sampling and to the validity and reliability of inferences. It therefore conveys the comprehensiveness of the inquiry and the confidence level related to multiple aspects of the interview component of any given research project.

We recognize that legitimate constraints imposed by Institutional Review Boards, by informants themselves, or by professional ethics may force researchers to keep some details of the interview confidential and anonymous. In certain cases, the Interview Methods Table might contain "confidentiality requested" and "confidentiality required" for every single interview. We do not seek to change prevailing practices that serve to protect informants. However, we believe that even in such circumstances, the interviewer can safely report many elements in an Interview Methods Appendix and Interview Methods Table—to the benefit of researcher and reader alike. A consistent set of expectations for reporting will give readers more confidence in research based on interview data, which in turn will liberate researchers to employ this methodology more often and with more rigor.

References

Berry, Jeffrey M. 2002. "Validity and Reliability Issues In Elite Interviewing." *PS: Political Science & Politics* vol. 35, no. 4: 679–682.

Bleich, Erik, and Robert Pekkanen. 2013. "How to Report Interview Data: The Interview Methods Appendix." In *Interview Research in Political Science*, edited by Layna Mosley. Ithaca: Cornell University Press, 84–105.

Brooks, Sarah M. 2013. "The Ethical Treatment of Human Subjects and the Institutional Review Board Process." In *Interview Research in Political Science*, edited by Layna Mosley. Ithaca: Cornell University Press, 45–66.

Elman, Colin, Diana Kapiszewski, and Lorena Vinuela. 2010. "Qualitative Data Archiving: Rewards and Challenges." *PS: Political Science & Politics* vol. 43, no. 1: 23–27.

Gerring, John. 2007. *Case Study Research: Principles and Practice*. 2nd ed. Cambridge: Cambridge University Press.

Guest, Greg, Arwen Bunce, and Laura Johnson. 2006. "How Many Interviews Are Enough? An Experiment with Data Saturation and Variability." *Field Methods* vol. 18, no. 1: 59–82.

Lieberman, Evan S. 2005. "Nested Analysis as a Mixed-Method Strategy for Comparative Research." *American Political Science Review* vol. 99, no 3: 435–452.

Lynch, Julia F. 2013. "Aligning Sampling Strategies with Analytical Goals." In *Interview Research in Political Science*, edited by Layna Mosley. Ithaca: Cornell University Press, 31–44.

MacLean, Lauren M. "The Power of the Interviewer." In *Interview Research in Political Science*, edited by Layna Mosley. Ithaca: Cornell University Press, 67–83.

Moravcsik, Andrew. 2010. "Active Citation: A Precondition for Replicable Qualitative Research." *PS: Political Science & Politics* vol. 43, no. 1: 29–35.

Mosley, Layna, ed. 2013. *Interview Research in Political Science*. Ithaca: Cornell University Press.

Mosley, Layna. 2013a. "Introduction." In *Interview Research in Political Science*, edited by Layna Mosley. Ithaca: Cornell University Press: 1–28.

Parry, Odette, and Natasha S. Mauthner. 2004. "Whose Data Are They Anyway?" *Sociology* vol. 38, no. 1: 132–152.

Transparency in Practice: Using Written Sources

Marc Trachtenberg
University of California, Los Angeles

Individual researchers, according to the revised set of guidelines adopted by the American Political Science Association two years ago, "have an ethical obligation to facilitate the evaluation of their evidence-based knowledge claims" not just by providing access to their data, but also by explaining how they assembled that data and how they drew "analytic conclusions" from it.[1] The assumption was that research transparency is of fundamental importance for the discipline as a whole, and that by holding the bar higher in this area, the rigor and richness of scholarly work in political science could be substantially improved.[2]

Few would argue with the point that transparency is in principle a good idea. But various problems arise when one tries to figure out what all this means in practice. I would like to discuss some of them here and present some modest proposals about what might be done in this area. While the issues that I raise here have broad implications for transparency in political research, I will be concerned here mainly with the use of a particular form of evidence: primary and secondary written sources.

