Working the Unworkable: A View of the IRB Process from the Inside

Kenneth R. Mayer

A good way to start – or perhaps stop – a conversation at any political science convention is to say "IRB." The likely reaction will be a recitation of horror stories of IRB imperiousness, nitpicking, delays, and a complete failure to recognize the burdens imposed on researchers. Most people will conclude that process is excessively formalized, does little to protect the welfare of research subjects, and is unable to prevent research controversies such as the recently retracted *Science* article on attitudes toward same sex marriage (LaCour and Green 2014; retracted June 2015), or the Montana judicial election field experiment in 2014.[1]

The complaints about IRBs are well known: mission creep, as boards move from protecting human subjects to micromanaging research projects and insisting on changes that have little to do with the actual federal rules; bureaucratic delays, as projects are held up by procedural snags and red tape; insistence that researchers comply with administrative requirements that are pure dead-weight, designed more to justify the procedures themselves than contribute to the protection of subjects; a failure (or refusal) to consider the effects of the IRB process on researcher time or the ability to conduct research altogether (Klitzman 2015; Schrag 2009; Citro et al., 2003; Gunsalus et al., 2007; Burris 2008). The lofty goal of protecting subjects has created a system that "has grown into a self-referential, unresponsive, and legalistic bureaucracy"

[1] In writing this article, I benefitted from helpful conversations with Nancy Kendall, Howard Schweber, Dorothy Farrar-Edwards, Lillian Larson, Robert Mattheiu, and Valerie Martinez-Ebers.

(Burris 2008, 65). None of this would be a surprise to those critical of the early efforts to draft regulations in the 1970s, who warned that the rules designed for biomedical research were ill-suited to the social sciences. A 2008 symposium in *PS* identified many issues (Hauck 2008), and little appears to have changed since then

My goal here is not to recapitulate these complaints, but to analyze what things look like from inside an IRB. My views are informed not just by my experience as a researcher who has gone through the process multiple times, but also from the vantage point of 5 years on the Educational, Social and Behavioral Sciences (ED/SBS) IRB at UW Madison, including 2 years as chair. During these 5 years, I spent considerable time talking to faculty and listening to their complaints, meeting with PIs to work through issues with their submitted protocols, calming graduate students certain that a notice of noncompliance for a minor administrative violation was a career-ender, and responding to frustrated colleagues infuriated at what they saw as unnecessary and even arbitrary demands and restrictions. I got into the habit of asking political scientists at other universities about their experiences with their own IRBs, and frequently discussed in the ED/SBS IRB here how we would have dealt with these projects. I am also part of a loosely organized group of faculty here working (very slowly) to promote some institutional reforms.

This experience led me to three conclusions. First, the IRB process is not as bad as it is frequently portrayed, at least as it operates at the UW (though I am certain that some of my colleagues here will dispute this). It is hardly ideal, to be sure, and we sometimes struggled to apply to social science research a framework designed for medical research when the projects we examined often presented very small risks to

subjects. But we made an effort to be judicious, were very careful to avoid mission creep, and focused our attention on those studies and activities that we felt were most likely to pose actual risks to subjects. We did not micromanage research projects or insist on methodological changes unless we saw a direct connect to a concrete risk that could be minimized. I do not recall any cases where we rejected a protocol or where a researcher abandoned a project rather than make the changes that we required. We frequently struggled with the fuzzy boundary between making administrative determinations and establishing policy, and were often worried that we were setting precedents that would create problems for future IRBs. And although we were aware that we had an obligation to implement institutional policies as best we could, there were many times where we questioned those policies.

We pushed back against demands from external actors, including auditors, that we felt were unreasonable. For example, The UW requires all PIs doing research involving non-English speaking subjects to submit consent forms in the relevant language. This is an appropriate accountability check, since even if there isn't an IRB member who speaks, say, Thai or Urdu, we have a written record and the ability to check the documents if any concerns arise. During one external review that took place when I was Associate IRB Chair, we faced pressure to impose a rule requiring all foreign language consent forms to be back-translated into English by independent professional translators. Such a rule, we were told, was necessary to insure that consent forms contained what PIs said they contained. The IRB Chair objected in the strongest terms, arguing that such a rule would do nothing protect subjects, would only serve to impose a monetary cost as a condition of submitting a protocol (it went without saying that PIs

would have to pay for it themselves), and that the UW trusted researchers to act ethically. The UW policy remains unchanged. Another review faulted us for sending too many protocols back for minor changes that could be reviewed by a subcommittee of two IRB members rather than deferrals that had to go back to the full IRB, and for having too many unanimous votes in the IRB. Still another objected to the fact that some consent forms were signed in pencil rather than pen, and wanted a requirement that all consent forms be signed in ink. Such a policy, in my view, would do nothing but cause PIs to roll their eyes in exasperation.

