



Q&A with Jake Bowers, co-author of 2014 Miller Prize Paper



Political Analysis

Political Analysis (<http://pan.oxfordjournals.org/>) chronicles the exciting developments in the field of political methodology, with contributions to empirical and methodological scholarship outside the diffuse borders of political science. It is published on behalf of [The Society for Political Methodology \(http://polmeth.wustl.edu/\)](http://polmeth.wustl.edu/). [Political Analysis](#)

(<http://pan.oxfordjournals.org/>) is ranked #5 out of 157 journals in Political Science by 5-year impact factor, according to the 2012 ISI Journal Citation Reports. Like Political Analysis on Facebook and follow [@PolAnalysis](https://twitter.com/polanalysis) (<https://twitter.com/polanalysis>) on Twitter.

[READ MORE](#)

(HTTF

BY R. MICHAEL ALVAREZ ([HTTPS://BLOG.OUP.COM/AUTHORS/R-MICHAEL-ALVAREZ/](https://blog.oup.com/authors/r-michael-alvarez/))

SEPTEMBER 25TH 2014

Despite what many of my colleagues think, being a journal editor is usually a pretty interesting job. The best part about being a journal editor is working with authors to help frame, shape, and improve their research. We also have many chances to honor specific authors and their work for being of particular importance. One of those honors is the Miller Prize (<http://polmeth.wustl.edu/about.php?page=awards>), awarded annually by the Society for Political Methodology (<http://polmeth.wustl.edu/>) for the best paper published in *Political Analysis* (<http://pan.oxfordjournals.org>) the proceeding year.

The 2014 Miller Prize was awarded to [Jake Bowers](http://www.jakebowers.org/) (<http://www.jakebowers.org/>), [Mark M. Fredrickson](http://www.markfredrickson.com/) (<http://www.markfredrickson.com/>), and [Costas Panagopoulos](http://www.costaspanagopoulos.com/) (<http://www.costaspanagopoulos.com/>), for their paper, “Reasoning about Interference Between Units: A General Framework.” (<http://oxford.ly/1rM85Jk>) To recognize the significance of this paper, it is available for free online access for the next year. The award committee summarized the contribution of the paper:

“..the article tackles an difficult and pervasive problem—interference among units—in a novel and compelling way. Rather than treating spillover effects as a nuisance to be marginalized over or, worse, ignored, Bowers et al. use them as an opportunity to test substantive questions regarding interference ... Their work also brings together causal inference and network analysis in an innovative and compelling way, pointing the way to future convergence between these domains.”

In other words, this is an important contribution to political methodology.

I recently posed a number of question to one of the authors of the 2014 Miller Prize paper, Jake Bowers, asking him to talk more about this paper and its origins.

R. Michael Alvarez: Your paper, “Reasoning about Interference Between Units: A General Framework” (<http://oxford.ly/1rM85Jk>) recently won the 2014 Miller Prize for the best paper published in Political Analysis in the past year. What motivated you to write this paper?

Jake Bowers: Let me provide a little background for readers not already familiar with randomization-based statistical inference.

Randomized designs provide clear answers to two of the most common questions that we ask about empirical research: *The Interpretation Question*: “What does it mean that people in group A act differently from people in group B?” and *The Information Question*: “How precise is our summary of A-vs-B?” (Or, more defensively, “Do we really have enough information to distinguish A from B?”).

If we have randomly assigned some A-vs-B intervention, then we can answer the interpretation question very simply: “If group A differs from group B, it is only because of the A-vs-B intervention. Randomization ought to erase any other pre-existing differences between groups A and B.”

In answering the information question, randomization alone also allows us to characterize other ways that the experiment might have turned out: “Here are all of the possible ways that groups A and B could differ if we re-randomized the A-vs-B intervention to the experimental pool while entertaining the idea that A and B do not differ. If few (or none) of these differences is as large as the one we observe, we have a lot of information against the idea that A and B do not differ. If many of these differences are as large as the one we see, we don’t have much information to counter the argument that A and B do not differ.”

