



The Experimental Political Scientist



In this issue

- [Triangle Table on GOTV](#)
- [Clarke on Neyman-Rubin...O'Reilly?](#)
- [Snowberg on Selective Trials](#)
- [Grossman on Lab in field](#)
- [Mentor Program Recap](#)
- [Announcements](#)

From the Editor

Welcome to the next issue of the Experimental Political Scientist. In this issue we have some exciting features. Leading off is commentary from several leaders in GOTV research. With primaries, elections, etc., coming up, the feature is highly relevant. Next we have some thoughts that challenge the increasing reliance on the Neyman-Rubin approach to causal inference. Hopefully the newsletter can serve as a forum for criticism of conventional practices, and my thanks to Kevin for offering his critical thoughts. Next Erik Snowberg and colleagues describe some exciting prospects for new types of experimental designs, "selective trials." Guy Grossman follows up with an excellent review of lab in the field work in comparative politics. Finally, we have several entries from participants in the new mentor program. Thanks to all the mentors. PLEASE send contributions (or ideas) for the next newsletter, deadline May 1st. Happy experimentation! Dustin Tingley, Harvard Government Department


Information on Joining or Contributing

The *Experimental Political Scientist* is the official newsletter of APSA Organized Experiments section 42. To receive the newsletters register for the section (\$8/yr!) at <http://www.apsanet.org> and visit us at <http://ps-experiments.ucr.edu/group>. Graduate students and professors are highly encouraged to submit to the newsletter. This and the previous issues provide examples for submission themes. Questions can be sent to the [editor](#).

The Emergence of the APSA Experimental Section

In my last letter for this newsletter, I reported on the many exciting section activities in advance of the APSA meeting. I am happy to report that all were very successful - including a set of fantastic panels, an extremely impressive mentor program, and the most engaged and well attended business meeting that I have ever seen. Indeed, over 100 people attended our business meeting. At the meeting, we distributed awards, nominated new officers (a list of which appears elsewhere in this newsletter), received updates on membership/finances, and heard from the junior scholars committee. We also discussed the ongoing work of the standards committee as they work on developing a set of experimental reporting standards. Perhaps most importantly, we continued a dialogue on an experimental political science journal. We have decided to move forward, asking the section to vote on whether to support such a journal. I hope you take the time to cast your vote in the very near future. This is an exciting time for experimental political science - already this year, a number of experimental books have appeared including Diana Mutz's *Population-Based Survey Experiments* and the *Cambridge Handbook of Experimental Political Science* (with contributions from a large number of section members). Section membership has increased at an impressive pace and I imagine that in a few years, we will be among the most popular sections. All of this success is due to the hard work and commitment of the section's members and so I thank you all for making experimental political science a key part of the discipline.

Jamie Druckman
Northwestern University



Forum on GOTV Experiments

Harnessing Elemental Psychological Forces to Promote Voter Turnout

Donald Green
Columbia University
dpg2110@columbia.edu

Given the vast number of field experiments on voter turnout since the New Haven study (Gerber and Green 2000), it's not unreasonable to wonder whether too much scholarly effort is being invested in a relatively narrow line of research. Study after study assesses the effects of voter mobilization tactics (e.g., mail, leaflets, email, text messages, automated phone calls, live phone calls, radio, canvassing). These experiments are replicated in election after election, sometimes with a focus on specific types of voters, such as minorities or young people. This experimental agenda has gradually spread to other countries, and researchers have recently examined the effectiveness of get-out-the-vote tactics in Benin, Brazil, Canada, China, Ghana, India, Mexico, Nigeria, Pakistan, and the United Kingdom. Although the number of field experiments outside the domain of voter mobilization has grown dramatically over the past decade, I suspect that the number of voter mobilization experiments nonetheless exceeds the number of field experiments on all other political topics.

What is so attractive about voter mobilization experiments, and do any untapped opportunities remain in this well-tilled field of research? Speaking only for myself, I find this line of research appealing because it provides an opportunity to test fundamental social science propositions on a grand scale. Although experimental studies of voting have generated many useful practical insights into how to run an efficient voter mobilization campaign (which Alan Gerber and I have catalogued in our book *Get Out The Vote: How to Increase Voter Turnout*), even more important has been their contribution to longstanding social science debates on such topics as persuasion, social norms, collective action, social networks, and habit formation. Although it is true that some field experiments (including some of my own) attempt nothing more than a rigorous evaluation of a particular GOTV tactic,

even these narrowly focused experiments sometimes contribute to theoretical insights. For example, the dozens of field experiments on the effects of partisan and nonpartisan direct mail showed quite convincingly that the typical piece of direct mail has negligible effects on turnout, regardless of whether it appeals to ethnic identities, provides information about when and where to vote, or endorses a candidate's qualifications or issue stances. On the surface, these failed efforts to increase turnout seem like an exercise in futility, but from a theoretical standpoint, it's instructive to learn what doesn't work. For example, the ineffectiveness of mail that provides information about when and where to vote suggests that "information costs" play a relatively minor role in explaining why registered voters abstain.

Where do we go from here? One of the exciting challenges is to marshal social psychological theory in order to come up with especially effective treatments. Whenever I hear scholars speak in generalities about the powerful effects of source credibility, descriptive norms, prescriptive norms, elite cues, social identities, selective incentives, networks, and so forth, I think to myself: would these forces prove powerful if used to stimulate voter turnout? Other scholars seem to be thinking along similar lines, and in the last few years, a growing number of research articles have attempted to test the effects of various theory-driven interventions. This turn has produced some important discoveries, and for the first time experimentally tested interventions are generating sizeable effects. (Social scientists can take some satisfaction from the fact that these theory-driven interventions tend to have much larger effects on turnout than the voter mobilization tactics that campaigns conventionally use.) We need more of this sort of research, based on a more systematic inventory and assessment of behavioral propositions.

Suppose these tests uncover ways of increasing turnout. Obviously, from a practical standpoint, highly effective interventions are useful because they generate votes. But what do effective interventions do for us as social scientists? From a theoretical standpoint, highly effective treatments provide researchers with a "downstream" opportunity to study an array of social phenomena. For example, suppose a random inducement to vote in the days leading up to the 2012

presidential election makes subjects in the treatment group substantially more likely to vote than subjects in the control group. The stage is now set for a variety of other investigations. Does voting appear to be habit forming – are subjects assigned to the treatment group more likely to vote in subsequent elections than subjects in the control group? Do the effects of a strong intervention spillover to other voters in the same household or neighborhood? Are public officials more interested in the political concerns of citizens whose voter turnout has increased exogenously?

Granted, the topic of voter turnout is considered a bit drab by many social scientists. But the fact remains that voter turnout is quite possibly the most attractive experimental testing ground anywhere in the social sciences. Experiments can be conducted on a grand scale and replicated in a variety of settings, with outcomes measured courtesy of government. As researchers have caught on to the opportunity that turnout research affords, research has become more nuanced both in terms of the psychological theories that are being used to mobilize voters and in terms of social psychological processes that are revealed once a successful intervention has taken place. Even so, researchers have barely scratched the surface. If you want to figure out what's next in this domain, open up an introductory textbook in social psychology.

Looking Beyond Election Day: Voter Mobilization Experiments and Pre-Election Day Voting

Christopher Mann
University of Miami
cmann@miami.edu

The literature on voter mobilization is the most developed field experiments literature and the literature showing the most divergence from current trends in politics. The central concern in GOTV field experiments has been how to get voters to the polling place on Election Day. Over the last 12 years, a variety of impressive experiments have investigated the effect of campaign tactics, political messages, and psychological mechanisms to get voters to overcome the costs of going to the polling place on Election Day.

During roughly the same period, there has been a revolution in the way Americans vote. In the 2008 and 2010 General Elections almost one-third of all ballots were cast without going to the polls on Election Day (Early Voting Information Center 2011). Instead, voters cast ballots long before "Election Day" by going to an early voting center or sending a ballot by mail. These changes in the way Americans vote have created a sizable deficit in the GOTV field experiments literature. Whether we measure the importance of pre-Election Day voting by the share of ballots (approximately 30% in 2008 & 2010) or the 32 states that allow it (National Council on State Legislatures 2011), it is quite clear that the attention to pre-Election Day voting in the current GOTV field experiments literature falls woefully short of representing what is happening in American politics. For example, I could find only one experiment on voting by mail Green & Gerber's (2008) review of the GOTV field experiments literature, one experiment on early in-person voting in Morton & Williams's (2010) review of the experimental method in political science, and no references to pre-Election Day voting in the forthcoming Cambridge Handbook on Experimental Political Science (Druckman, Green, Kuklinski and Lupia Forthcoming).

The current GOTV field experiments literature is essentially a search for what makes voters overcome the cost of going to the polls on Election Day. This narrow focus is somewhat understandable, since the process of casting a ballot was a constant for most of the 20th Century following the adoption of the secret ballot.

Pre-Election Day allows new research on variation in the cost, timing, and location of voting. Early in-person voting allows voters to choose when it is convenient to vote over a period of days or weeks and where it is convenient for them to vote. These choices appear to reduce the direct costs and opportunity costs associated with going to a particular polling place during particular hours on a particular Tuesday. For voters who cast their ballot by mail, the voting process is more significantly transformed: "Election Day" is any time day or night over a period of weeks and "polling places" are literally anywhere the voter can fill in the ballot. As a result, the direct costs and opportunity costs of casting a ballot are lower, the social benefits altered (e.g. showing up at the neighborhood polling isn't observed by the neighbors), and vote choices are

an open-book test (plus open-browser, open-newspaper, open-conversation with friends, etc).