What I told my frustrated colleagues is that they cannot use a "reasonable person" standard to judge the IRB process. They needed to apply an *unreasonable bureaucrat* standard: that someone with the authority to punish the University or deny a certification might question the IRBs work years later, perhaps after a serious adverse event that calls into question our entire process; who demands justification for each decision the IRB has made, and who is no mood to give us the benefit of any doubts. When something goes wrong, the individual researcher is not the one who alone bears the cost. The institution and other researchers there will be seriously harmed if the federal government imposes the nuclear option of shutting down all research. On the other side of the ledger, most compliance costs are borne by the individual PIs, who see little general benefit from the IRB process.[2]

[2] Schrag identified a handful of examples in which PIs said that IRB review improved their projects by identifying problems, or where review headed off research that was potentially harmful. Still, he concludes, "it is vastly easier to find abuses by IRBs than abuses by social scientists. And rarer still are abuses that could have been prevented by prospective review" (2009, 171).

Second, the process can be improved, primarily by recognizing that many procedural requirements are designed to address risks that are so unlikely, and so slight if they do occur, that it would be reasonable to dispense with them. Do minor alterations to interview questions need to be reported and justified? If a survey has a higher than expected response rate and winds up with 850 respondents instead of the originally intended 600, is that really a problem? At the institutional level, policy makers should expand the categories of research that are exempt from review.[3] It could involve, for example, a policy that the risk of a small "dignitary harm" to a subject is so transitory or otherwise innocuous that it will not be considered a substantive harm. The difficulty here, of course, is what happens at the margins, where benign ethnographic or survey research spills over into more sensitive areas, or where seemingly trivial harms turn into something more serious..

During my 5 years on the IRB, I was aware of only a few problems that *could* have resulted in harm to subjects, all involving actual or potential breaches of confidentiality. Most of these occurred through lost or stolen laptops or disks, or file cabinets left unlocked. The most serious involved the unauthorized sharing of data with another researcher, including information about personality disorders, in a study of a vulnerable population. Even in these cases, the latter of which we deemed serious enough to report to the federal OHRP, there was no evidence that I saw of an actual harm to a subject. Complaints that came into the office tended to be minor, such as

[3] Examples are all online or telephone surveys that do not involve face to face interaction with a researcher, and certain types of ethnographic research into benign subjects.

people irritated at receiving an unsolicited mail invitation to participate in a survey that they regarded as intrusive.

Third, IRBs are frequently blamed for problems that do not have anything to do with IRBs, either because research became controversial for reasons other than risks to subjects, or because it was never subjected to IRB review at all.[4] This kind of elision blurs important distinctions between different types of research, and causes spillover effects when problems in one area are used to justify restrictions imposed on all research.

An Overview of the Educational and Social Sciences IRB at Wisconsin

The UW received $879 million in external research funding in 2013-2014, including $660 million from the federal government (UW Madison 2015, 62). In 2012-2013, the UW was ranked 14th nationally among all universities in federal funding for science and engineering R&D (UW Madison 2015, 66-67).

We have 3 IRBs – a Health Sciences IRB, an Educational and Social Sciences IRB, and a Greater than Minimal Risk IRB which considers protocols involving more significant risks. The UW is at the more stringent end of the spectrum of review, as in practice we subject all human subjects research to the IRB process irrespective of funding source. We have two Federal Wide Agreements (FWAs), one with the Department of Health and Human Services, and a separate FWA for research funded by

[4] There may well be an indirect effect operating with research not submitted to IRBs by PIs who consider the process sufficiently onerous they feel justified in ignoring it. These PIs might have been willing to work with a more responsive IRB.

the Department of Defense. We are accredited by AAHRP (Association for the Accreditation of Human Research Protection Programs).