Let me begin by talking about the first of the three points in the APSA guidelines, the one having to do with access to data. The basic notion here is that scholars should provide clear references to the sources they use to support their claims—and that it should be easy for anyone who wants to check those claims to find the sources in question. Of the three points, this strikes me as the least problematic. There's a real problem here that needs to be addressed, and there are some simple measures we can take to deal with it. So if it were up to me this would be the first thing I would focus on.

What should be done in this area? One of the first things I was struck by when I started reading the political science literature is the way a scholar would back up a point by citing, in parentheses in the text, a long list of books and articles, without including particular page numbers in those texts that a reader could go to see whether they provided real support for the point in question. Historians like me didn't do this kind of thing, and this practice struck me as rather bizarre. Did those authors really expect their readers to plow through those books and articles in their entirety in the hope of finding the particular passages that related to the specific claims being made? Obviously not. It seemed that the real goal was to establish the author's scholarly credentials by providing such a list. The

Marc Trachtenberg is Research Professor of Political Science at the University of California, Los Angeles. He is online at trachten@polisci.ucla.edu and http://www.sscnet.ucla.edu/polisci/faculty/trachtenberg.

[1] American Political Science Association 2012.

[2] See especially Moravcsik 2014.

whole practice did not facilitate the checking of sources; in fact, the inclusion of so much material would deter most readers from even bothering to check sources. It was amazing to me that editors would tolerate, and perhaps even encourage, this practice. But it is not unreasonable today to ask them to insist on precise page-specific references when such references are appropriate. The more general principle here is that citing sources should not be viewed as a way for an author to flex his or her academic muscles; the basic aim should be to allow readers to see, with a minimum of trouble on their part, what sort of basis there is for the claim being made. This is something journal editors should insist on: the whole process should be a lot more reader-friendly than it presently is.

A second easily-remediable problem has to do with the "scientific" system of citation that journals like the *American Political Science Review* use. With this system, references are given in parentheses in the text; those references break the flow of the text and make it harder to read. This problem is particularly serious when primary, and especially archival, sources are cited. The fact that this method makes the text less comprehensible, however, was no problem for those who adopted this system: the goal was not to make the argument as easy to understand as possible, but rather to mimic the style of the hard sciences. (When the APSR switched to the new system in June 1978, it noted that that system was the one "used by most scientific journals."[3]) It was obviously more important to appear "scientific" than to make sure that the text was as clear as possible. One suspects, in fact, that the assumption is that real science *should* be hard to understand, and thus that a degree of incomprehensibility is a desirable badge of "scientific" status. Such attitudes are very hard to change, but it is not inconceivable that journal editors who believe in transparency would at least permit authors to use the traditional system of citing sources in footnotes. It seems, in fact, that some journals in our field do allow authors to use the traditional system for that very reason.

The third thing that editors should insist on is that citations include whatever information is needed to allow a reader to find a source without too much difficulty. With archival material especially, the references given are often absurdly inadequate. One scholar, for example, gave the following as the source for a document he was paraphrasing: "Minutes of the Committee of Three, 6 November 1945, NARA, RG 59."[4] I remember thinking: "try going to the National Archives and putting in a call slip for that!" To say that this particular document was in RG59, the Record Group for the records of the State Department, at NARA—the National Archives and Records Administration—was not very helpful. RG59, as anyone who has worked in that source knows, is absolutely enormous. To get the Archives to pull the box with this document in it, you need to provide more detailed information. You need to tell them which collection within RG59 it's in, and you need to give them information that would allow them to pull the right box in that collection. As it turns out, this particular document is in the State Department Central (or Decimal) Files, and if you gave them the decimal citation for this document—which in this case happens to be 740.00119 EW/11-645—they would know which box to pull. How was I able to find that information? Just by luck: at one point, this scholar had provided a two-sentence quotation from this document; I searched for one of those sentences in the online version of the State Department's *Foreign Relations of the United States* series, and sure enough, an extract from that document containing those sentences had been published in that collection and was available online. That extract gave the full archival reference. But a reader shouldn't have to go through that much trouble to find the source for a claim.