Of course, these are not the only questions one should ask about research, and interpretation should not end with knowing that an input created an output. Yet, these concerns about meaning and information are fundamental and the answers allowed by randomization offer a particularly clear starting place for learning from observation. In fact, many randomization-based methods for summarizing answers to the information question tend to have validity guarantees even with small samples. If we really did repeat the experiment all the possible ways that it could have been done, *and* repeated a common hypothesis test many times, we would reject a true null hypothesis no more than $\alpha\%$ of the time even if we had observed only eight people (Rosenbaum 2002, Chap 2).

In fact a project with only eight cities impelled this paper. [Costa Panagopoulos \(http://blog.oup.com/2013/12/q-a-with-costas-panagopoulos-about-2013-election-cycle/\)](http://blog.oup.com/2013/12/q-a-with-costas-panagopoulos-about-2013-election-cycle/) had administered a field experiment of newspaper advertising and turnout to eight US cities, and he and I began to discuss how to produce substantively meaningful, easy to interpret, and statistically valid, answers to the question about the effect of advertising on turnout. Could we hypothesize that, for example, the effect was zero for three of the treated cities, and more than zero for one of the treated cities? The answer was yes.

I realized that hypotheses about causal effects do not need to be simple, and, furthermore, they could represent substantive, theoretical models very directly. Soon, Mark Fredrickson and I started thinking about substantive models in which treatment given to one city might have an effect on another city. It seemed straightforward to write down these models. We had read Peter Aronow’s and Paul Rosenbaum’s papers on the sharp null model of no effects and interference, and so we didn’t think we were completely off base to imagine that, if we side-stepped estimation of average treatment effects and focused on testing hypotheses, we could learn something about what we called “models of interference”. But, we had not seen this done before. So, in part because we worried about whether we were right about how simple it was to write down and test hypotheses generated from models of spillover or interference between units, we wrote the “Reasoning about Interference” paper to see if what we were doing with Panagopoulos’ eight cities would scale, and whether it would perform as randomization-based tests should. The paper shows that we were right.

R. Michael Alvarez: In your paper, you focus on the “no interference” assumption that is widely discussed in the contemporary literature on causal models. What is this assumption and why is it important?

Jake Bowers: When we say that some intervention, (Z_i) , caused some outcome for some person, (i) , we often mean that the outcome we would have seen for person (i) when the intervention is not-active, $(Z_{i=0})$ — written as $(y_{\{i,Z_{i=0}\}})$ — would have been different from the outcome we would have seen if the intervention were active for that same person (at that same moment in time), $(Z_{i=1})$, — written as $(y_{\{i,Z_{i=1}\}})$. Most people would say that the treatment had an effect on person (i) when (i) would have acted differently under the intervention than under the control condition such that $y_{\{i,Z_{i=1}\}}$ does not equal $y_{\{i,Z_{i=0}\}}$. David Cox (1958) noticed that this definition of causal effects involves an assumption that an intervention assigned to one person does not influence the potential outcomes for another person. (Henry Brady’s piece, “Causation and Explanation in Social Science” in the *Oxford Handbook of Political Methodology* provides an excellent discussion of the no-interference assumption and Don Rubin’s formalization and generalization of Cox’s no-interference assumption.)

As an illustration of the confusion that interference can cause, imagine we had four people in our study such that (i in $\{1,2,3,4\}$). When we write that the intervention had an effect for person ($i=1$), ($y_{\{i=1, Z_1=1\}}$ does not equal $y_{\{i=1, Z_1=0\}}$), we are saying that person 1 would act the same when ($Z_{\{i=1\}}=1$) regardless of how intervention was assigned to any other person such that

$$(Y_{\{i=1, \{Z_1=1, Z_2=1, Z_3=0, Z_4=0\}\}} = Y_{\{i=1, \{Z_1=1, Z_2=0, Z_3=1, Z_4=0\}\}} = Y_{\{i=1, \{Z_{i=1}, \dots\}\}})$$

If we do not make this assumption then we cannot write down a treatment effect in terms of a simple comparison of two groups. Even if we randomly assigned the intervention to two of the four people in this little study, we would have six potential outcomes per person rather than only two potential outcomes (you can see two of the six potential outcomes for person 1 in above). Randomization does not help us decide what a “treatment effect” means and six counterfactuals per person poses a challenge for the conceptualization of causal effects.