The Educational, Social and Behavioral Sciences IRB reviews about 800 new protocols each year and another 900 or so as continuing review of existing protocols. Of the new protocols, only about 1 in 8 goes to the full committee, with the rest reviewed administratively or by a subcommittee of 2 IRB members. About half of the new protocols are from students (mostly graduate students doing dissertation research).

A back of the envelope calculations provides a rough idea of the time costs of the process. It is a reasonable estimate that preparing and submitting a medium-complexity protocol for IRB review takes about 10 hours, with more detailed projects taking longer.[5] Factor in another 2-5 hours responding to IRB questions for even an expedited protocol, and another 10 or so for protocols that are deferred or require modifications. A conservative estimate is an average of 15-20 person-hours for each new protocol, although complex protocols with multiple PIs or sites could easily be 4 or 5 times that, which suggests about 12,000-16,000 hours per year in researcher time for initial submissions. Submitting materials for the protocols undergoing continuing review or involving changes will take at least 1-2 hours each, adding as much as 1,800 more hours. IRB committee meetings and proposal reviews add perhaps another 750

[5] The tasks specific to IRB review would be crafting an explanation of the research, drawing up a consent form, thinking about risk and procedures to minimize it, working through the web application, and going through training for researchers making their first submission. . An experienced PI who has gone through the process many times might be able to complete it in a few hours, but not many researchers fall into that category. Most graduate students will go through the process only once, and will put in more time.

hours per year for the 14 non-staff IRB members collectively. Add in another 10,000 hours per year for 5 full time IRB staff, and allowing for additional time for consultations, and it is plausible to think that the university spends, conservatively, in the neighborhood of at least 25,000-30,000 person-hours each year on the basics, with most of that load falling on researchers.[6] But we do not know the actual number, as there is no mechanism for tracking how much time is spent on the process.

Does that investment in time and effort – not to mention the social costs of possibly discouraging certain kinds of useful but difficult research, self-censoring of possible projects, or an overall reduction in research activity – produce anything? Can we say with any confidence that these costs make research subjects better off? Or that an alternative model might achieve the same protections with far lower costs? Or that the current system might have the perverse effect of *reducing* protections by diverting attention from real risks as IRBs are buried under administrative trivia or encouraging researchers to take their chances and skip the process altogether?

Hard Cases Make Bad IRB Law

Consider the following case: a PI proposes a study of near-death elderly in Tibet about what happens when they die. The local religious and cultural custom is that death occurs in stages over several hours, even after what in the West would be considered irreversible physical death. The research involved interventions and rituals performed on individuals who under *our* definition of death were dead, but under *their* definition of death were still alive.

[6] This rough estimate does not include the ongoing time devoted to maintaining records and managing compliance over the course of a research project.

Is this human subjects research? If the subjects are dead, then they are clearly no longer human subjects and the IRB had no reason to be involved. If they are alive, they clearly are human subjects.

We struggled with the question of whether we even had the authority to look at this. One the one hand, the subjects were dead at the time the research would take place, at least by our conceptions. On the other, the Common Rule, UW policy, and basic ethics required us to give special consideration to local cultural norms and practices. If the subjects were still alive according to the local religious and social practices, then we had an obligation to respect those practices. There was an additional consideration, which is that the research involved what could reasonably be described as vulnerable subjects (who were elderly, did not speak English, may not have been able to read and write, and may be terminally ill). We concluded that this was human subjects research, and relied on the consent process to respect the autonomy of subjects and address any risks.

This may seem to be an obvious case. An IRB critic might think it was absurd for us to spend any time pondering what it means to be dead. I wonder what *Monty Python* or Jon Stewart could have done with our deliberations. Our interpretation of the rules and application to this protocol were reasonable in my view, but a different IRB could have concluded that this did not constitute human subjects research.

All of the tough cases we dealt with involved this type of ambiguity, including the following examples from protocols we reviewed:

(1) UW Faculty and staff are mandated reporters for child abuse and neglect, and sexual assault, whether observed or revealed. Under this policy, an

otherwise benign home visit or interview could trigger a mandated report if a PI or other investigator observes or learns of evidence of abuse, neglect, or sexual assault. Quite obviously, this is a risk to a research subject. How should the risk be addressed in a consent form? How much detail is necessary? Does this requirement and risk mean that *all* social science research that involves children or people who may have children must be subject to full review?