How would it be possible to get scholars to respect rules of this sort? To a certain extent, simply pointing out, in methodological discussions like the one we're having now, how important this kind of thing is might have a positive effect, especially if people reading what we're saying here are convinced that these rules make sense. They might then make sure their students cite sources the right way. And beyond that people who do not respect those rules can be held accountable in book reviews, online roundtables, and so on. That's how norms of this sort tend to take hold.

Finally, let me note a fourth change that can be readily made. Scholars often set up their analyses by talking about what other scholars have claimed. But it is quite common to find people attributing views to other scholars that go well beyond what those scholars have actually said. Scholars, in fact, often complain about how other people have mis-paraphrased their arguments. It seems to me that we could deal with this problem quite easily by having a norm to the effect that whenever someone else's argument is being paraphrased, quotations should be provided showing that that scholar has actually argued along those lines. This would go a long way, I think, toward minimizing this problem.

Those are the sorts of changes that can be made using traditional methods. But it is important to note that this new concern with transparency is rooted, in large measure, in an appreciation for the kinds of things that are now possible as a result of the dramatic changes in information technology that have taken place over the past twenty-five years or so. All kinds of sources—both secondary and primary sources—are now readily available online, and can be linked directly to references in footnotes. When this is done, anyone who wants to check a source need only click a link in an electronic version of a book or article to see the actual source being cited. In about ten or twenty years, I imagine, we will all be required to provide electronic versions of things we publish, with hypertext links to the sources we cite. A number of us have already begun to

[3] Instructions to Contributors, *American Political Science Review* vol.72 no.2 (June 1978), 398 (http://www.jstor.org/stable/1954099, last accessed 6/27/2015).

[4] The citation appeared in notes 8, 10, 17, and 19 on pp. 16-18 of Eduard Mark's comment in the H-Diplo roundtable on Trachtenberg 2008 (roundtable: https://h-diplo.org/roundtables/PDF/Roundtable-X-12.pdf; original article: http://dx.doi.org/10.1162/jcws.2008.10.4.94). For my comment on Dr. Mark's use of that document, see pp. 52–54 in the roundtable.

move in that direction: I myself now regularly post on my website electronic versions of articles I write with direct links to the sources I have cited in the footnotes.[5] For the scholarly community as a whole, perhaps the most important thing here is to make sure that we have a single, unified set of standards that would govern how we adjust to, and take advantage of, the digital revolution. A possible next step for the Data Access and Research Transparency project would be to draft guidelines for book publishers and journal editors that might give them some sense for how they should proceed so that whatever norms do emerge do not take shape in a purely haphazard way.

But what about the two other forms of transparency called for in the guidelines? Individual researchers, the APSA *Guide* says, "should offer a full account of the procedures used to collect or generate" their data, and they "should provide a full account of how they draw their analytic conclusions from the data, i.e., clearly explicate the links connecting data to conclusions." What are we to make of those precepts?

Let's begin with the first one—with what the guidelines call "production transparency," the idea that researchers "should offer a full account of the procedures used to collect or generate the data." The goal here was to try to counteract the bias that results from people's tendency to give greater weight to evidence that supports their argument than to that which does not.[6] And it is certainly true that this problem of "cherry-picking," as it is called, deserves to be taken very seriously. But I doubt whether this guideline is an effective way of dealing with it. People will always say that their sources were selected in an academically respectable way, no matter how good, or how bad, the process really is. Forcing people to explain in detail how they have collected their data will, I'm afraid, do little to improve the situation.