These fundamental transformations in the voting process present challenges and opportunities to study the behavioral, psychological, and institutional aspects of voting:

Do the (hypothesized) lower cost of pre-Election Day voting increase turnout? Or is the cost of acquiring information about these new modes of voting a barrier to use, particularly for low propensity voters? What happens if civic and political organizations educate and recruit voters to use the new modes of voting?

What mobilization tactics are the most successful at increasing participation for voting by mail? For early in person voting? Do the same social psychological mechanisms that increase turnout on Election Day work in pre-Election Day voting?

Which mode offers the greatest opportunities for civic groups, political organizations, and/or election officials to increase voting participation? Do voters make separate decisions about whether to vote using each mode or just make one decision despite multiple opportunities? Should civic and political organizations attempt to mobilize voters for different modes as they arise sequentially (i.e. by mail, then early in person, and finally Election Day) or is this a waste of money?

What is the effect of changing voting institutions? How does allowing pre-Election Day voting - especially providing choice among different modes of voting - affect voting behavior? While it is not possible to randomly assign voting laws to jurisdictions, we can get some leverage on the impact of voting reforms by randomly assigning treatments that vary the information provided about voting alternatives.

What happens when jurisdictions conduct elections entirely by mail? Traditional Get-Out-The-Vote is replaced by Get-The-Vote-In-The-Mailbox-On-Time-And-Signed. How do voters learn to correctly use new methods of voting? How can election officials and/or campaigns motivate voters to overcome different barriers from different modes of voting?

The questions above have motivated the research that I (and invaluable co-authors) have done in the last couple of years using more than a dozen large scale

field experiments about voting by mail and early in person voting. Others have begun to conduct field experiments about pre-Election Day voting as well (e.g. Arceneaux, Kousser, and Mullin 2011; Monroe and Sylvester 2011). The results of these experiments provide valuable insight about these modes of voting and voting behavior in general. However, we have just begun to learn from applying the experimental method to the sweeping changes in pre-Election Day voting.

The process of voting is changing rapidly in the United States, and - as the Spanish and German vote by mail campaigns posters in my office attest - in many other countries as well. These changes are a challenge and an opportunity to study the behavioral, psychological, and institutional aspects of voting. A decade ago, field experiments fundamentally changed the study of voting. Now, the far-reaching changes in voting call for a shift in the focus of GOTV field experiments so that we can understand why and how tomorrow's voters will decide to vote.

References

- Arceneaux, Kevin, Thad Kousser and Megan Mullin, 2011. "Get Out the Vote-by-Mail? A Randomized Field Experiment Testing the Effect of Mobilization in Traditional and Vote-by-Mail Precincts", *Political Research Quarterly*, Published online September 30, 2011 (doi:10.1177/1065912911421013). Druckman, James N., Donald P. Green, James H. Kuklinski and Arthur Lupia, editors. Forthcoming. *Cambridge Handbook of Experimental Political Science*, Cambridge University Press. Early Voting Information Center at Reed College, 2011. "Pre-Election Day Voting Statistics", <http://earlyvoting.net/>. Green, Donald P. and Alan S. Gerber, 2008. *Get Out the Vote: How to Increase Voter Turnout* (2nd Edition), Brookings Institution Press. Monroe, Nathan W. and Dari E. Sylvester, 2011. "Who Converts to Vote-By-Mail? Evidence From a Field Experiment", *Election Law Journal*, March 2011: 15-35. Morton, Rebecca B. and Kenneth C. Williams, 2010. *Experimental Political Science and the Study of Causality: From Nature to the Lab*, Cambridge University Press. National Conference of State Legislatures, 2011. "Absentee and Early Voting", <http://www.ncsl.org/default.aspx?tabid=16604>.

When the Client Owns the Data

David Nickerson
University of Notre Dame
dnickers@nd.edu

Under the tutelage of Don Green and Alan Gerber, I began using field experiments to study voter mobilization in 2000. For the first six years of my career, experiments were a hard sell. Both the non-profit and partisan worlds were dominated by consultants who used year-to-year comparisons, simple accounting of voter contacts, client surveys, focus groups, and qualitative case studies to establish the effectiveness of programs. These traditional modes of evaluation nearly always showed that the campaign or civic organization was very successful in mobilizing voters. Experiments were viewed suspiciously because the creation of a randomly selected control group to serve as a baseline for comparison set the bar much higher and could easily show a program to be ineffective. Thus, groups routinely refused to conduct experiments or scuttled them at the last minute.

This environment created an interesting dynamic tension. Groups were desperate for reliable information on how best to mobilize constituents, but unwilling to participate in them. A childhood friend who worked for a high powered DC consulting firm explained to me in a 2002 email,

“I love your work. Send me any results you have, even if they aren’t up to publication standards. It is better information than I have to work with. But there is no way we will let you study us. Everyone thinks our campaigns work. The only thing you can do is show people our program doesn’t work.”

The fear of being shown to be ineffective was not unfounded. Most of the civic organizations whose evaluations reported null findings watched helplessly as funders declined to fund future proposals.

By 2006, the voter engagement culture had completely reversed itself and experiments became a product that sold itself. Groups are now eager to collaborate with academics and get access to expertise cheaply. On balance, the collaboration between civic engagement groups, campaigns, and academics has been beneficial. However, the fear of being punished for null findings

remains and creates the possibility for publication bias.

The advantages of working with a campaign are self-evident and compelling. When academics conduct research, there is always the fear that the treatments do not accurately reflect actual campaign activity. This concern about construct validity is largely eliminated when the treatment is designed and administered by the organization to which the treatment is meant to generalize. Furthermore, since campaigns spend huge sums to contact voters, experiments conducted in collaboration with campaigns can also be much larger. The most I have ever spent on an experiment out of my modest research account is \$16,000, which wiped out all my available resources. In contrast, I have experimented on campaign interventions that cost millions of dollars. The ability to estimate precise treatment effects on a wide scale requires this type of investment that is typically well beyond what academics can provide. In fact, working with the actual organizations conducting the voter outreach may be the only way to convincingly study the effectiveness of many campaign activities of interest to political scientists.

There are two downsides to conducting experiments with campaigns, one of which is relatively minor and the other is more worrisome. The minor concern is that researcher control over the experiment is minimized and poses two chief problems. First, the need for the experiment to tightly test the proposed theory may be trumped by specific goals and logistical limitations of the campaign. Second, it is also possible that important aspects of the experiment such as the assignment of treatment or measurement of the dependent variable were conducted poorly by someone on the campaign. In one sense, the correct response to the first problem is to shrug one’s shoulders: the experiment is what it is. As long as the author describes the experiment accurately, readers can evaluate how well the data speaks to the theory. Similarly, the second problem can be solved by training and overseeing the people managing data on campaigns better. In the decade that I have worked with practitioners, I can attest to the staggering evolution from low sophistication campaigns to widespread knowledge of how to manipulate and analyze data. So the lack of researcher control of the experiment itself does not concern me.

The bigger concern is the lack of control researchers have over the release of the experiments. Groups still fear (legitimately) that "disappointing" findings will negatively affect their organization. This dynamic could lead to the selective disclosure of experimental results. For illustrative purposes, suppose 20 organizations were testing the effectiveness of direct mail for altering vote choice with researchers. Suppose further that 19 of these experiments came back with null (or even negative) findings and only one showed a positive and statistically significant result, which would be perfectly consistent with a true effect of zero for mail on persuasion. However, if the organization only let the researcher publish the results of the one "positive" result, readers would be left with the mistaken impression that mail was very effective in persuading voters. The gate keeping authority of the organization is what drives this version of publication bias rather than the researcher's personal file drawer problem (Gerber et al. 2001). The one researcher releasing her results may be completely unaware her findings are spurious, however, the net bias in the published literature is the same.

So what can be done to combat this publication bias? My general practice has been to negotiate the right to publish the results of the experiment up front and promise the organization that their identity will not be revealed. This practice allows me to publish less than flattering results (to the extent that null findings are publishable, see Gerber and Malhotra 2008) but is not a cure all. First, while most readers will not be able to guess the identity of the organization studied, funders will almost assuredly come across the paper and figure out the identity of the organization since they are familiar with the work and researcher / organization networks. Second, organizations are increasingly viewing the results of experiments as valuable and secrets to be kept from competitors. Thus, many simply will not agree to joint ownership of the data and the researcher has to hope the organization will grant release of the proprietary data.

The medical sciences handle the problem of publication bias with proprietary data by requiring all experiments to be registered if they are to be published or used for approval. This solution does not quite work in political science since campaigns do not require FDA approval to implement techniques learned through experiment.

Furthermore, registering experiments may tip off competitors on the activities of competitors putting the campaign at a strategic disadvantage. Thus, such a registry would serve as a strong disincentive for campaigns collaborating with academic researchers and curtail the growth of that literature. Perhaps it would be possible to link the experiment to the researcher and make the existence of each experiment public after a period. However, many organizations will not trust that the experiments will not be linked back to the organization, so the chilling effect will only be slightly dampened.

A more modest proposal would be to require researchers to disclose when experiments have a proprietary provenance. A simple statement that the campaign owns the data and allowed the researcher to publish the results will allow the reader to update the probability of publication bias. The magnitude of this bias will be difficult to assess without many publicly available replications, but the researcher may not even know the distribution of treatment effects. Such ownership disclosures may not solve the problem presented by organizations selectively releasing proprietary results, but by alerting readers to the problem it puts experiments conducted by campaigns on the same level of confidence as other experiments.