(2) Studies dealing with risky behaviors or sensitive topics may apply for a Certificate of Confidentiality (CoC) from either the National Institutes of Health or the Department of Health and Human Services. These are legal instruments designed to protect researchers from being forced to divulge identifying information about subjects "in any civil, criminal, administrative, legislative, or other proceeding, whether at the federal, state, or local level."[7] A CoC enhances our confidence that subject data will remain confidential, but does it offer meaningful protection to a research subject if a prosecutor or federal investigator went to court to demand it? The answer, we were advised, is that no one is sure, since there have been few legal tests.

(3) A PI proposed a study involving interviews with people who have either been charged with a crime or are the target of a criminal investigation, about the activities that led to the charges or investigations. Should the IRB be concerned about the risk that their participation could subject them to prosecution or hurt their ability to defend themselves if a district attorney

[7] National Institutes of Health, "Certificates of Confidentiality: Background Information." http://grants.nih.gov/grants/policy/coc/background.htm, accessed July 31, 2015.

demands the researcher's notes? Is informed consent sufficient to address these risks?

(4) How should the IRB assess risks in a project that studies the experiences, activities, and expertise of consultants in a particular school district with the goal of determining whether they improve student outcomes? What if the researcher's hypothesis is that the consultants are not effective (or even competent)? Such a study could subject all of those individuals to the risk of losing their jobs, whether or not they participated in the study. Is this a risk to a *group*, which is not, strictly speaking, a matter of IRB concern, or is it a risk to individuals that must be addressed and minimized?

(5) How high does compensation have to be before it becomes coercive? Compensation of, say, $1,000 for a 10 minute survey would be rejected as absurdly high. But what about $100 for a 5 hour study with a home visit? If $100 is acceptable, is $200? $225? Should researchers be *required* to offer compensation for such a study as matter of respect for individuals, and if so, how much time can a study take before a PI is required to offer it? Does it make a difference whether the subject population is wealthy suburban homeowners, or people recruited from a downtown homeless shelter?

(6) Should survey research be exempt from review, even when the surveys are conducted by private contractors who may want to keep the data they collect, along with identifiers, and use that data for their own research?

(7) How should researchers protect vulnerable LGBT teenagers who are keeping their status secret from their families, or HIV-positive subjects who live in a country with no legal protection against losing their jobs?

(8) For research that we think poses some risk, should we require researchers to modify their studies to reduce those risks, or is it sufficient to merely disclose the risks in the consent form? If we rely on consent forms to address the problem, how much detail must be included? At what point does a consent form morph into a legalistic document that, like an end user license agreement, nobody reads?

(9) How should we evaluate risks in a project that studies attitudes toward government in a country with an authoritarian regime? Can we rely on the proposition that people in those regimes understand the risks of criticizing the government, and that those risks are therefore (as the Common Rule defines it) "not greater in and of themselves than those ordinarily encountered in daily life" [45 C.F.R. 102(j)]?

(10) A researcher who wants to study a African tribal region says that it is necessary to obtain permission from tribal leaders to work in a particular village. Does that permission have any impact on residents' ability to decline to participate in the research (since they know that the elders have approved it)? Can we ignore the risk of coercion because local customs conceptualize permission and consent differently than we do?

These examples have three things in common. The first is that they all required judgment, because it was often unclear what policy or ethical standard applied, or

because multiple regulatory requirements and ethical imperatives conflicted with each other. The second is that they all have more than one reasonable answer. Despite the complexity of the regulations, many critical terms are subject to different interpretations, from the basic properties of "risk" even to the definition of "research" or what a "human subject" is. My own view was that we did not have to be perfect, only that we make the best decision we could with the information that we had, after careful consideration of risks and benefits. But another IRB might, quite reasonably, have come to different conclusions.

The third is that these cases are not easily addressed by simple rules or policies. Even in an institution that is prepared to accept the risk of exempting a wider range of studies or practices, or declare a policy of exempting all survey research, it is not difficult to construct cases that are at the margins of what qualifies as exempt or as risky. *Most* of those risks are very small, others are so unlikely that we can ignore them, and still others fall into categories that perhaps ought not to trigger close review. *Most* could be exempted, but not all, and even then there must be a process for resolving the close cases and for addressing adverse events that are unlikely, but not impossible.

