at all*?"). The temptation to weaponize the language of ethics in the service of old-fashioned interest group politics is very real.¹²

The only field experiment I have personally overseen was conducted in Georgia in 2008, in collaboration with another then-graduate student (Daniel Hidalgo, now an Assistant Professor at MIT). In brief, it was an information dissemination RCT, intended to lower the cost of election-day malfeasance reporting. The randomization protocol worked exactly as we expected it would, and the subjects behaved in ways that were consistent with our theoretical priors. If we had stuck to our pre-committed data analysis plan, we would have written up an optimistic narrative of citizen empowerment. However, the experimental design forced us to confront an unexpected downstream effect of the intervention: For every person contacted, two people stayed home on election day. When we "went fishing" for heterogeneous treatment effects, the most plausible narrative to explain the mechanism was fear. Many Georgian citizens, especially in rural areas where the political machine was contested, were alienated by the study. We speculate that they found the increased attention to their district unnatural.¹³

For myself as a researcher, it the experiment represented an important turning point in my relationship with Georgia as a field site. I had imagined myself to be working in solidarity with NGOs against a state apparatus—but many of the subjects clearly assumed that I was working for (or with) the state. And I don't really think that the problem was a lack of local knowledge: I knew enough about shady Georgian electoral practices to design an experiment that successfully disrupted those practices. But the motives for doing this were suspect for a large percentage of the study population. Ultimately the study found its way into print.¹⁴ But noticing that our experimental interventions were interacting with authoritarian legacies really did open a can of worms. It became clear that our donors had no

idea what they were supposed to do with these findings, and that there actually is no consensus in our discipline on what should happen if experimental outcomes turn out to be normatively bad and/or genuinely unexpected.

I am not sure anyone hearing this story thinks that we ought to have stuck to our data analysis plan and wound a self-congratulatory yarn about how we increased citizen activism, just burying the unexpected suppression finding. But I am still not 100% sure how we should have reported the suppression finding. This really does bother me sometimes.

Is Your Experiment Mostly an Excuse to Engage in Confrontational Activism?

Thomas Schelling famously compared bargaining in the shadow of violence to a game of chicken, where two cars speed towards each other to see who swerves first. Most of the time at least one player swerves. But sometimes players both miscalculate the resolve of their opponent and a tragic collision occurs. In game theory, it demonstrates the idea that there are circumstances in which each player prefers not to yield to the other, but both structure their strategies to avoid the worst possible outcome that occurs when neither yields. So far, when social scientists have worked in solidarity with local activists—on election monitoring and on "get out the vote" campaigns to help opposition candidates—it has been state security entities that have swerved. We should not expect this to continue indefinitely.

My early graduate school experiences with the IRB emphasized that scholarly work can be akin to a diplomatic passport. As members of a scholarly community, we get visas, we get a number of invisible social protections, and we are expected to methodically conduct the research we set out to conduct. While this model works well for historians working in archives, it does not work nearly as well for political scientists working on contemporary party politics, contemporary voting behavior, or contemporary counterinsurgency. I am afraid that it works less well the longer a political scientist stays in the field, makes friends, and begins to see herself in solidarity with her subjects. If a scholar is inclined to fly into a country, stay for two to three weeks, conduct a few elite interviews, eat at nice restaurants, and then fly home, the diplomatic immunity model works fine. But if the researcher engages in sustained fieldwork over many years—learning local languages, unraveling the ironies of local political representation, allowing herself to cultivate empathy with human subjects and "go native"—at some point it occurs to the researcher that she might be able to actually effect social change. Once they feel that they *can* cause change, some of our students will decide that they *should* cause change. Many of us are interested in processes of social mobilization, party formation, and opposition voting, and continue to struggle with the question of whether we ought to do more than just watch on the sidelines.