References

- Gerber, Alan S., Donald P. Green and David Nickerson. 2001. "Testing for Publication Bias in Political Science," *Political Analysis* 9(4):385-392. Gerber, Alan S. and Neil 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.

Kevin Clarke
Department of Political Science
University of Rochester
kevin.clarke@rochester.edu

Neyman-Rubin...O'Reilly?

Counterfactuals seem to be very much on the minds of political scientists these days. Whether in the back of

dingy hotel conference rooms, the august pages of the *American Political Science Review*, or in faculty meetings amongst the weak coffee and damp danish, there are demands to consider the counterfactual and arguments over what the appropriate counterfactual actually is. The roots of our obsession can be traced to two beliefs that permeate nearly every corner of political science. The first is that Causal Inference should be, and is, the goal of all inquiry in political science. (This belief always puts me in mind of that E*Trade commercial where the voice over intones “No one ever woke up one morning and said ‘I just want to be an ordinary scientist,’ or ‘an ordinary writer,’ or even ‘an ordinary runner,’” while images of Stephen Hawkings, Ernest Hemingway, and Jackie Joyner Kersee scroll past. I imagine an ad that goes “No political scientist ever woke up one morning and said ‘I just want to do description,’” while images of Bill Riker and Warren Miller scroll past.) The second is that the problem of causality has been solved, and the solution involves counterfactuals. There is little that I can do about the first belief in these pages, but there may be something I can do about the second.

My purpose here is not to say that counterfactual accounts of causation are wrong and that some other kind of account is right. Rather, I intend to show that counterfactual analyses are plagued by the same problems that beset all other analyses, namely, they simultaneously admit too little and exclude too much. Philosophers often do this type of work through the use of counterexamples, and I will do the same. For political scientists uncomfortable with this mode of analysis, think of it as assessing the validity of a variable. Any operationalization captures some parts of the concept in question and misses others. At the same time, any operationalization includes parts of other concepts that should have been excluded. A first step in assessing the validity of a measure is to check its face validity; a useful measure of ideology should place Tom Coburn to the right of Chuck Schumer. Analyzing counterfactual accounts of causation works in the same way. We have commonsense notions of what is, and is not, a cause, and valid accounts of causation should be able to distinguish between the two. After all, if

counterfactuals do not perform well when we have good intuition, why would we expect them to perform well when we are bereft of intuition and must turn to statistics?

Let us begin with an example in which the counterfactual account admits too little. One of the mantras of those who adopt the counterfactual approach is that there exists “no causation without manipulation.” That is, no event is to be considered a cause unless it can be experimentally manipulated. Thus, Holland concludes that attributes of units, such as race or gender, cannot be causes because they could not serve as a treatment in an experiment.¹

Time for our first counterexample. Most schoolchildren will tell you that tides are caused by the gravitational force of the moon.² The moon’s gravitational force, however, cannot be manipulated, and it is difficult to see how it could be used as a treatment in an experiment. The counterfactualist must therefore conclude that he does not know what causes the tides, which puts him in league with Bill O’Reilly, who infamously stated that the tides cannot be explained by science. (I have to assume that aligning oneself with the forces of darkness and ignorance is not a comfortable place for most political scientists.)

Astronomy is pretty far removed from political science, but analogous situations are easily found. A large proportion of the rational choice literature concerns the effects of institutions on behavior. Whether of the brick and mortar kind or not, most institutions are not amenable to experimental manipulation. Institutions are generally integral parts of systems in the same way that the gravitation force of the moon is part of the interplanetary system, and for the counterfactualist, attributes of systems cannot be causes. Nonetheless, institutions often play a significant role in shaping and constraining the behavior of actors, and few of us would hesitate to label institutions as causes.

Now let us consider cases where counterfactual accounts admit too much. The following three counterfactuals are political science variations on examples given by Kim (1973).

¹It is not enough for the attribute to be in principle manipulable. Holland argues that changing an attribute of a unit means that it is no longer the same unit.

²Technically, the interaction between the gravitational forces of the moon and the sun and the rotation of the Earth.

1. If German troops had not crossed the German-Polish border, Germany would not have invaded Poland.
2. If we had not held a vote, we would not have passed the budget.
3. If I had not pulled the lever in the voting machine, I would not have voted for Obama.

All three examples are perfectly acceptable counterfactuals, and most counterfactual accounts of causation would admit the premise of each as being causal. In all three examples, however, our intuition tells us that there is no causal relationship. In example (1), the dependency is tautological; troops crossing a border is the definition of an invasion. The real cause of the invasion lies in the strategic situation facing Germany in 1939. In (2), one event, holding a vote, is a constituent part of the second event, passing the budget. Here the cause lies in balancing the interests of the chamber. In (3), an agent, the voter, does an action, voting for Obama, by doing another action, pulling the lever of the voting machine. The cause of voting for a candidate, however, is not in the action of voting, but the real or expected stream of benefits received by the voter. A more useful account of causation would rule (1)-(3) out.

There are still other problems that plague counterfactual accounts. Consider the following:

If Archduke Ferdinand had not been assassinated, World War I would not have started.

Even if we assume that this counterfactual is true, it requires a number of other counterfactual “causes” to be true as well. The assassination would not have caused the war if the alliance system was not in place, if there had been no arms race, or if nationalism had a weaker hold on the continent. The point is that the context within which the assassination occurred determines, in part, whether the assassination is a cause of the war.

Overdetermination is another problem for counterfactuals, one that proponents have

acknowledged. An event is overdetermined if there are two or more sufficient causes of it. Mackie (1974) provides the following illustration:

Lightning strikes a barn in which straw is stored, and a tramp throws a burning cigarette butt into the straw at the same place and at the same time: the straw catches fire.

The counterfactual “if lightning had not struck, the straw would not have caught fire” is false; the cigarette butt would have started the fire. Similarly, the counterfactual “if the tramp had not thrown the cigarette butt, the straw would not have caught fire” is also false. A counterfactual analysis therefore concludes that neither the lightning nor the cigarette butt is the cause of the fire. Our intuition, however, tells us that both caused the fire. If you think over determination does not occur in political science, consider the causes of voting democratic by an overeducated son of New England, with professors for parents, who attended Oberlin College and the University of Michigan.

To reiterate, my purpose in rehearsing these problematic examples is not to claim that counterfactual accounts of causation should be abandoned in favor of some other account. I want to emphasize that counterfactual accounts are, in some sense, no better or worse than other accounts, and the best known of these accounts continues to evolve in attempts to deal with these and other problems (see Lewis 2000)). All accounts of causation are caught between the Scylla of ruling in events we normally dismiss as causes and the Charybdis of ruling out events we normally think of as causes.³ There is little reason to treat counterfactual accounts as privileged, and proponents who sniff that counterfactuals provide the only *principled* way of thinking about causation should be dismissed out of hand (see Eells (1991) on probabilistic causation for an alternative). If we cannot privilege counterfactual accounts of causation, we cannot privilege the statistical frameworks, such as the Neyman-Rubin model, that depend on them. So next time someone on the other side of the damp danish demands to know the counterfactual, ask why he or she agrees with Bill O’Reilly about the tides.

³The analogy is not perfect; in Greek mythology, Charybdis was considered the greater danger.

Eells, Ellery. 1991. *Probabilistic Causality*. New York: Cambridge University Press.
Kim, Jaegwon. 1973. "Causes and Counterfactuals." *Journal of Philosophy* 70 (October): 570-572.
Lewis, David. 2000. "Causation." *Journal of Philosophy* 97 (April): 182-197.
Mackie, J.L. 1974. *The Cement of the Universe*. New York: Oxford University Press.

Selective Trials

Sylvain Chassang
Princeton

chassang@princeton.edu

Gerard Padro i Miquel
LSE

g.padro@lse.ac.uk

Erik Snowberg
Caltech

snowberg@caltech.edu

1 Introduction

Political scientists use randomized controlled trials (RCTs) to address concerns that people who opt into a certain treatment—say, subscribing to a newspaper—have unobserved, systematic differences from those who do not. An RCT randomly assigns treatment, avoiding selection bias.

However, it is difficult to interpret the results of an RCT in the presence of unobserved subject effort. Those who are, say, randomly assigned to receive a newspaper subscription may be unlikely to read the newspaper when it arrives. If an experimenter then observes that those who received a free newspaper subscription have the same level of political knowledge as those who did not, she might wrongly conclude newspapers are ineffective at conveying political knowledge, when it may be that no one in the experiment actually read the newspaper.

In Chassang, Padró i Miquel and Snowberg (forthcoming) we introduce a new framework for experimental design. The designs identified using this framework, which we call selective trials, allow experimental subjects to express preferences over treatment—to control for unobserved effort—while maintaining randomization—to control for selection. Thus, our approach creates designs for analyzing heterogeneous treatment effects, rather than ex-post statistical tools. This article introduces the basic concepts of selective trials through a simple example, briefly discusses what is possible in more general environments, and concludes with some simple guidelines about how to determine whether selective trials may be useful in an experiment.

2 A Simple Example

To understand how selective trials can improve inference, we must first formalize the example above of an experiment to evaluate whether newspapers increase political knowledge. This can be seen as a highly stylized version of the experiment in the excellent study by Gerber et al. (2009), in which we consider only a single newspaper—rather than two—and a single outcome measure—political knowledge.