So, interference is a problem with the definition of causal effects. It is also a problem with estimation. Many folks know about what Paul Holland (1986) calls the “Fundamental Problem of Causal Inference” that the potential outcomes heuristic for thinking about causality reveals: we cannot ever know the causal effect for person (i) directly because we can never observe both potential outcomes. I know of three main solutions for this problem, each of which have to deal with problems of interference:

- Jerzy Neyman (1923) showed that if we change our substantive focus from individual level to group level comparisons, and to averages in particular, then randomization would allow us to learn about the true, underlying, average treatment effect using the difference of means observed in the actual study (where we only see responses to intervention for some but not all of the experimental subjects).
- Don Rubin (1978) showed a Bayesian predictive approach — a probability model of the outcomes of your study and a probability model for the treatment effect itself allows you can predict the unobserved potential outcomes for each person in your study and then take averages of those predictions to produce an estimate of the average treatment effect.
- Ronald Fisher (1935) suggested another approach which maintained attention on the individual level potential outcomes, but did not use models to predict them. He showed that randomization alone allows you to test the hypothesis of “no effects” at the individual level. Interference makes it difficult to interpret Neyman’s comparisons of observed averages and Rubin’s comparison of predicted averages as telling us about causal effects because we have too many possible averages.

It turns out that Fisher’s sharp null hypothesis test of no effects is simple to interpret even when we have unknown interference between units. Our paper starts from that idea and shows that, in fact, one can test sharp hypotheses about effects rather than only no effects.

Note that there has been a lot of great recent work trying to define and estimate average treatment effects recently by folks like Cyrus Samii and Peter Aronow, Neelan Sircar and Alex Coppock, Panos Toulis and Edward Kao, Tyler Vanderweele, Eric Tchetgen Tchetgen and Betsy Ogburn, Michael Sobel, and Michael Hudgens, among others. I also think that interference poses a smaller problem for Rubin’s approach in principle — one would add a model of interference to the list of models (of outcomes, of intervention, of effects) used to predict the unobserved outcomes. (This approach has been used without formalization in terms of counterfactuals in both the spatial and networks models worlds.) One might then focus on posterior distributions of quantities other than simple differences of averages or interpret such differences reflecting the kinds of weightings used in the work that I gestured to at the start of this paragraph.

R. Michael Alvarez: How do you relax the “no interference” assumption in your paper?

Jake Bowers: I would say that we did not really relax an assumption, but rather side-stepped the need to think of interference as an assumption. Since we did not use the average causal effect, we were not facing the same problems of requiring that all potential outcomes collapse down to two averages. However, what we had to do instead was use what Paul Rosenbaum might call Fisher's solution to the fundamental problem of causal inference. Fisher noticed that, even if you couldn't say that a treatment had an effect on person (i), you could ask whether we had enough information (in our design and data) to shed light on a question about whether or not the treatment had an effect on person (i). In our paper, Fisher's approach meant that we did not need to define our scientifically interesting quantity in terms of averages. Instead, we had to write down hypotheses about no interference. That is, we did not really relax an assumption, but instead we directly modelled a process.

Rosenbaum (2007) and Aronow (2011), among others, had noticed that the hypothesis that Fisher is most famous for, the sharp null hypothesis of no effects, in fact does not assume no interference, but rather implies no interference (i.e., if the treatment has no effect for any person, then it does not matter how treatment has been assigned). So, in fact, the assumption of no interference is not really a fundamental piece of how we talk about counterfactual causality, but a by-product of a commitment to the use of a particular technology (simple comparisons of averages). We took a next step in our paper and realized that Fisher's sharp null hypothesis implied a particular, and very simple, model of interference (a model of no interference). We then set out to see if we could write other, more substantively interesting models of interference. So, that is what we show in the paper: one can write down a substantive theoretical model of interference (and of the mechanism for an experimental effect to come to matter for the units in the study) and then this model can be understood as a generator of sharp null hypotheses, each of which could be tested using the same randomization inference tools that we have been studying for their clarity and validity previously.