Many professional political scientists self-selected into the discipline because of an interest in understanding how political structures change. Political science attracts students interested in sustained thinking about the constraints on political change, but also in testing theories about how systems can adapt to perturbations.¹⁵ Particularly in graduate school, when the time horizons stretch out towards infinity, many choose dependent variables that can be squared with the kinds of changes they would like to see in the world. Advisors often encourage this attitude, up to a point.

But the impulse to tinker can be dangerous. A colleague once proposed putting up pro-Aliyev and anti-Aliyev stickers in randomly selected neighborhoods of Baku (the capital of Azerbaijan) and measuring how long it would take for them to get taken down. We hoped the research project would reveal something about expectations of regime autosurveillance without drawing attention to shady election practices (which we both agreed would be too dangerous to study). Over several weeks, I slowly operationalized the idea. Having sketched a formal research proposal, I finally shared the idea with my advisor, David Laitin. After careful consideration and 40 minutes of back-and-forth to reveal that I had done my due diligence on issues of identification and research design, he finally paused, looked me in the eye, and delivered the quote that provides the epigraph to this chapter. I stopped immediately and I am glad that I did. But the sequencing of events is telling and representative of the way that scholar-NGO collaborations actually unfold: First I came up with something that was do-able and interesting, then I thought about relevant literatures, next I began to plan the logistics of getting local permissions and implementation allies, then I consulted with a trusted advisor, and, had I proceeded with the project, only at the end would I have begun the process of reverse-engineering a self-righteous moral justification for IRB paperwork. I doubt that I am the only one for whom this is true.

In the best-case scenario, the rise of RCTs and the experimental ethos can serve as a tool for researchers to get a lot of early advice, interface early and often with advisors and home IRB institutions, learn the relevant local laws, acquire local allies across the spectrum of relevant civil society actors about best practices, and pre-commit to data analysis plans. But, in the context of games of chicken, all of this early work *must also be understood as a mechanism for locking-in one's commitment*. With a large organization mobilized, tenure clock pressures, and a donor waiting for results, the PI could easily come to see herself as fully committed to a confrontation with the authorities. In the worst-case scenario, certain of us will be complicit in blurring the line between anti-regime activism and the march of science. In the event of tragedy, I suspect it would take hours, not days, for certain conservative voices to opportunistically rush to claim in hindsight that "anyone could have predicted" the risks. The language of "playing God" will be invoked. "Nazi doctor" analogies are sure to follow. It would be bad for the entire discipline.

There are two kinds of disadvantages to these sorts of confrontational projects. The first have already been alluded to: As a community, we may be overdue for a high-profile incident in which we will not be able to falsify the hypothesis that it was the researchers' intervention that *caused* someone to be injured. Just off the orderly equilibrium path that we observe in authoritarian regimes, there is more violence than most people can easily conceive. Large-scale field experiments in badly governed states are already interacting with local political equilibria in ways that are well beyond our capacity to predict. My hunch is that if a cascade of unexpected events ends in violence, it will be due to the entrepreneurial labor of a very bright, ambitious, and idealistic young scholar: Someone vulnerable and hungry, trying to get out ahead of the curve, fighting the tenure clock, who has only selectively internalized the advice that has been heaped upon him or her, but who *earnestly believes* that they are working on behalf of a brutalized population that needs helping, that their position is functionally unique, that their labor is changing the world for the better.

But a second disadvantage is in many ways just as serious. Affiliation with activists risks gradually eroding our political neutrality, even without a "black swan" crisis.¹⁶ The observer status that undergirds the diplomatic passport analogy (above) is worth preserving. My university business cards did not give me diplomatic immunity or unlimited access, but they weren't exactly cheap talk, either. When I was understood to be a neutral, scholarly observer, I think I received more access and better data. Affiliations with activists or NGOs who are "against the state" can have liabilities. A simple one is the possibility that the side you are on may change—either because the political environment changes or because your ideological predilections change. You may be compromised without your assent. (Field notes can be confiscated. Emails can be read without your knowledge or permission.) Working on politically sensitive topics invites scrutiny about your true motives—and that scrutiny can have real, tangible, negative effects on the lives of your interview subjects. I eventually stopped working in Tajikistan completely because I could not, in good conscience, answer these questions to my own satisfaction. When we can credibly present ourselves as scholars, and not activists or spies, we are at our safest—and so are our human subjects. It is very important to remember that.