Gerber et al. (2009) finds that giving a free newspaper subscription to someone who does not already subscribe to a paper does not change political knowledge. This may be because most people who do not subscribe to a newspaper will not read a newspaper, even if it is given to them for free. However, there may still be some people in this group who value newspaper subscriptions, just less than their price, and would read the newspaper if they

had one. Intuitively, a selective trial determines who, among non-subscribers, may value and read a newspaper by offering it to them at a small price, while at the same time randomizing who is offered the subscription at a small price. The formalization makes it possible to precisely show this intuition, and how to calculate treatment effects from the resulting data.

For all trial formats, subjects $i \in \mathbb{N}$ are assigned to the treatment group $\tau_i = 1$ and receive a newspaper subscription, or are in the control group $\tau_i = 0$, and do not receive a subscription. The outcome of interest y_i —a subject’s political knowledge—is determined by whether a subject gets a newspaper subscription, τ , how much effort $e \in [0, 1]$ the subject puts into reading the newspaper, how effective the newspaper is at changing political knowledge R , and a random baseline q_i , independently drawn from some distribution which may be affected by time, location, and political environment:

$$y_i = q_i + R\tau_i e_i.$$

The amount of effort each subject puts into reading the newspaper (if treated) is determined by his or her value for reading the newspaper, and the cost of doing so,

$$V_i e_i - c_i e_i,$$

where both the value and the cost of reading the newspaper may be determined by a number of idiosyncratic features, such as whether or not a subject’s favorite sport is in season during the experiment. Normalizing each subject’s cost of reading the newspaper to one, $c_i = 1, \forall i$, we have that if a subject is treated, he or she will read the newspaper ($e_i = 1$) if $V_i \geq 1$. We assume that V_i is distributed according to some distribution F_{V_i} , which need not be known to the experimenter or subjects.

Consistent with the experimental design in Gerber et al. (2009), we assume all experimental subjects are not currently subscribed to a newspaper. Thus, all subjects must have a low value, or high cost, for reading the newspaper. That is, if a newspaper subscription costs p_0 , then for all subjects, $V_i - c_i < p_0$. However, this does not mean that a treated subject would not read a free paper, provided that $V_i - 1 > 0$.

We contrast two ways of running the experiment:

1. An RCT, where each subject is randomly, and independently, assigned to receive treatment or not, with 50% probability.
2. A selective trial, which adds a second arm to the RCT above. The subjects that are randomly assigned to the second arm are given the opportunity to purchase a newspaper subscription at a deeply discounted rate $p \ll p_0$.

Inference from an RCT. The RCT allows for the identification of the average treatment effect:

$$\begin{aligned} \mathbb{E}[y_i | \tau = 1] - \mathbb{E}[y_i | \tau = 0] &= \mathbb{E}[q_i + R \times \mathbf{1}_{V_i \geq 1} | \tau = 1] - \mathbb{E}[q_i | \tau = 0] \\ &= R \times \text{Prob}(V_i > 1) = R \times (1 - F_{V_i}(1)). \end{aligned}$$

If the experimenter knows $F_{V_i}(1)$, then she can determine the effect of reading the newspaper on political knowledge. However, this will not generally be the case. Thus, if, as in Gerber et al. 2009, the average treatment effect is statistically indistinguishable from zero, it may be because newspapers are ineffective at changing political knowledge (R is low), or no-one in the experiment read their free newspaper ($F_{V_i}(1)$ is high).

Selective Trials and the Value of Self-Selection. Now consider data from the arm added by the selective trial. Note that a subject will pay p for the newspaper subscription if and only if $V_i - 1 \geq p$. As this implies that $V_i \geq 1$, everyone who pays for the subscription will read it.

As everyone who pays p for the subscription is treated, and all those who don't are untreated, the second arm identifies $\mathbb{E}[y_i|\tau = 0, V_i \leq 1 + p]$ and $\mathbb{E}[y_i|\tau = 1, V_i \geq 1 + p]$. Moreover, note that the proportion of people who are willing to pay p identifies $\text{Prob}(V_i \geq 1 + p) = 1 - F_{V_i}(1 + p)$. However, one cannot simply compare subjects that paid and did not pay for the paper, as those that paid may systematically differ from those that did not. In particular, q_i is likely correlated with V_i .

To recover the effect of reading the paper on political knowledge, we must add data from the RCT. Noting that by the law of iterated expectations

$$\mathbb{E}[y_i|\tau = 0] = \text{Prob}(V_i \leq 1 + p)\mathbb{E}[y_i|\tau = 0, V_i \leq 1 + p] + \text{Prob}(V_i \geq 1 + p)\mathbb{E}[y_i|\tau = 0, V_i \geq 1 + p],$$

we can thus identify

$$\mathbb{E}[y_i|\tau = 0, V_i \geq 1 + p] = \frac{\mathbb{E}[y_i|\tau = 0] - (1 - P(V_i \geq 1 + p))\mathbb{E}[y_i|\tau = 0, V_i \leq 1 + p]}{P(V_i \geq 1 + p)},$$

where all the terms in the left hand side are identified from data. Putting these pieces together we have that

$$\begin{aligned} & \mathbb{E}[y_i|\tau = 1, V_i \geq 1 + p] - \mathbb{E}[y_i|\tau = 0, V_i \geq 1 + p] \\ &= \mathbb{E}[q_i + R \times \mathbf{1}_{V_i \geq 1}|\tau = 1, V_i \geq 1 + p] - \mathbb{E}[q_i|\tau = 0, V_i \geq 1 + p] \\ &= R \end{aligned}$$

recovers the effect of reading the newspaper on political knowledge. As the selective trial contains a randomized controlled trial, this information is *in addition to* what an experimenter would obtain with a standard RCT. This is by design: all selective trials are at least as informative as RCTs.

Note that the second arm of the selective trial mimics an observational study. This suggests that if, for some exogenous reason, part of an (otherwise identical) population had no access to newspapers while the other part did, then running an RCT in the part that did not have access would allow similar identification. However, as such opportunities are likely to be rare, this situation can be replicated by the experimenter randomly assigning which part of the population will be in an RCT and which part will be in an observational study.

3 Generalized Selective Trials

The example above is obviously too stark. Chassang, Padró i Miguel and Snowberg (forthcoming) generalizes the analysis above and allows for many realistic elements. In particular, this general framework allows for arbitrary heterogeneity among agents, including heterogeneous preferences, beliefs, and returns. Moreover, the general framework allows for multidimensional effort in both the treatment and control group. This allows the model to accommodate, among other things: complex technologies, dynamic effort expenditure, and attempts by agents in the control group to obtain substitute treatments. Throughout, we maintain the criteria that any experimental design must be able to recover at least the data an experimenter would recover from a randomized controlled experiment. Thus, in many environments, selective trials may be implemented with little additional risk.

Moreover, our paper suggests a number of theoretically equivalent ways to implement selective trials. For example, the design above is equivalent to giving all subjects a single lottery ticket for a newspaper subscription, and allowing each subject the opportunity to purchase a second lottery ticket at a small cost. In addition, we show how our designs may be implemented using the BDM mechanism of (Becker et al. 1964). While a full description of this mechanism is beyond the scope of this article, we note that it has been extensively used in laboratory settings, and is gaining traction in the field (Berry et al. 2011).

Finally, while our paper is written largely with examples from development economics, our analysis contributes to the medical literature by proposing blind selective trials. These trials, where treatment status is hidden through the use of a placebo, allow the experimenter to tease apart the pure effect of treatment from the effect of changes in behavior due to the perception of treatment. As placebos are generally unavailable in social science experiments, we show how similar data can be elicited without a placebo by using incentives.

4 Concluding Thoughts

Selective trials are designed for situations where subjects can significantly affect experimental outcomes through unobserved actions. That is, they will be most useful when there is significant heterogeneity in treatment effects driven by subject behavior.

Moreover, in designing a particular selective trial, it is important to have a well-specified behavioral model of how subjects will choose which actions to engage in, and how experimental manipulations may affect those choices. While we use a behavioral model that relies heavily on rationality, there is no particular reason why other, well-specified, behavioral theories cannot be embedded in our framework. In particular, any experimental design that relies on the manipulation of a mediating variable could be considered in our framework, as the mediating variable is likely specified by a behavioral theory (Imai et al. (2011), indeed, Imai, Tingley, and Yamamoto, forthcoming, designs experiments using manipulation of mediators in a statistical, rather than decision theoretic, framework). Thus, to the extent that understanding how one would design a selective trial in a particular setting pushes an experimenter to clearly specify her behavioral model, considering a selective trial may be useful even if, in the end, an RCT is more appropriate or convenient.

If the above guidance seems vague, there is a reason for this: selective trials currently exist only in theory, and, to a limited extent, in the lab. However, many of the elements of selective trials, such as the randomized prices used in the example, have already been usefully exploited in field settings (Karlan and Zinman 2009, Cohen and Dupas 2010, Ashraf et al. 2010). Ultimately, if you are thinking about running an RCT and want to consider a selective trial, send us an email. We'll do our best to help.

References

- Ashraf, Nava, James Berry and Jesse M. Shapiro. 2010. "Can Higher Prices Stimulate Product Use? Evidence from a Field Experiment in Zambia." *American Economic Review* 7 100(6):2383-2413.
- Becker, Gordon M., Morris H. DeGroot and Jacob Marschak. 1964. "Measuring Utility by a Single-Response Sequential Method." *Behavioral Science* 9(3):226-232.
- Berry, James, Greg Fischer and Raymond Guiteras. 2011. "Eliciting and Utilizing Willingness to Pay: Evidence from Field Trials in Northern Ghana." London School of Economics, mimeo.
- Chassang, Sylvain, Gerard Padro i Miquel and Erik Snowberg. forthcoming. "Selective Trials: A Principal-Agent Approach to Randomized Controlled Experiments." *American Economic Review*.
- Cohen, Jessica and Pascaline Dupas. 2010. "Free Distribution or Cost-Sharing? Evidence from a Randomized Malaria Prevention Experiment." *Quarterly Journal of Economics* 125(1):1-45.
- Gerber, Alan S., Dean Karlan and Daniel Bergan. 2009. "Does the Media Matter? A Field Experiment Measuring the Effect of Newspapers on Voting Behavior and Political Opinions." *American Economic Journal: Applied Economics* 1(2):35-52.
- Imai, Kosuke, Dustin Tingley and Teppei Yamamoto. forthcoming. "Experimental Designs for Identifying Causal Mechanisms." *Journal of the Royal Statistical Society-Series A*.