R. Michael Alvarez: What are the applications for the approach you develop in your paper?

Jake Bowers: We are working on a couple of applications. In general, our approach is useful as a way to learn about substantive models of the mechanisms for the effects of experimental treatments.

For example, [Bruce Desmarais](http://people.umass.edu/bruced/) (<http://people.umass.edu/bruced/>), Mark Fredrickson, and I are working with [Nahomi Ichino](http://sites.lsa.umich.edu/nichino/) (<http://sites.lsa.umich.edu/nichino/>), [Wayne Lee](http://math.unc.edu/people/graduate-students/grad-students) (<http://math.unc.edu/people/graduate-students/grad-students>), and [Simi Wang](https://sites.google.com/site/wangsimi88/) (<https://sites.google.com/site/wangsimi88/>) on how to design randomized experiments to learn about models of the propagation of treatments across a social network. If we think that an experimental intervention on some subset of Facebook users should spread in some certain manner, then we are hoping to have a general way to think about how to design that experiment (using our approach to learn about that propagation model, but also using some of the new developments in network-weighted average treatment effects that I referenced above). Our very early work suggests that, if treatment does propagate across a social network following a common infectious disease model, that you might prefer to assign relatively few units to direct intervention.

In another application, Nahomi Ichino, Mark Fredrickson, and I are using this approach to learn about agent-based models of the interaction of ethnicity and party strategies of voter registration fraud using a field experiment in Ghana. To improve our formal models, another collaborator, [Chris Grady](http://www.pol.illinois.edu/people/cdgrady2) (<http://www.pol.illinois.edu/people/cdgrady2>), is going to Ghana to do in-depth interviews with local party activists this fall.

R. Michael Alvarez: Political methodologists have made many contributions to the area of causal inference. If you had to recommend to a graduate student two or three things in this area that they might consider working on in the next year, what would they be?

Jake Bowers: About advice for graduate students: Here are some of the questions I would love to learn about.

- How should we move from formal, equilibrium-oriented, theories of behavior to models of mechanisms of treatment effects that would allow us to test hypotheses and learn about theory from data?
- How can we take advantage of estimation-based procedures or procedures developed without specific focus on counterfactual causal inference if we want to make counterfactual causal inferences about models of interference? How should we reinterpret or use tools from spatial analysis like those developed by Rob Franzese and Jude Hayes or tools from network analysis like those developed by Mark Handcock to answer causal inference questions?
- How can we provide general advice about how to choose test-statistics to summarize the observable implications of these theoretical models? We know that the KS-test used in our article is pretty low-powered. And we know from Rosenbaum (Chap 2, 2002) that certain classes of test statistics have excellent properties in one-dimension, but I wonder about general properties of multi-parameter models and test statistics that can be sensitive to multi-way differences in distribution between experimental groups.
- How should we apply ideas from randomized studies to the observational world? What does adjustment for confounding/omitted variable bias (by matching or “controlling for” or weighting) mean in the context of social networks or spatial relations? How should we do and judge such adjustment? Would what Rosenbaum-inspired sensitivity analysis or Manski-inspired bounds analysis might mean when we move away from testing one parameter or estimating one quantity?

R. Michael Alvarez: You do a lot of work with software tool development and statistical computing. What are you working on now that you are most excited about?