Are You Really Planning to "Go Native," or Are You Collecting Exotic Passport Stamps?

Imagine a spectrum. At one end, you have non-American graduate students who have gravitated to the subdiscipline of comparative politics in order to write credibly about their home societies. They do not need to "go native"—they *are* native, and cannot jettison this status. These scholars enjoy substantial comparative advantages over their American counterparts in terms of

fieldwork start-up costs: They already speak the language(s), know the history, have a “feel” for how to get things done, and may already have a network of contacts that will help them collect data. At the other end are pure technologists. These academics would never be confused with having any interest at all in going native. They may have no real knowledge of the language spoken by the experimental subjects or even be able to produce a credible map of the country. But they *do* know exactly what the paper needs to look like in order to appease reviewers, and the value of replicating experiments in different settings.

Where do you lie on this spectrum, vis-à-vis your current field site? Do you really see yourself as someone who will be back here in 5, 10, or 20 years? Or are you, in your private moments, very very uncomfortable and counting down the minutes until you can go home? If you don't exactly know, it is good to be able to admit that. Your answer will probably change with time. But it is important to appreciate that the kinds of experiments that you will consider worthwhile—worth writing grants and recruiting labor for—are different depending on where you are on this spectrum. Performances of “going native” are an unattractive result that may be inevitable in your mid-20s. As a rule: You should never forget who you actually are, because your study population certainly won't. If you are really planning on going native, it may make more sense to self-censor certain opinions so that you can go-along-and-get-along over the long haul. By contrast, it may *only* make sense to engage in confrontational activist politics in a society that you have embraced as your own. Experimentally demonstrating uncomfortable social facts is, after all, an important facet of how political science works. But you should notice, for your own ability to sleep at night, that as one gets closer to opposite ends of the spectrum *the same RCT* can be described as urgently necessary or a stupid stunt; *the same experimental study* can be described as normal science or a gratuitous waste of valuable energy.

In my experience, it is easier to talk people into collaborative research if you can credibly present yourself as someone with a real stake in getting the local story right. But once you secure local collaboration, it is very tempting to simply “parachute in” for the minimum amount of time necessary to get your name on the paper. Different approaches work for different people, but I will warn you that locals, at least in my experience, become very cautious—and very nationalistic on behalf of their co-nationals in the “control group”—when they feel that they are being experimented on. One of my graduate students, who had spent years living in Kyrgyzstan, was rebuffed in her efforts to organize a randomized roll-out of Internet service to rural areas. I have no doubt that she was capable of articulating (in fluent Russian) the canonical justifications for a lottery: a scarce public good, more transparent mechanism of selection, and all the rest. But the phrase that she kept hearing was *ne-naturalnaya*: not natural. I am speculating, but my guess is that that her Kyrgyz interlocutors understood why it would be good for this PhD's

career if she could manipulate the provision of a public good to real people, but still thought the whole arrangement being proposed was . . . *nenaturalnaya*. As Jarvis Cocker observed: “Everybody hates a tourist, especially one who thinks it's all such a laugh.”

5.3 Conclusion

In a decentralized self-regulating environment, academic advisors and senior scholars will have to shoulder the responsibility of policing this new frontier. Indeed, this is just an acknowledgement of a burden that they already bear. Projects are already being judged by the ethics of the research design. What is missing now is for academics to push their students to articulate an ethically defensible frame for their ongoing research. “I had IRB for this” is not going to be a sufficient answer at all. In contrast, students ought to be expected to articulate a positive case for their interventions into the lives of their subjects. (And note that this is not exactly the same thing as saying what most people currently say, which is “Well, it was going to happen anyway, so there was no additional harm to randomizing.”) Senior scholars have a comparative advantage in the production of credible “area studies” wisdom—an intuitive sense for what makes some studies worth the risk, and others less so. But ultimately the arguments must be voiced by the graduate student as she learns what it means to be a PI. Senior scholars can play the guiding role, perhaps with some gentle assistance from junior scholars and allies in the IRB.