To show what I mean, let me talk about an article I did a few years ago dealing with audience costs theory—that is, with the claim that the ability of a government to create a situation in which it would pay a big domestic political price for backing down in a crisis plays a key role in determining how international crises run their course.[7] I identified a whole series of crises in which one might expect the audience costs mechanism to have played a role. I then looked at some historical sources relating to each of those cases to see whether that mechanism did in fact play a significant role in that particular case; my conclusion was that it did not play a major role in any of those cases. If I had been asked to explain how I collected my data, I would have said "I looked at all the major historical sources—important books and articles plus easily available collections of diplomatic documents—to see what I could find that related to the issue at hand." That I was using that method should already have been clear to the reader—all he or she would have had to do was look at my footnotes—so including an explanation of that sort would contribute little. But maybe I would be expected to go further and explain in detail the process I used to decide which sources were important and which were not. What would this entail? I could explain that I searched in JSTOR for a number of specific search terms; that after identifying articles that appeared in what I knew were well-regarded journals, I searched for those articles in the Web of Science to see how often they were cited and to identify yet further articles; that I supplemented this by looking at various bibliographies which I identified in various ways; and that having identified a large set of studies, I then looked at them with the aim of seeing how good and how useful they were for my purposes. I could go further still and present a detailed analysis of the intellectual quality of all those secondary works, including works I had decided not to use. That sort of analysis would have been very long—much longer, in fact, than the article itself—but even if I did it, what would it prove? The works I liked, a reader might well think, were the works I happened to find congenial, because they supported the argument I was making. And even if the reader agreed that the studies I selected were the best scholarly works in the area, how could he or she possibly know that I had taken into account all of the relevant information found in that material, regardless of whether it supported my argument? I have my own ways of judging scholarly work, and I doubt whether my judgment would be much improved if authors were required to lay out their methods in great detail.

I also wonder about how useful the third guideline, about "analytic transparency," will be in practice. The idea here is that researchers "should provide a full account of how they draw their analytic conclusions from the data"—that is, that they should "clearly explicate the links connecting data to conclusions." It sometimes seems that the authors of the guidelines think they are asking scholars to do something new—to provide a statement about method that would be a kind of supplement to the book or article in question. But my basic feeling is that scholarly work, if it is any good at all, should *already* do this. And it is not just that a scholarly work should "explicate the links connecting data to conclusions," as though it is just one of a number of things that it should do. My main point here is that this is what a scholarly work *is*, or at least what it should be. The whole aim of a scholarly article, for example, should be to show how conclusions are drawn from the evidence.

It is for that reason that a work of scholarship should have a certain formal quality. The goal is *not* to describe the actual process (involving both research and thinking) that led to a set of conclusions. It is instead to develop an argument (necessarily drawing on a set of assumptions of a theoretical nature) that shows how those conclusions follow from, or at least are supported by, a certain body of evidence. It is certainly possible to explain in detail how in practice one reached those conclusions—I spend a whole chapter in my methods

[5] I have now posted five such articles: Trachtenberg 2005; 2011; 2013a; 2013b; and 2013c. I generally include a note in the published version giving the location of the electronic version and noting that it contains links to the materials cited. The texts themselves, in both versions, are linked to the corresponding listing in my c.v. (http://www.sscnet.ucla.edu/polisci/faculty/trachtenberg/cv/cv.html).

[6] Moravcsik 2014, 49.

[7] Trachtenberg 2012.

book showing how in practice one does this kind of work—but normally this is not what a work of scholarship does.[8] It is normally much leaner, and has a very different, and more formal, structure. One never talks about everything one has looked at; one tries instead to distill off the essence of what one has found and to present it as a lean and well-reasoned argument. Nine-tenths of the iceberg will—and indeed should—lie below the surface; it is important to avoid clutter and to make sure that the argument is developed clearly and systematically; the logic of the argument should be as tight as possible.

So a lot of what one does when one is analyzing a problem is bound to be missing from the final text, and it is quite proper that it should not be included. Let me again use that audience costs paper as an example. After it came out, *Security Studies* did a forum on it, and one of the criticisms that was made there had to do with what could be inferred from the historical evidence. The fact that one could not find much evidence that political leaders in a particular crisis deliberately exploited the audience costs mechanism, Erik Gartzke and Yonatan Lupu argued, does not prove much, because there is no reason to suppose that intentions would be revealed by the documentary record; absence of evidence is not evidence of absence.[9]

The answer here is that one can in certain circumstances infer a great deal from the "absence of evidence." What one can infer depends on the kind of material we now have access to—on the quality of the documentation, on what it shows about how freely political leaders express themselves when they were talking about these issues in private to each other at the time. One can often reach certain conclusions about what their intentions were in choosing to react the way they did on the basis of that material. Those conclusions might not be absolutely rock-solid—one never knows for sure what is in other people's minds—but one can often say more than zero, and sometimes a lot more than zero, about these questions. But should that point have been explained in the original paper? It would be okay to deal with it if one was writing a piece on historical methodology, but these methodological points should not have to be explained every time one is dealing with a substantive issue. For it is important to realize that you always pay a price in terms of loss of focus for dealing with ancillary matters and drifting away from the issue that is the real focus of the analysis.