Jake Bowers: I am working on two computationally oriented projects that I find very exciting. The first involves using machine learning/statistical learning for optimal covariance adjustment in experiments (with Mark Fredrickson and [Ben Hansen](http://www.psc.isr.umich.edu/people/profile/703/Ben_Hansen) (http://www.psc.isr.umich.edu/people/profile/703/Ben_Hansen)). The second involves collecting thousands of hand-drawn maps on Google maps as GIS objects to learn about how people define and understand the places where they live in Canada, the United Kingdom, and the United States (with [Cara Wong](http://www.pol.illinois.edu/people/carawong) (<http://www.pol.illinois.edu/people/carawong>), [Daniel Rubenson](http://www.politics.ryerson.ca/rubenson/) (<http://www.politics.ryerson.ca/rubenson/>), Mark Fredrickson, [Ashlea Rundlett](http://www.pol.illinois.edu/people/rundlet2) (<http://www.pol.illinois.edu/people/rundlet2>), [Jane Green](http://www.manchester.ac.uk/research/jane.green/) (<http://www.manchester.ac.uk/research/jane.green/>), and [Edward Fieldhouse](http://www.manchester.ac.uk/research/ed.fieldhouse/) (<http://www.manchester.ac.uk/research/ed.fieldhouse/>)).

When an experimental intervention has produced a difference in outcomes, comparisons of treated to control outcomes can sometimes fail to detect this effect, in part, because the outcomes themselves are naturally noisy in comparison to the strength of the treatment effect. We would like to reduce the noise that is unrelated to treatment (say, remove the noise related to background covariates, like education) without ever estimating a treatment effect (or testing a hypothesis about a treatment effect). So far, people shy away from using covariates for precision enhancement of this type because of every model in which they soak up noise with covariates is also a model in which they look at the p-value for their treatment effect. This project learns from the growing literature in machine learning (aka statistical learning) to turn specification of the covariance adjustment part of a statistical model over to an automated system focused on the control group only which thus bypasses concerns about data snooping and multiple p-values.

The second project involves using Google maps embedded in online surveys to capture hand-drawn maps representing how people respond when asked to draw the boundaries of their “local communities.” So far we have over 7000 such maps from a large survey of Canadians, and we plan to have data from this module carried on the British Election Study and the US

Cooperative Congressional Election Study within the next year. We are using these maps and associated data to add to the “context/neighborhood effects” literature to learn how psychological understandings of place by individuals relates to Census measurements and also to individual level attitudes about inter-group relations and public goods provision.

Headline image credit: Abstract city and statistics. [CC0 via Pixabay \(http://pixabay.com/en/business-finance-statistics-graphic-361126/\)](http://pixabay.com/en/business-finance-statistics-graphic-361126/).

Jake Bowers (<http://www.jakebowers.org/>) is an Associate Professor in Political Science and Statistics at the University of Illinois at Urbana-Champaign. He began to work on randomization inference and causal inference while he was an assistant professor in the political science department at the University of Michigan, and while a post-doctoral scholar with the Robert Wood Johnson Health Policy Scholars Program at Harvard University. His interest in research designs with interference and dependence grew from his experiences studying multilevel models and time series processes while pursuing a Ph.D at the University of California, Berkeley.

R. Michael Alvarez is a professor of Political Science at Caltech. His research and teaching focuses on elections, voting behavior, and election technologies. He is editor-in-chief of *Political Analysis* (<http://pan.oxfordjournals.org/>) with Jonathan N. Katz.

POSTED IN: [JOURNALS \(HTTPS://BLOG.OUP.COM/CATEGORY/SUBTOPICS/OXFORDJOURNALS/\)](https://blog.oup.com/category/subtopics/oxfordjournals/) / [POLITICS \(HTTPS://BLOG.OUP.COM/CATEGORY/SOCIAL_SCIENCES/POLITICS-POLITICAL-SCIENCE/\)](https://blog.oup.com/category/social_sciences/politics-political-science/)

Subscribe to the OUPblog via email:

Or subscribe to articles in the subject area by [email \(http://blog.oup.com/subscribe-oupblog-via-email/\)](http://blog.oup.com/subscribe-oupblog-via-email/) or [RSS \(http://blog.oup.com/follow/\)](http://blog.oup.com/follow/)

Recent Comments

There are currently no comments.