If trends persist, there is a non-zero probability that an experimental social science intervention will alter facts on the ground enough to literally change who gets shot and who does the shooting. If this happens, the discipline will confront the big question that this essay has danced around: Are there certain kinds of inputs that we simply should not randomize? If the applications of a certain kind of violence are inevitable, but the efficacy is in question, is there not a case for randomization? How about using social media, or collaboration with state security services, to apply different kinds of fear-inducing treatments to populations? It is fine to laugh all of this off as something akin to a movie plot, but if randomistas are not wrong on the philosophy of science, isn't limiting the size or kind of intervention arbitrary? On the one hand, I suspect that many in our community would balk at endorsing the kinds of large-scale social engineering projects that Soviet social scientists engaged in.¹⁷ On the other hand, when pressed, I am forced to admit that I would like to know many things that we do not know but are forced to act as if we do in order to make policy at present, such as whether or not UN Peacekeepers need to be armed.¹⁸ Policy entrepreneurs would rush to assert they know the answer already, but they may be mistaken in their belief that their observational data can provide insights that are actually analogous to having run a proper experiment with a real control group.

Some of us may decide that research is too dangerous to sign off on, or too damaging to the interests of the study population (particularly for populations that cannot easily opt out of the study). If we articulate these concerns loudly and publicly, agents of risk-averse universities may find themselves obligated to stop the science. But in general there are not many mechanisms to hold risk-acceptant colleagues to risk-averse standards, even in the same department at the same institution. In the short term, I do not see any reason that referee reports, combined with old-fashioned naming and shaming at academic meetings, cannot steer our ship. Disagreements should be aired. In this give-and-take, graduate students will have the chance to hear different points of view and then vote with their feet, “forum shopping” for advisors who are optimally paired with the students’ risk profile. In the worst-case scenario, our discipline may balkanize into do-ers and scolders. I have no doubt that the discipline will survive.

I firmly believe that inquiry into the ugly facets of politics, such as political violence, civil war, intergenerational poverty, and racial discrimination, is normatively defensible. Most of us believe that the rigorous application of the scientific method to social diseases provides a host of direct and indirect goods to liberal society. All of the potential disadvantages sketched above could easily be outweighed by the social benefits associated with pushing the frontiers of social science research. And all of the potential disadvantages sketched above could easily be outweighed by the creation of transnational civil society linkages, enabled by technology and a great deal of volunteerism from within our discipline. Many of us will continue to do experimental science while imagining that in the future citizens in badly governed states will feel a bit more protected, and governments a bit less likely to act with impunity, because of our labor. It is a reassuring sentiment, even if the evidence is mixed.

Notes

- 1 Driscoll (2015).
- 2 Driscoll and Lidow (2014).
- 3 IRB practices have evolved in a way that are designed to shield research institutions from legal liability, and tend to be staffed with people who want to help researchers think through potentially negative downstream consequences of their projects. The scholar must submit to a set of lengthy trainings, occasionally change the research design, and in general provide evidence that the PI has thought about a standardized battery of questions. As a quid-pro-quo, it is the institution that will be liable if anything goes wrong, not the researcher.
- 4 “Black swan” events are a catch-all metaphor for a high-profile cascade of unforeseeable events. Taleb (2007) developed the theory of black swan events, suggesting they are characterized by rarity, extreme “impact,” and retrospective (though not prospective) predictability (xvii–xviii). The thrust of the work, in my reading, is summarized well on page 77: “We worry about the wrong ‘improbable’ events.” In this essay, a black swan nuclear event is a high-profile cascade of terrible events that could be traced to a social science PI.