Imai, Kosuke, Luke Keele, Dustin Tingley and Teppei Yamamoto. forthcoming. "Unpacking the Black Box: Learning about Causal Mechanisms from Experimental and Observational Studies." *American Political Science Review*.

Karlan, Dean S. and Jonathan Zinman. 2009. "Observing Unobservables: Identifying Information Asymmetries with a Consumer Credit Field Experiment." *Econometrica* 77(6):1993-2008.

Lab-in-the-field Experiments

Guy Grossman

Department of Political Science

Columbia University

gsg2102@columbia.edu

Political scientists are increasingly taking a class of experiments that is believed to measure social preferences, typically conducted in a laboratory environment, to various 'field' settings. However, whereas laboratory experiments (McDermott, 2002) and field experiments (de Rooij, Green and Gerber, 2009, Humphreys and Weinstein, 2009) have been the object of recent reviews, lab-in-the-field experiments (LITFEs) have received, thus far, relatively little theoretical coverage (*cf.* Morton and Williams (2010, 296-300)). In this brief review I consider some of the advantages and limitations of LITFEs, demonstrate their utility in one research area, and conclude by pointing to promising future research avenues.⁴

Morton and Williams (2010, 296) define LITFEs as experiments in which *subjects participate in a common physical location but the experiment brings, to some degree, the laboratory experiments to the subject's natural environments more than the subjects come to the laboratory*. This definition captures the main attraction of LITFEs: their ability to offer a balance between the control that a researcher has in the laboratory and the benefits of the natural setting of the field.

In most general terms, LITFEs can be categorized into two types. First are "Manipulation LITFEs", which measure subjects' behavioral differences under different manipulations, in situations where subjects or groups are randomly assigned to different treatments by the researcher. Here, as in any laboratory experiment, manipulations can take many forms, such as subjects' information set (Habyarimana et al., 2009), the size of the social return and a group's gender composition (Fearon, Humphreys and Weinstein, 2009), or subjects' level of participation in selecting a group leader (Grossman and Baldassarri, 2011).

Second are "Measurement LITFEs", where no manipulation takes place and subjects are merely randomized into decision-making roles (e.g., player A and B in dictator, trust, and ultimatum games). Outside political science, Measurement LITFEs have been mainly used to test whether theoretical models stand up to cultural variation. The Foundations of Human Sociality project (Henrich et al., 2004) being a paradigmatic example. In political science, Measurement LITFEs are increasingly being used when researchers are interested in juxtaposing the behavior of comparable populations in different locations that have been exogenously exposed to different conditions of interests. As in Henrich et al. (2004), the source of variation of interest in this class of studies is outside the control of the researcher.

Recent examples include Whitt and Wilson (2007) who use the dictator game as a measurement tool to understand legacies of discrimination in post-conflict Bosnia. Similarly, Gilligan, Pasquale and Samii (2011)

⁴See Cardenas and Carpenter (2008) for a recent review of substantive findings of behavioral experiments in development economics.

exploit exogenous variation in exposure to violence during Nepal's civil war to test the impact of violence exposure on post-war levels of social capital. Using a set of LITFEs and original survey data they find that members of communities with greater exposure to violence exhibit significantly greater levels of social capital, measured by subjects' willingness to invest in trust-based transactions and contribute to a collective good. Finally, in a recent study (Grossman, 2011), I use a plausibly exogenous variation in the rules for selecting leaders of community organizations in rural Uganda to test its impact on leaders' responsiveness. Using a dictator and third-party punishment experiments, I find that compared to appointed leaders, elected leaders are more likely to exhibit pro-social behavior towards group members as well as be more responsive to the signaled preferences of members of their group.

Pros and Cons of Lab-in-the-field Experiments

Given the high costs and logistical hurdles, why take the lab to the field? One major benefit of (variant) LITFEs is that they maintain the advantages of laboratory experiments, while addressing some of their weaknesses. Just as any laboratory experiment, LITFEs (a) allow to derive causal inferences, (b) enable to break down and investigate complex processes into smaller tractable units, (c) permit the experimenter to introduce a wider range of variation than is usually possible in field experiments, while better controlling the measurement of subjects and variables, and (d) more easily ensure that desired designs and treatments are administered consistently.⁵ By contrast, in field experiments researchers often lack control over the full context within which subjects make decisions.

In addition, LITFEs address some of the weaknesses of laboratory experiments. First, by moving away from western universities, LITFEs extend the scope conditions of experimental findings to new cultures and areas of the world. This is especially important given the growing evidence suggesting that in the context of behavioral experiments that measure social-preferences, subjects from developing countries behave differently than their counterparts in the developed world (Bahry and Wilson, 2006, Cardenas and Carpenter, 2008).

Second, though laboratory experiments have greatly contributed to the understanding of strategic behavior, they are poorly positioned to test theories in which group-level characteristics are consequential. For example, Burnham and Johnson (2005) question the use of laboratory experiments to inform the study of cooperation in social dilemmas. Since laboratory experiments strip context away (Henrich et al., 2004), they are limited in their ability to replicate the mutual trust, past experience, shared norms, and group identity that are central for balancing tension between private and group interests. Recent LITFEs in political science began addressing this limitation, by using subjects who are members of *pre-existing groups*: rural communities (Baldwin and Mvukiyeye, 2011, Gilligan, Pasquale and Samii, 2011), voluntary associations (Grossman and Baldassarri, 2011) and schools (Alexander and Christia, 2011).

Third, laboratory experiments tend to use convenient samples, partly because their main goal is to test general causal statements and not necessarily to determine the probability that a certain event will occur in a given population (Berkowitz and Donnerstein, 1982). However, convenient samples raise serious questions regarding the generalizability of laboratory experiments. By contrast, many LITFEs use random samples (Gilligan, Pasquale and Samii (2011), Grossman and Baldassarri (2011), Habyarimana et al. (2007), Whitt and Wilson (2007), to name but a few). Drawing a representative sample from a population of interest has several important benefits.

First, it allows the researcher to generalize the experimental results to larger populations of interests, i.e., define the target population more broadly. This is important, not least because of the growing evidence suggesting that the behavior of students — the quintessential laboratory experiments' subjects — might not be representative of

⁵See McDermott (2002, 3839) for a discussion of the advantages of laboratory experiments.

the larger society.⁶ Second, drawing representative samples, LITFEs are better suited to address the issue of heterogeneous treatment effects within the population. As Bahry and Wilson (2006) forcefully argue, LITFEs allow to test the extent to which using homogeneous samples masks key elements of strategic behavior. Third, when drawing a random sample, LITFEs can increase the accuracy of an important class of laboratory experiments where subjects' choices are interactive (e.g., bargaining, trust, third-party punishment and public goods games). This is because drawing a random sample allows subjects to form consistent beliefs about the behaviors of the individuals with whom they are interacting (see discussion in Habyarimana et al. (2007)).

Naturally, LITFEs are not a panacea for all ills. Specifically, LITFEs suffer from many of the limitations of laboratory experiments.⁷ First are concerns regarding the nature and extent of scrutiny of one's actions by others, which Bardsley (2009) refers to as the "contamination problem". In field experiments, interventions are (ideally) designed such that they will interact rather seamlessly with the subjects' everyday life experiences and the outcome measures are behavioral or institutional consequences of real world significance (Harrison and List, 2004). By contrast, in LITFEs subjects enter a more 'artificial' environment in which they are keenly aware that their behavior is being monitored, recorded and scrutinized, and where behavior is strongly influenced by instructions, induced role-playing and framing. Levitt and List (2007) argue that such scrutiny not only exaggerates the importance of pro-social behaviors relative to environments without such scrutiny, but also raises serious concerns of experimenter effects (Hawthorne effects).

The limited 'naturalness' of the context that LITFEs create also raise concerns of poor ecological validity. Are there important factors in the naturally occurring decision-making context that are excluded from the lab-in-the-field experiment? Is the degree of anonymity conferred upon experimental participants fitting? Are the chosen parameters reasonable? Is the information that subjects have appropriate? Are the stakes being used, fundamentally distort the relation between monetary and pro-social preferences? Does the induced strategic environment capture the theoretical constructs of interests? Are the measured variables equivalent to the variables the theory is concerned with? Moving from the laboratory to the field does not relieve researchers from addressing the central validity question — "how do the lab-in-the-field experimental results relate to 'real-world' behavior?"

LITFEs and Theories of Cooperation

In order to demonstrate their usefulness, in this section I discuss three recent LITFEs that contribute to the study of the relationship between political institutions and public goods contributions. In many developing countries, one outcome of the failure of the public sector to deliver effective social services has been the increasing reliance on local initiatives to provide public goods and social services. Ample evidence suggests that communities and various self-help groups exhibit variation in their ability to overcome the social dilemmas inherent in public goods provision. The extent to which local institutions can mitigate the tension between private and public interests occupies a central role in the political economy of development. As is well known, identifying the causal effect of institutions on political behavior — in our case, cooperation — is complicated by the fact that institutions are self-selected endogenously and deeply historically entrenched.