IRBs and Research Controversies: The retracted *Science* Study and the Montana Judicial Elections Field Experiment

As a thought experiment, I will consider whether our IRB would have caught and prevented two recent controversial political science studies. This exercise will

highlight the limitations of IRB review, and raise questions about whether IRBs ought to consider broader ethical issues separate from risk to subjects.

The *Science* Study

In 2014, *Science* published a study (LaCour and Green 2014) demonstrating that canvassers could through short conversations with voters change attitudes towards gay rights and same sex marriage, and could convert same-sex marriage opponents into supporters. It involved almost 12,000 respondents taking part in 7 separate surveys over 9 months, and nearly 1,000 in-person interviews. The in-person effect was strongest among canvassers who were openly gay, suggesting that these conversations had a lasting impact on attitudes. It was a remarkable finding that received widespread attention in the *Washington Post*, the *New York Times*, and on NPR's *This American Life* and *Science Friday* (Konnikova 2015).

Science retracted the study in 2015 at the request of Donald Green, one of the authors. Green had expressed concerns about the validity of the data, the main author's (LaCour) inability to provide the raw data, and discrepancies in claims about funding. When other researchers tried to extend the study, they wondered whether LaCour had even conducted the surveys at all, and concluded it was likely that he had simply fabricated the data (Broockman et al. 2015). These authors found numerous irregularities, as well as evidence that the data were actually taken from an entirely different survey, which "[lead] us to be skeptical that the data were collected as described" (Broockman et al., 2015, 2). The irregularities included an unusually high response rate, denials from the survey firm that appeared to have done the original surveys that they had any involvement in the project, and Lacour's refusal to name the

firm that he says was involved. Lacour has admitted to "lying about aspects of the study, including who paid for it and methodology" (Carey 2015).

If this study had gone through our IRB, I am confident that we would have noted many of the irregularities. We require disclosure of all funding sources, as well as any contracts or agreements between a PI and outside firms involved in any research activities. We would also have required information about any survey firm's data use policies, and would have rejected outright any study that failed to disclose the firm's name. A PI would have to justify the number of subjects to be enrolled, submit survey instruments and interview questions, and demonstrate that the in-person interviewers had completed human subjects training before going into the field. In a project this large and involved, IRB review is not a one-shot process. Graduate students are not permitted to serve as Principal Investigators on either research grants or IRB protocols; a faculty member or other authorized individual is listed as the PI of record on any graduate student protocol, and we require certification that the faculty PI has reviewed it. Substantively, we would have reviewed carefully a study involving in-home visits and conversations on a sensitive topic.

Of course, all of this assumes that there was a protocol. An endnote to the original article stated that the research had been approved by the UCLA IRB (LaCour and Green 2014, 1369), but UCLA later said no such approval had been given (Malakoff 2015). IRBs cannot catch fraudulent research that simply ignores the process altogether. Whether we would catch fraudulent documentation – fake data use agreements, or forged contracts, for example – is a separate question. I am confident

that we would, and the disciplinary consequences to the submitter would be severe. But this case is more about researcher ethics and honesty than about the IRB process.

The Montana Field Experiment

In October 2014, political scientists at Stanford and Dartmouth conducted a field experiment to test whether providing information about candidate ideology would affect turnout in nonpartisan races. Over 100,000 registered voters in Montana received postcards which placed candidates for the nonpartisan State Supreme Court on a liberal-conservative scale and compared them to the placements of Barack Obama and Mitt Romney. The mailings were titled "2014 Montana General Election Voter Information Guide," and included the Montana state seal. Similar mailings went out to voters in California and New Hampshire (Michelson 2014).