- 5 They could be the experiment’s implementers, who end up jailed for high treason. They could be people who are just at the wrong place at the wrong time during the violent outbursts that take place in the wake of a contested election.
- 6 I was once asked point-blank by the head of the Somali Youth League of San Diego: “Are we putting together a predator drone list?” I explained that we were not—we were coding the identities of public figures in Somalia. I do not know if he really believed me.
- 7 And for the record: These arguments are very persuasive and self-replicating. No matter how many times critics repeat the assertion that they are a passing “fad,” I doubt randomized control trials (RCTs) are actually a passing phenomenon the way critics wish that they were.
- 8 The correlation with the off-the-shelf democracy and governance indices would probably be high.
- 9 Neither effort resulted in a study, but if they had I am certain that I would have talked myself into doing the research. I think it is also important to note for posterity that in the preliminary “scouting” stage of the research no one—no one at Yale, at Stanford, at Harvard, at UCSD, no one, not once—suggested that randomizing the placement of police cameras in Newark was anything other than a NATURE paper, or that the Abkhazia work would be anything but historic.
- 10 It is unlikely that any individual social scientist is having a uniquely negative impact on the study population, but it is also not clear exactly where it is appropriate to construct the rhetorical “analogous to a prison population” boundary. Are citizens living in certain authoritarian police states (e.g., Pyongyang) analogous, in some way, to a prison population? How about native populations living on reservations? How about populations of refugee camps? How about populations living under military occupation by a country that is known to monitor cell phone and email communications?
- 11 This is only to say that there are certain frontiers of comparative politics where I anticipate that it will be very difficult for researchers to live up to the best practices articulated in Green and Gerber (2002, 829): to “aspire to leave no footprints . . . like visitors to nature preserves.”
- 12 “Of course I’m ethical. So are my friends. You want to be my friend, don’t you?” Cumulatively, this is probably going to lead to a lot of frustration in the next generation of scholars, many of whom are investing a great deal of energy, and doing everything they are being told they are supposed to do, but still aren’t going to get the recognition they feel they deserve. I expect crocodile tears as these arguments filter through the sieve of department, subdisciplinary, and disciplinary politics. The tone and tenor of anonymous referee reports probably matter as much as anything in determining the future of field experimental research in comparative politics.
- 13 Although I felt genuinely guilty about this for a while, I managed to remain focused on the fact that if we had not conducted the work experimentally we would never have even *known* about the troubling downstream finding. The guilt finally passed when it sunk in that doing a pre-election voter attitude survey—e.g., just showing up to ask questions in randomly sampled Georgian villages—had a voter suppression effect that was analogous to distributing information. No one thinks that we should stop doing surveys in Georgia and allow elite pontification from the capital to replace systematic data on public opinion.
- 14 Driscoll and Hidalgo (2014).
- 15 Many members of our community put our training to work to do things other than just publish papers. Many political scientists consult. Many more volunteer their labor to political parties, write op-eds, blog, or “do politics” in a way that Max Weber would have immediately recognized.

- 16 All political scientists are guilty, to lesser or greater degrees, of being disingenuous about our actual political neutrality. We simply present ourselves and our research differently at different times. As I have struggled with the ethical dilemmas discussed in this chapter (and others not discussed), it has been extremely comforting to know that my home institution has "had my back." I am increasingly appreciative of how shared cultural understandings of what professors *are* shielded me while in the field. Some of my subjects understood this better than I did at the time.
- 17 Recall that the Soviet Union was administered by a class of fully self-aware social scientists, confident that there was no normatively defensible alternative to testing their theories. They were very interested in development, order, and social transformation, just as we are.
- 18 I am extremely grateful to Don Green for articulating this point in this way. The short discussion in Green and Gerber (2002, 828–831) is highly valuable.