Alexander and Christia (2011) conducted N-person public goods experiments with costly sanctions in the ethnically divided city of Mostar in Bosnia-Herzegovina. The study's main goal is to examine how environmental context conditions affect the role ethnic diversity plays in public goods production. Specifically they take advantage of a natural experiment — which has randomly placed students in either ethnically integrated or ethnically segregated secondary schools — to test how institutions of integration affect cooperation within and

⁶See Henrich, Heine and Norenzayan (2010) for an exceptionally illuminating recent review.

⁷See Bardsley (2009), Levitt and List (2007) for insightful critiques.

across ethnic groups. This unique setup allowed the researchers to examine directly the effect of real-life institutions on the participants' willingness to contribute to the production of a public good. Obviously, a laboratory experiment in a western university is unable to mimic the war and post-war experiences that inform the cooperative behavior of the experimental subjects.

Alexander and Christia (2011) major experimental finding is that institutions of integration can mitigate the negative effect that ethnic diversity has on students' cooperative behavior. Importantly, they find that institutions of integration drive up public goods contributions by reinforcing the use of costly sanctions. In other words, social sanctioning is a mediator in the relationship between institutions of integration and cooperation in ethnically diverse groups. The authors stipulate that the presence of institutions of integration drives up contributions because it makes the threat of punishment more credible and, consequently, renders sanctions more effective. Interestingly, their LITFE results indicate that though sanctioning increases contributions in homogenous groups, it may be an ineffective tool for promoting cooperation in ethnically mixed groups, unless institutions of integration are present. By combining experimental and field methods, the authors make an important contribution to the existing debate on the emergence of cooperation in ethnically diverse societies.

Grossman and Baldassarri (2011) similarly use a series of N-person public goods experiments to investigate how governance institutions may mediate the relationship between sanctioning and cooperative behavior. The study's starting point is that most of our current theories of how sanctioning induces cooperation in social dilemmas are based on models of diffused and decentralized punishment. However, in practice, because peer sanctioning is only effective under restrictive conditions, groups and communities commonly develop forms of self-regulation, in which the power to sanction defectors is transferred to a local centralized authority. Our main interest is in examining whether the *political process* through which centralized authorities obtain their sanctioning powers causally impacts cooperative behavior. Specifically, we test whether elections have an independent positive effect on members' levels of cooperation.

To answer these question, we designed an adaptation of the public goods game in which sanctioning power was given to a single monitor, and the process by which the monitor is chosen was experimentally manipulated. In one treatment variant the monitor was randomly selected, whereas in the second treatment, game participants elected the monitor using a secret ballot. The experimental setup allows us to attest the impact of centralized-sanctioning institutions on cooperative behavior as well as to demonstrate that the size of this effect depends on the political process through which these institutions are established. Testing several possible causal mechanisms we find evidence supporting the hypothesis that allowing subjects to take part in the design of governing institutions changes their disutility from defection. In other words, we demonstrate that preferences are not fixed – they change from before and after the process of participation, even when the incentive structure itself has remained unchanged.

This LITFE was conducted in rural Uganda with a random sample of members of Ugandan farmer associations. Armed with survey-based data on the cooperative behavior of producers as members of the farmer cooperative, we were able to evaluate the LITFEs findings against comparable observational data. We find that farmers' cooperative behavior in the controlled setting predicts cooperative behavior in their natural environment, in which they face a similar social dilemma on a regular basis. These findings strengthen the validity of the experimental results.

Finally, in the context of a community driven reconstruction (CDR) project in northern Liberia, Fearon, Humphreys and Weinstein (2009) designed a public goods game with the goal of observing whether communities exposed to the CDR treatment behave differently from control communities after the project came to an end. CDR projects support the establishment of new local institutions in order to promote social reconciliation. Specifically the Liberia project attempted to build democratic, community-level institutions for making and implementing decisions about local public goods. Whether this form of assistance can bring about these desired effects is, however, an open and debatable question (Casey, Glennerster and Miguel, 2011). Using N-person

public goods LITFEs, the researchers set to examine whether a brief, foreign-funded efforts to build local institutions has in fact positive effects on local patterns of cooperation.

18 months into the CDR project, the participating communities (41 treatment and 41 control communities) were told that they would receive additional development funds, and that the specific amount received would depend on how much money a random sample of 24 people contributed to the project in a community-wide public goods game. In addition to the exogenous variation of the CDR treatment, the researchers ran cross-cutting experimental treatments varying the size of the social return as well as the gender composition of the sampled villagers. The researchers find that levels of cooperation to the public goods were significantly higher in communities that had participated in the CDR project. These findings suggest that “changes in community cohesion can take place over a short period of time; can occur in response to outside intervention; and can develop without fundamental changes either to the structure of economic relations or to more macro-level political processes” (pp. 291). Combining a field experiment with a LITFE, the paper makes an important contribution to a thorny debate about whether patterns of social cooperation are responsive to new institutions, even when underlying demographic, economic, and political factors remain unchanged.

Future Avenues of Research

I conclude this short review by pointing to three avenues for future work that can make important methodological contributions to the larger research community. The first avenue is related to the use of LITFEs in field experiments as an alternative measures to survey responses, when the outcome of interests are relatively abstract concepts such as social cohesion, trust, or fairness. In an oft-cited study, Glaeser et al. (2000) use a series of trust games to demonstrate that survey-based attitudinal measures of trust are weakly correlated with behavioral measures of trust. Since LITFEs that measure social preferences require subjects to make costly monetary decisions, such measures seem more reliable than survey responses that may be ‘cheap-talk’. In a recent study of a large community driven development project in Sierra Leone, Casey, Glennerster and Miguel (2011) use a series of innovative LITFEs to test the impact of new participatory institutions. Going beyond off-the-shelf protocols, designing new behavioral measures for theoretical construct of interests, and examining their correlation with widely used survey measures is a promising avenue for future work.

The extent to which findings from laboratory experiments in political science are simply a function of WEIRD experimental subjects — drawn from Western, Educated, Industrialized, Rich and Democratic countries — will surely be an important line of future research.⁸ Future studies should be designed to compare experimental findings from (a) Western industrialized societies with those from poor developing countries; (b) Western societies with those from non-Western industrialized societies; (c) Americans with people from other Western societies; (d) university educated subjects with nonuniversity-educated subjects, or university students with non-student adults in countries of all income level.

Addressing ecological validity concerns — the level of similarity between the controlled experiment and the natural environment, in which subjects are making similar real world decisions — is another promising avenue for future research. It is commonly argued by non-experimentalists that the data generating process in a controlled laboratory environment and in nature are so different that the artificiality of the experiment cannot be considered a valid analogue for real world decision-making processes. To address this conceptual concern, Guala (2002) and Harrison and List (2004) propose an omnibus test for ecological validity: using the same subjects to check whether data generated in laboratory conditions correlates in expected ways with data generated from nature. There are currently a small number of studies following this suggestion.

⁸See the recent debate in the special volume of *Behavioral and Brain Sciences* (2010), 61–135. The WEIRD acronym comes from Henrich, Heine and Norenzayan (2010).

For example, Karlan (2005) shows for borrowers in a microcredit program in Peru that behavior in a trust game predicts repay rates of subjects' loans one year later. He does not find, however, that pro-social behavior in a public goods game is correlated with repay probabilities. Carpenter and Seki (2004) find that the productivity of fishermen in Japan can be predicted by social preferences exhibited in a public good game. Benz and Meier (2008) found correlation between how individuals behave in donation experiments and how the same individuals behave in a naturally occurring decision situation on charitable giving. Similarly, as mentioned above, Grossman and Baldassarri (2011) find a correlation between the cooperative behavior of members of Ugandan farmer cooperatives in a public goods experiment and the behavior of the same subjects in a natural setting, in which they face a similar social dilemma. In contrast, Laury and Taylor (2006) find very weak correlation between contributions to a laboratory public goods experiment and voluntary contributions to a naturally occurring public good. Finally, Barr and Zeitlin (2010) find that primary teachers' allocations to parents in a dictator game in Uganda are positively but weakly correlated with their time allocations to teaching. Previous work therefore suggests that whether individuals' pro-social behavior in experiments correlates with their pro-social behavior in the field is still largely an open question.