The study generated outrage – from Montanans angry about what they saw as interference in their electoral processes, from the Montana Secretary of State who filed a complaint with the state Commissioner on Political Practices (COPP) alleging that the mailing was illegal (Willis 2014), and from Senator John Tester (D-MT), who objected to the presidents of Dartmouth and Stanford about a "so-called research project that uses Montana elections as a political laboratory experiment, at the expense of free and fair judicial elections in our state. . . which was apparently undertaken without the consent of the people of Montana or its state and local election officials" (Tester 2014). Both Stanford and Dartmouth publicly apologized for the study. In May 2015, the Montana Commissioner of Political Practices concluded that the mailing broke the law because it constituted express advocacy and the researches failed to register as a political

committee. He referred the case to prosecutors for possible further action (COPP 2015).

What role should an IRB have played here? The research did not undergo IRB review at Stanford, and appears to have gone through an abbreviated review at Dartmouth. The COPP finding concluded that Dartmouth reviewed only a study of primary voters in New Hampshire, and that they flyer sent out in Montana was never submitted to an IRB (COPP 2015, 4,6).

If the Montana study had gone through our IRB, we would have reviewed it at the full committee because of its size and the lack of consent in the field experiment setting. We would have insisted on seeing the actual materials that went out to voters, and would have noticed the use of the state seal (I expect that our legal counsel, who is an active presence at all IRB meetings, would have flagged it immediately). We would have objected to the use of the seal without permission not just on legal grounds, but because it was deceptive in suggesting that the study had official sanction. If the PIs had the approval of the state election authorities we might have permitted its use. We would have asked the PIs to justify the size of the mailing as well as the unequal treatment of Democratic and Republican respondents (62% of the flyers were sent to Democratic-leaning areas), although we may have been satisfied with a response noting that it was necessary to achieve enough statistical power when the expected effects are small. A PI would probably object that our questions about sample size and the content of mailings constituted unacceptable meddling into research design and methodology.

I know that some IRB members would have objected to this study on the grounds that it might have an effect on an actual election outcome, or because those who received the flyers might be angry if they learned that they were part of a study to which they had not consented. We would have considered the possible psychological harm that a participant may experience after learning they had been subjected to a research study without their consent. But because other individual risks to research subjects were relatively small, it is possible that we would have approved it in modified form.

The objections to the Montana study are, once again, more about ethics than about risk. Jeremy Johnson, an Associate Professor at Carroll College who reviewed the case for the COPP, argues that IRBs must do more than "[focus]narrowly on individual human subjects" and that the *Belmont Report* requirement of respect for "persons, beneficence, and justice" mandates a broader review of research ethics (Johnson 2015, 4). He believes that voters have a right to object to a study "[intending] to manipulate vote totals that could potentially change the outcome of an election" (Johnson 2015, 5). The lack of similar reactions in California and New Hampshire suggest that the problem was not inherent in the research method, but to the specifics of what occurred in Montana.

These are important issues that highlight the difficulties of grappling with ethical questions in the context or IRB review. Insisting that political scientists should do nothing that might affect an election would prevent virtually all research conducted during an election season. Even a narrower view, that research ought not to have an immediate *partisan* effect could preclude research on turnout among marginalized

populations that have identifiable voting patterns, among other things (Bedolla and Michelson 2012, for example). Must we avoid research that might affect an individual's vote choice in any way? At the same time, surely we can draw distinctions between research intending to understand political attitudes and vote choice, and studies that attempt to alter the decision making process or (as in this case) manipulate that process through unconsented treatments.

Eschewing research just because people may object to it would be a mistake, however, subjecting much of what we do to a heckler's veto. Many different groups will object to many different kinds of research on many different grounds. If we are to avoid research that may be controversial, or because people may be upset that a study took place at all, we would have to abandon huge swaths of studies: sex education, school vouchers, drug use, religion in politics, climate change, illegal immigration, HIV-prevention, right-wing political mobilization, left-wing political mobilization; the list is endless. Here be dragons.

Conclusion

If I were to offer one pragmatic suggestion for improving the IRB process, it would be to plead for better data. Even decades after the regulations have been in place, for social science research we lack the most basic information about what we are doing and whether it actually does any good. The 2007 Illinois White Paper put it this way:

> [E]xamination reveals that virtually no scientific evidence is brought to bear on any aspect of the debate about how IRBs function. Unrealistic

> and untested assertions abound. At the microlevel, this includes, for example, how IRBs decide what is adequate and respectful informed consent, what subjects perceive as risk, and what kinds of benefits to subjects and their communities make the relationship fair. At the macrolevel, there is virtually no research on the functioning of IRBs (numbers of protocols reviewed, numbers of serious abuses by discipline, common turnaround time, etc.) or of the effectiveness of the IRB system in protecting human subjects (Gunsulas et al. 2007, 639).