Leave a Comment

Your email address will not be published. Required fields are marked *

COMMENT

NAME *

EMAIL *

WEBSITE

POST COMMENT

- [About Us \(https://blog.oup.com/about/\)](https://blog.oup.com/about/)
- [Follow Us \(https://blog.oup.com/follow/\)](https://blog.oup.com/follow/)
- [Contact Us \(https://blog.oup.com/contact/\)](https://blog.oup.com/contact/)
- [Terms and Conditions of Use \(https://blog.oup.com/terms_and_condi/\)](https://blog.oup.com/terms_and_condi/)
- [Reuse of Blog Content \(https://blog.oup.com/reuse-of-oupblog-content/\)](https://blog.oup.com/reuse-of-oupblog-content/)
- [Cookie Policy \(http://global.oup.com/cookiepolicy/\)](http://global.oup.com/cookiepolicy/)
- [Privacy Policy \(https://blog.oup.com/privacy_policy/\)](https://blog.oup.com/privacy_policy/)

Categories

- [Arts & Humanities \(https://blog.oup.com/category/arts_and_humanities/\)](https://blog.oup.com/category/arts_and_humanities/)
- [History \(https://blog.oup.com/category/history/\)](https://blog.oup.com/category/history/)
- [Language \(https://blog.oup.com/category/language-words/\)](https://blog.oup.com/category/language-words/)
- [Law \(https://blog.oup.com/category/law-legal/\)](https://blog.oup.com/category/law-legal/)
- [Science & Medicine \(https://blog.oup.com/category/science_medicine/\)](https://blog.oup.com/category/science_medicine/)
- [Social Sciences \(https://blog.oup.com/category/social_sciences/\)](https://blog.oup.com/category/social_sciences/)
- [Multimedia \(https://blog.oup.com/category/multimedia/\)](https://blog.oup.com/category/multimedia/)
- [Series & Columns \(https://blog.oup.com/category/series-columns/\)](https://blog.oup.com/category/series-columns/)

Social Media

- [Facebook \(https://www.facebook.com/OUPAcademic/\)](https://www.facebook.com/OUPAcademic/)
- [Twitter \(https://twitter.com/OUPAcademic/\)](https://twitter.com/OUPAcademic/)
- [Google Plus \(https://plus.google.com/+OUPAcademic/\)](https://plus.google.com/+OUPAcademic/)
- [Tumblr \(http://oupacademic.tumblr.com/\)](http://oupacademic.tumblr.com/)
- [Pinterest \(http://www.pinterest.com/oupacademic/\)](http://www.pinterest.com/oupacademic/)
- [YouTube \(https://www.youtube.com/OUPAcademic/\)](https://www.youtube.com/OUPAcademic/)
- [Vimeo \(http://vimeo.com/oupacademic/\)](http://vimeo.com/oupacademic/)
- [Soundcloud \(https://soundcloud.com/oupacademic/\)](https://soundcloud.com/oupacademic/)
- [Sina Weibo \(http://weibo.com/oupacademic/\)](http://weibo.com/oupacademic/)
- [OxfordWords blog \(http://blog.oxforddictionaries.com/\)](http://blog.oxforddictionaries.com/)

Gateways

- [Oxford Academic \(http://global.oup.com/academic/\)](http://global.oup.com/academic/)
- [Oxford Journals \(http://www.oxfordjournals.org/\)](http://www.oxfordjournals.org/)
- [Oxford Online \(http://global.oup.com/online/\)](http://global.oup.com/online/)
- [Higher Education \(http://global.oup.com/ushe/\)](http://global.oup.com/ushe/)
- [Oxford Index \(http://oxfordindex.oup.com/\)](http://oxfordindex.oup.com/)
- [Oxford Dictionaries \(http://www.oxforddictionaries.com/\)](http://www.oxforddictionaries.com/)
- [OUP Worldwide \(http://global.oup.com/\)](http://global.oup.com/)

- [University of Oxford \(http://www.ox.ac.uk/\)](http://www.ox.ac.uk/)

Oxford University Press is a department of the University of Oxford. It furthers the University's objective of excellence in research, scholarship, and education by publishing worldwide.

OXFORD
UNIVERSITY PRESS

Copyright © Oxford University Press 2017