- Alexander, Marcus and Fotini Christia. 2011. Context Modularity of Human Altruism. MIT Working Paper .
- Bahry, Donna L. and Rick K. Wilson. 2006. Confusion or fairness in the field? Rejections in the ultimatum game under the strategy method. *Journal of Economic Behavior & Organization* 60(1):3754.
- Baldwin, Kate and Eric Mvukiyehe. 2011. When Participation is Not Novel: the Effects of Participatory Clan Chief Selection in Liberia. Working Paper .
- Bardsley, Nicholas. 2009. *Experimental economics: rethinking the rules*. Princeton, NJ: Princeton University Press.
- Barr, Abigail and Andrew Zeitlin. 2010. Dictator games in the lab and in nature: External validity tested and investigated in Ugandan primary schools. Centre for the Study of African Economies Series Working Paper.
- Benz, Matthias and Stephan Meier. 2008. Do people behave in experiments as in the field? evidence from donations. *Experimental Economics* 11(3):268281281.
- Berkowitz, Leonard and Edward Donnerstein. 1982. External validity is more than skin deep: Some answers to criticisms of laboratory experiments. *American Psychologist* 37(3):245257.
- Burnham, Terence C. and Dominic Johnson. 2005. The biological and evolutionary logic of human cooperation. *Analyse & Kritik* 27(2):11335.
- Cardenas, Juan Camilo and Jeffrey Carpenter. 2008. Behavioural Development Economics: Lessons from Field Labs in the Developing World. *Journal of Development Studies* 44(3):311338.
- Casey, Katherine, Rachel Glennerster and Edward Miguel. 2011. Reshaping Institutions: Evidence on External Aid and Local Collective Action. Working Paper .
- de Rooij, Eline A., Donald P. Green and Alan S. Gerber. 2009. Field Experiments on Political Behavior and Collective Action. *Annual Review of Political Science* 12(1):389395.
- Fearon, James D., Macartan Humphreys and Jeremy M. Weinstein. 2009. Can Development Aid Contribute to Social Cohesion after Civil War? Evidence from a Field Experiment in Post-conflict Liberia. *American Economic Review* 99(2):287291.
- Gilligan, Michael J., Benjamin J. Pasquale and Cyrus D. Samii. 2011. Civil War and Social Capital: Behavioral-Game Evidence from Nepal. Working Paper .
- Glaeser, Edward L., David I. Laibson, Jose A. Scheinkman and Christine L. Soutter. 2000. Measuring Trust. *The Quarterly Journal of Economics* 115(3):811846.
- Grossman, Guy. 2011. The Impact of Leader Selection Rules on Monitoring Institutions and Leader Responsiveness Exploiting Exogenous Variation in the Formation of Ugandan Farmer Associations. Working Paper .
- Grossman, Guy and Delia Baldassarri. 2011. Centralized Sanctioning and Legitimate Authority in the Provision of Public Goods: Evidence from a Lab in Field Experiment In Uganda. Working Paper .

- Guala, Francesco. 2002. On the scope of experiments in economics: comments on Siakantaris. *Cambridge Journal of Economics* 26(2):261267.
- Habyarimana, James, Macartan Humphreys, Daniel N. Posner and Jeremy M. Weinstein. 2007. Why Does Ethnic Diversity Undermine Public Goods Provision? *The American Political Science Review* 101(4).
- Habyarimana, James, Macartan Humphreys, Daniel N. Posner and Jeremy M. Weinstein. 2009. *Coethnicity: diversity and the dilemmas of collective action*. New York: Russell Sage Foundation.
- Harrison, Glenn W. and John A. List. 2004. Field Experiments. *Journal of Economic Literature* 42(4):10091055.
- Henrich, Joseph, Robert Boyd, Samuel Bowles, Colin Camerer, Ernst Fehr and Herbert Gintis. 2004. *Foundations of Human Sociality: Economic Experiments and Ethnographic Evidence from Fifteen Small-Scale Societies*. Oxford University Press.
- Henrich, Joseph, Steven J. Heine and Ara Norenzayan. 2010. Beyond WEIRD: Towards a broadbased behavioral science. *Behavioral and Brain Sciences* 33(2-3):111135.
- Humphreys, Macartan and Jeremy M. Weinstein. 2009. Field Experiments and the Political Economy of Development. *Annual Review of Political Science* 12(1):367378.
- Karlan, Dean S. 2005. Using Experimental Economics to Measure Social Capital and Predict Financial Decisions. *The American Economic Review* 95(5):16881699.
- Levitt, Steven D. and John A. List. 2007. What Do Laboratory Experiments Measuring Social Preferences Reveal about the Real World? *The Journal of Economic Perspectives* 21(2):153174.
- McDermott, Rose. 2002. Experimental Methods in Political Science. *Annual Review of Political Science* 5(1):3161.
- Morton, Rebecca B. and Kenneth C Williams. 2010. *Experimental political science and the study of causality: from nature to the lab*. Cambridge: Cambridge University Press.
- Whitt, Sam and Rick K. Wilson. 2007. The Dictator Game, Fairness and Ethnicity in Postwar Bosnia. *American Journal of Political Science* 51(3):655668.

Section Mentor Program

Yanna Krupnikov

Department of Political Science, University of Indiana

ykrupnik@indiana.edu

During the 2011 APSA the Experiments Section ran a mentoring program which matched graduate students with faculty mentors. The goal of this program was to give graduate students a chance to meet faculty mentors with similar interests. In total, 30 graduate students and 29 faculty mentors participated in the program. Below, four of the graduate student participants discuss their experiences. A similar program is being planned for APSA 2012.

Jaime Settle, UCSD

I am most interested in how innate differences between people (genetic, physiological, and psychological differences) moderate the effects of geographic contexts and social interactions on political behavior. My dissertation examines 1) how exposure to political contention affects the way people respond emotionally to political stimuli and whether this consequently affects voter turnout, and 2) whether individual heterogeneity in sensitivity to threat moderates these relationships. As part of the mentoring program, I hoped to get feedback on an experimental design where I plan to manipulate the level of political contention in a laboratory environment and measure people's physiological response to it.

I met with Adam Seth Levine, an assistant professor at Cornell. We had a great discussion about our overall research agendas and the particular challenges we face doing experimental work. The experiments community is a small world, and we discovered that we had many friends and colleagues in common. Adam also had very helpful advice for me about going out on the job market this fall. The meeting was most useful in making the Experiments Section feel accessible to me as a young scholar. Adam's friendly demeanor and candor encouraged me to become more involved in the section and more engaged with the experiments community.

LaFleur Stephens, University of Michigan

Funded in part by a grant with Time Sharing Experiments for the Social Sciences (TESS), my dissertation examines various electoral strategies used by African Americans campaigning in non-majority-minority jurisdictions. The focus of my meeting with my APSA Experiments mentor, Lynn Vavreck, was an experiment that challenges earlier research on racial appeals. Previous scholarship has held that, because of the widespread adoption of the norm of equality, explicit racial appeals of the sort that were common prior to the 1960s are now rejected by the vast majority of Americans. In fact, my results suggest that respondent's acceptance of the norm of equality is contingent on the race of the messenger. However, since previous research on racial appeals was never conducted in the context of a black messenger, to date, this important caveat remained unknown.

Lynn's feedback was invaluable, as she provided advice on how to market my research to a broader audience beyond race and politics scholars. She also helped me to think about how this project fits into my overall research agenda. Finally, our meeting was short and informal, which is a low-stakes way for junior scholars to receive feedback.

Danielle Thomsen, Cornell University

I was very excited when Professor Laurel Harbridge was assigned to be my mentor at the APSA conference. My dissertation project explores the effect of partisan polarization on the gendered makeup of the U.S. Congress. I presented a related paper at APSA, which suggests that women candidates are perceived to be competent in masculine policy issues when they adopt hawkish positions. We discussed this paper and the experimental design at length during our meeting.

Professor Harbridge not only helped me to think about my paper in new ways, but she also gave me concrete suggestions for improving future iterations of the experiment. Beyond the session, we have since stayed in touch regarding my research, and her advice continues to guide the evolution of the paper into a hopeful publication. Even after the project runs its course, I look forward to having another mentor in the discipline that I can turn to for advice. I am grateful to Professor Harbridge for her willingness to participate in this program.

Aleks Ksiazkiewicz, Rice University

Given my interest in the intersection of biology and political behavior, I was quite excited when I found out that Doug Oxley was going to be my mentor. His work on the physiology of ideology and behavioral genetics is close to my own dissertation research on the heritability of political phenotypes. The program coordinator made a fantastic mentor-mentee fit in my case.

In our conversation, Doug shared his experiences on the job market as a young scholar in the field of biopolitics, provided advice on how to successfully conduct interdisciplinary research (including tips on developing successful collaborative relationships across disciplinary boundaries), and gave me invaluable feedback on my dissertation ideas. For example, we discussed empirical issues regarding the twin study methodology and theoretical concerns about how genetic populations of interest should be defined.

In all, I would strongly recommend other graduate students participate in the mentoring program. It helped me to develop a new contact, to get extensive feedback on my research ideas from someone outside my department, and to receive practical advice on positioning myself in a challenging job market. This meeting contributed to the quality of my research and provided an opportunity to begin a new collaborative relationship.

Section News and Announcements

- New Newsletter Section Starting

Starting in the next edition of the *Experimentalist*, we will have a new section on "Tricks of the Trade: Dos and Don'ts for Conducting Experiments". The closest thing to a guaranteed result in experiments is that the researcher will discover a new way that the experiment could be done better or avoid problems. Unfortunately, these lessons tend to accrue only to the individual researcher. Thus, many experimental researchers are reinventing the same practical lessons about how to run experiments. The purpose of "Tricks of the Trade" is to share these lessons for the benefit of the community. Think of it as the collective wisdom of the Section that we would like to pass along to colleagues and especially to a newcomer to conducting experimental research (e.g. graduate students or colleagues new to using the experimental method). Items in this section should be brief, practical, and - we hope - leavened with a sense of humor about the things that can go wrong when conducting experiments in political science. Christopher Mann of the University of Miami suggested the "Tricks of the Trade" section and will serve as its editor. Chris is looking for submissions about lab, survey, field, and other experiments. Entries can be submitted for attribution, anonymously, or using a pseudonym at <https://sites.google.com/site/christopherbmann/tricks-of-trade-submissions>.

- Existing Newsletter Sections Need Authors

I need future authors!! Book reviews (see Gailmard's of Norton/Williams, nominate your own book and reviewer too!), theme sections (e.g., the bio-politics and GOTV sections), co-author/inter-disciplinary experiences....Submit! Next issue is May 2012

- Experiment Section Journal: Feedback Needed

From Jamie Druckman:

Hi Members of the Experimental Section!