Are we really protecting subjects from harm? Do all of the compliance and reporting requirements make a material contribution to the welfare of research subjects? Does the IRB process impede research that would improve society? My answers are probably, no, and possibly. But the real answer is "nobody really knows."

Bedolla, Lisa García and Melissa R. Michelson. 2012. *Mobilizing Inclusion: Tranforming the Electorate Through Get-Out-the-Vote Campaigns*. New Haven: Yale University Press.

Broockman, David, Joshua Kalla, and Peter Aronow. 2015. "Irregularities in LaCour (2014)." May 29.

Burris, Scott. 2008. "Regulatory Innovation in the Governance of Human Subjects Research: A Cautionary Tale and Some Modest Proposals." *Regulation and Governance* 2:65-84 (No. 2).

Citro, Constance F., Daniel R. Ilgen, and Cora B. Marrett, eds. 2003. *Protecting Participants and Facilitating Social and Behavioral Sciences Research.*

Commissioner Of Political Practices (COPP) of the State of Montana. 2015. *Decision Finding Sufficient Facts to Demonstrate a Violation of Montana's Campaign Practice Law,* McCulloch v. Stanford and Dartmouth, No. COPP 2014-CFP-046. May 11.

Gunsalus, C.K., Edward M. Bruner, Nicholas C. Burbules, Leon Dash, Matthew Finkin, Joseph Goldberg, William T. Greenough, Gregory A Miller, Michael G. Pratt, Masumi Iriye, and Deb Aronson. 2007. "The Illinois White Paper - Improving the System for Protecting Human Subjects: Counteracting IRB 'Mission Creep'." *Qualitative Inquiry* 13:617-649 (No. 5, July).

Hauck, Robert J-P. 2008. "Protecting Human Research Participants, IRBs, and Political Science Redux: Editor's Introduction." *PS: Political Science and Politics* 41:475-476 (No. 3, July)

Klitzman, Robert L. 2015. *The Ethics Police? The Struggle to Make Human Research Safe*. New York: Oxford University Press.

Johnson, Jeremy B. 2015. *McCullough v. Stanford and Dartmouth*. Annex to COPP of the State of Montana 2015.

Konnikova, Maria 2015. "How a Gay-Marriage Study Went Wrong." *The New Yorker*, May 22.

LaCour, Michael J., and Donald P. Green. 2014. "When Contact Changes Minds: An Experiemt on Transmission of Support for Gay Equality." *Science* 346:1366-1369 (December 12).

Malakoff, David. 2015. "Gay Marriage Study Author LaCour Issues Defense, But Critics Aren't Budging." Science Insider, May 30. http://news.sciencemag.org/scientific-community/2015/05/gay-marriage-study-author-lacour-issues-defense-critics-arent-budging, accessed July 28, 2015.

Michelson, Melissa R. 2014. "Messing with Montana: Get-out-the-Vote Experiment Raises Ethical Questions." *The New West* (official blog of the Western Political Science Association). https://thewpsa.wordpress.com/2014/10/25/messing-with-montana-get-out-the-vote-experiment-raises-ethics-questions/, accessed August 1, 2015.

Schrag, Zachary M. 2010. *Ethical Imperialism: Institutional Review Boards and the Social Sciences, 1965-2009.* Baltimore: Johns Hopkins University Press.

Tester, John. 2014. Letter to John L. Hennessy and Philip J. Hanlon. October 24. Reprinted in *Billings Gazette*. http://billingsgazette.com/tester-s-letter-to-

stanford-and-dartmouth/pdf_79c534eb-75ed-50f8-9c67-883e7beb6fa4.html. Accessed July 28, 2015.

UW Madison. 2015. *Data Digest 2014-2015*. Office of the Provost, Academic Planning & Institutional Research, and Vice Chancellor for Finance & Administration (March). http://apir.wisc.edu/datadigest/201415Digest/dd15_web.pdf.

Willis, Derek. 2014. "Professors' Research Project Stirs Political Outrage in Montana." *New York Times*, October 28.