My point here is that a good deal more goes into an assessment of what one is to make of the evidence than can be summed up in the sort of statement this third guideline seems to call for. Philosophers of science point out that in reaching conclusions "logical rules" are less important than "the mature sensibility of the trained scientist."[10] In our field this basic point applies with particular force. In making these judgments about the meaning of evidence, one brings a whole sensibility to bear—a sense one develops over the years about how the political system works, about how honest political leaders are when talking about basic issues in private, about what rings true and what is being said for tactical purposes.

So I doubt very much the second and third guidelines will be of any great value, and I can easily imagine them being counter-productive. This does not mean, of course, that we don't have to worry about the problems that the authors of these guidelines were concerned with. How then should those problems be dealt with? How, in particular, should the problem of cherry-picking be dealt with?

To begin with, we should consider how the current system works. We maintain standards basically by encouraging scholarly debate. People often criticize each others' arguments; those debates help determine the prevailing view. It is not as though we are like trial lawyers, using every trick in the book to win an argument. We use a more restrained version of the adversarial process, where a common interest in getting things right can actually play a major role in shaping outcomes.

This is all pretty standard, but there is one area here where I think a change in norms would be appropriate. This has to do with the prevailing bias against purely negative arguments, and with the prevailing assumption that an author should test his or her own theories.

It is very common in the political science literature to see an author lay out his or her own theory, present one or two alternatives to it, and then look at one or more historical cases to see which approach holds up best. And, lo and behold, the author's own theory always seems to come out on top. We're all supposed to pretend that the author's obvious interest in reaching that conclusion did not color the analysis in any way. But that pose of objectivity is bound to be somewhat forced: true objectivity is simply not possible in such a case. I personally would prefer it if the author just presented the theory, making as strong a case for it as possible, and did not pretend that he or she was "testing" it against its competitors. I would then leave it to others to do the "testing"—and that means that others should be allowed to produce purely negative arguments. If the "test" shows that the theory does not stand up, the analyst should be allowed to stop there. He or she should not be told (as critics often are) that a substitute theory needs to be produced. So if I were coming up with a list of rules for journal editors, I would be sure to include this one.

[8] See Trachtenberg 2006, ch.4 (https://www.sscnet.ucla.edu/polisci/faculty/trachtenberg/cv/chap4.pdf).

[9] Gartzke and Lupu 2012, 393 n7, 394f.

[10] Kuhn 1970, 233. Note also Stephen Toulmin's reference to the "judgment of authoritative and experienced individuals" (1972, 243).

References

American Political Science Association, Committee on Ethics, Rights and Freedoms. 2012. *A Guide to Professional Ethics in Political Science*, 2nd edition. Washington: APSA. (http://www.apsanet.org/Portals/54/APSA%20Files/publications/ethicsguideweb.pdf, last accessed 7/1/2015).

Gartzke, Erik and Yonatan Lupu. 2012. "Still Looking for Audience Costs." *Security Studies* vol. 21, no. 3: 391–397. (http://dx.doi.org/10.1080/09636412.2012.706486, last accessed 6/27/2015)

Kuhn, Thomas. 1970. "Reflections on My Critics." In Imre Lakatos and Alan Musgrave, eds., *Criticism and the Growth of Knowledge*. Cambridge: Cambridge University Press, 231–278.