I hope everyone's fall is off to a good start. I apologize for the lengthy e-mail but PLEASE read it carefully as it discusses the LAUNCHING OF AN EXPERIMENTAL JOURNAL. I imagine such a journal would affect many if not all of us in terms of providing a novel publication venue.

We have been in the process of trying to adopt a journal for the section for about a year. The idea of a journal was initially considered, in detail, by a committee of John Geer (chair), Rose McDermott, Don Green, Rick Wilson, Scott Gartner, and Dan Posner. The committee decided it was a worthwhile endeavor. The next - and critically important step - is to have the section membership vote on adopting a journal (more on that in a moment).

In anticipation of the vote, a committee has written a set of journal governance rules that detail how the journal will work and relate to the section. The committee consisted of Jamie Druckman (chair), Chris Larimar, Macartan Humphreys, Don Green, Kevin Arceneaux, Rick Matland, Rose McDermott, Tom Palfry, and Rick Wilson.

These rules have now been posted (thanks to Kevin Esterling) and we are providing a one month period for discussion. Along these lines we have set up a forum where section members can post comments on the rules (and discuss them). Details on how to access the rules and comment appear at the bottom of this e-mail.

After approximately a month, we will hold an on-line vote. Members will be asked to vote against or in favor of the following:

Adoption of an official journal of the APSA experimental section. All section members would receive the journal. Section fees would not exceed \$35.

I will e-mail everyone when voting opens (probably around the start of November, such as the second week...). Members will have at least 30 days to cast a vote and I'll send many reminders; sorry in advance for that! We need at least 100 members to cast a vote for it to be valid. To move forward, a majority of voters must support the proposal.

If the journal is accepted, we will then continue the conversations we have begun with several publishers (many of whom are very interested in publishing the journal, as a regular, paper journal). We also will formally apply to the APSA for recognition of the journal.

Also, note, as discussed at the APSA business meetings, after a careful selection process, the inaugural editors of the journal would be Becky Morton and Josh Tucker. They will separately post their plans for operating the journal in the near future.

To access the rules and post comments, follow these instructions

1. Navigate to www.apsanet.org
2. Click the login box at the top-right and log in
3. Click on "go to apsa connect"
4. On the "my profile" menu, click on "my communities"
5. Click on "Section 42: Experimental Research"
6. In the box that says "latest shared files," click on "Final Journal Governance Rules" (i.e., you need to click on the title of the document itself, on the words)
7. In the section that says, "Attachment(s)," click on the document entitled "Final Journal Governance Rules.pdf."
8. After viewing the document, press the back arrow to return (JUST ONCE)
9. The next section below is labeled "Comments." In the comments box, press the link that says "+Add a comment" and then type in a comment and then press the save button.
10. If you like, you can also rate the document using the stars in the upper right hand side of the page.
11. When done, press log out.

APSA Experimental Section Officers

- President: Rebecca Morton (2012)
- President-elect: Alan Gerber (2012)
- At-large Council: Andre Blais (2012-13), Jennifer Jerit (2012-13), Jamie Druckman (2012-13), Don Green (2011-12)
- Treasurer: Jennifer Merolla (2012)
- Secretary: Rikhil Bhavnani (2012)
- Newsletter editor: Dustin Tingley (2010-12)
- Webmaster: Kevin Esterling (2010-12)
- APSA Program Chair: Ken Williams (2012)
- Nominating Committee (2012): Costas Panagopoulos (chair), John Bullock, Rick Matland, Rose McDermott
- Outreach Committee (2012): Yanna Krupnikov (chair), Adam Levine, Spencer Piston, Monica Schneider, Cheryl Boudreau
- Standards Committee: Alan Gerber (chair), Kevin Arceneaux, Tom Palfrey, Cheryl Boudreau, Conor Dowling, Sunshine Hillygus
- Journal Advisory Committee: John Geer (chair), Kevin Arceneaux, Macartan Humphreys, Tom Palfrey, and Lynn Vavreck.
- Best Paper at prior APSA: Daniel Rubenson (chair), Dan Myers, Chris Weber
- Best Dissertation in prior calendar year: Jason Reifler (chair), Elizabeth Suhay, Reuben Kline
- Best Book in prior calendar year: Melissa Michelson (chair), Adam Berinsky, Thad Dunning

We are pleased to announce awards and policy committees. Awards are made possible through joining the APSA section.

APSA Organized Section on Experimental Research Best Paper Award

We are now soliciting nominations for the Award for Best Paper Presented at the Previous Year's APSA featuring experimental analyses. The criteria for being nominated are simply that (a) the paper was presented at APSA and (b) that it features experimental analysis. Chairs and discussants are especially encouraged to nominate papers, but nominations from anyone who saw an interesting paper (as well as self-nominations) are welcome as well. Please email a copy of the paper in .pdf format along with a brief sentence or two about why you are nominating the paper to all four members of the committee at the following email addresses: Daniel Rubenson (chair) (rubenson@ryerson.ca), Dan Myers (myerscd@umich.edu), Chris Weber (crweber@lsu.edu). Please include "APSA Experiments Best Paper Nomination" in the subject heading of your email. Nominations will be accepted through May 1, 2012.

APSA Organized Section on Experimental Research Best Dissertation Award The APSA Organized Section on Experimental Research invites nominations for the 2011 Best Dissertation prize. Eligible nominees will have completed a dissertation in the 2011 calendar year that utilizes experimental methods on substantive political science research, or makes a fundamental contribution to experimental methods. Nominations should come from faculty members but they need not be on the student's dissertation committee. Nominations and copies of the completed dissertations are due April 1, 2012 and should be emailed to the selection committee: Jason Reifler (chair) (poljar@langate.gsu.edu), Elizabeth Suhay (suhaye@lafayette.edu), Reuben Kline (reubkline@gmail.com).

APSA Organized Section on Experimental Research Best Book Award For the best book published in 2011 that either uses or is about experimental research methods in the study of politics. A copy of the book should be sent to each member of the selection committee at the addresses provided below no later than April 1, 2012.

Melissa R. Michelson (chair) Menlo College 1000 El Camino Real Atherton, CA 94027-4301

Adam Berinsky Department of Political Science Massachusetts Institute of Technology 77 Massachusetts Ave., E53-457 Cambridge, MA 02139

Thad Dunning FOR POSTAL DELIVERIES: Dept. of Political Science Yale University P.O. Box 208301 New Haven, CT 06520-8301

FOR FED EX DELIVERIES: 115 Prospect St. Rosenkranz Hall, Room 420 New Haven, CT 06511 Phone: 203-432-5230

- **Call for book proposals**

Call for Proposals: Routledge Studies on Experimental Political Science www.routledge.com/politics

Series Editors: Kevin T. Arceneaux, Temple University and Cindy D. Kam, Vanderbilt University

Advisory Board: Howard Lavine, University of Minnesota; Joshua Tucker, New York University; Rick Wilson, Rice University; and Elizabeth Zechmeister, Vanderbilt University

Political scientists are increasingly using experiments to study important political and social phenomena. The logic of experimentation makes it an appealing and powerful methodological tool that enables scholars to

establish causality and probe into the mechanisms underlying observable regularities. Experiments, because of their transparency, also enable researchers to communicate their findings to a broad audience. Although highly technical knowledge is not necessary for understanding the gist of experiments, experiments must be designed, administered, and analyzed with care and attention to detail.

The Routledge Studies on Experimental Political Science was developed to publish books that educate readers about the appropriate design and interpretation of experiments and books that showcase innovative and important applications of experimental work. We are particularly interested in scholarly monographs, but proposals for edited volumes will also be considered.

The series will showcase experimental work in political science in at least two ways:

Pedagogy: Books that provide pedagogical guidance on the design, administration, and analysis of experiments, particularly tailored for a political science audience. Such books would be targeted at advanced undergraduates, graduate students, and faculty members.

Applications: Books that use experimental methods, particularly innovative experimental methods, to understand important causal relationships in political and social settings.

If you have a book proposal or idea in mind which might be suitable for the series, please do not hesitate to contact:

Series Editors Kevin Arceneau: arceneau@temple.edu Cindy Kam: cindy.d.kam@vanderbilt.edu

Acquisitions Editor Natalja Mortensen, Political Science Research: natalja.mortensen@taylorandfrancis.com

Upcoming Events

- **Fifth Annual NYU-CESSE Conference on Experimental Political Science call for papers**

Please mark your calendars for next year's conference, which will take place on March 2-3, 2012. Joshua Tucker and Eric Dickson will be co-directing the conference, and we expect to post a call for papers early in the fall. Thanks to the generosity of NYU's Department of Politics and the NYU's Center for Experimental Social Science we continue to be able to offer the conference without a registration fee and to provide free food throughout the weekend, so we hope to see as many of you there as possible!

Further details about the call for papers can be found here:

<http://themonkeycage.org/blog/2011/09/12/5th-annual-nyu-cess-experiments-in-political-science-call-for-papers/>

Paper proposals will be accepted until November 15th, 2011.

- **WCE.2012, UC Berkeley**

The Fifth Annual West Coast Experiments Conference will be held at UC Berkeley, on May 11, 2012. This will be an all day conference, and will bring together researchers interested in advances in experimental methods for political science. This years co-organizers are Jas Sekhon, Laura Stoker, Sean Gailmard, Mat McCubbins and

Kevin Esterling. There are no registration fees for the west coast conference, and all meals for the day will be provided to all registered attendees. Space will be limited. Nominations for paper presentations can be sent to kevin.esterling@ucr.edu. More details to come.

[Back to Contents](#)