Moravcsik, Andrew. 2014. "Transparency: The Revolution in Qualitative Research." *PS: Political Science and Politics* vol. 47, no. 1:

48–53. (http://dx.doi.org/10.1017/S1049096513001789, last accessed 6/27/2015)
Trachtenberg, Marc. 2005. "The Iraq Crisis and the Future of the Western Alliance," in *The Atlantic Alliance Under Stress*, edited by David M. Andrews. Cambridge: Cambridge University Press. Extended version with links to sources at http://www.sscnet.ucla.edu/polisci/faculty/trachtenberg/useur/iraqcrisis.html.
———. 2006. *The Craft of International History: A Guide to Method*. Princeton University Press. (http://graduateinstitute.ch/files/live/sites/iheid/files/sites/international_history_politics/users/stratto9/public/Trachtenberg%20M%20The_Craft_of_International_His.pdf)
———. 2008. "The United States and Eastern Europe in 1945: A Reassessment." *Journal of Cold War Studies* vol. 10, no. 4: 94–132. (http://dx.doi.org/10.1162/jcws.2008.10.4.94, last accessed 6/27/2015)
———. 2011. "The French Factor in U.S. Foreign Policy during the Nixon-Pompidou period, 1969–1974." *Journal of Cold War Studies* vol. 13, no. 1: 4–59. Extended version with links to sources at https://www.sscnet.ucla.edu/polisci/faculty/trachtenberg/ffus/FrenchFactor.pdf.
———. 2012. "Audience Costs: An Historical Analysis." *Security Studies* vol. 21, no. 1: 3–42. Long version at http://www.sscnet.ucla.edu/polisci/faculty/trachtenberg/cv/audcosts(long).doc.
———. 2013a. "Dan Reiter and America's Road to War in 1941." *H-Diplo/ISSF Roundtable* vol. 5, no. 4 (May). Extended version with links to sources at http://www.polisci.ucla.edu/faculty/trachtenberg/cv/1941.doc.
———. 2013b. "Audience Costs in 1954?" *H-Diplo/ISSF*, September. Extended version with links to sources at http://www.sscnet.ucla.edu/polisci/faculty/trachtenberg/cv/1954(13)(online).doc.
———. 2013c. "Kennedy, Vietnam and Audience Costs." Unpublished manuscript, November. Extended version with links to sources at http://www.sscnet.ucla.edu/polisci/faculty/trachtenberg/cv/viet(final).doc.
Toulmin, Stephen. 1972. *Human Understanding*. vol.1. Princeton: Princeton University Press.

Transparency In Field Research

Transparent Explanations, Yes. Public Transcripts and Fieldnotes, No: Ethnographic Research on Public Opinion

Katherine Cramer
University of Wisconsin, Madison

I am a scholar of public opinion. My main interest is in examining how people understand, or interpret, politics. For that reason, much of my work involves listening to people talk with others with whom they normally spend time. When I listen to the way they sort out issues together, I am able to observe what they perceive to be important, the narratives they use to understand their world, the identities that are central to these understandings, and other important ingredients of public opinion.

My work is therefore primarily qualitative, and usually interpretivist. By interpretivist, I mean that I am trying to capture how people perceive or attribute meaning to their worlds. I treat that effort as a necessary part of trying to understand why they express the opinions that they do.

Across the course of my career, transparency has been a professional necessity. My methods are rather unusual in the field of public opinion research, so the burden is on me to teach my readers and my reviewers what I am doing, why I am doing it, and how my work should be judged. Usually, public opinion scholars focus on individuals' preferences and how to predict them, not on the process of understanding. In addition, we tend to be well versed in the strengths and weaknesses of polling data, but basically unfamiliar with conversational data.

Put another way, reviewers are likely to dive into my papers looking for the dependent variable and the strength of evidence that my results can be generalized to a national population. But my papers usually do not provide information on either of these things. Unless I explain why my work has different qualities to be judged, the typical reviewer will quickly tune out and give the paper a resounding reject after the first few pages.

So the first thing I have had to be transparent about is the fact that much of my work is not attempting to predict preferences. My work typically does not describe how a set of variables co-vary with one another to bring about particular values on a dependent variable. Indeed, I'm not usually talking about causality. These characteristics are just not what scholars typically come across in political science public opinion research. I have had to go out of my way to explain that my work uses particular cases to help explain in detail the process of a group of people making sense of politics. I have had to be up front about the fact that my goal is to help us understand how it is that certain preferences are made obvious and appropriate when objective indicators about a person's life would suggest otherwise.

For example, in a piece I published in the *APSR* in 2012,[1] reviewers helpfully pointed out that I had to bluntly state that my study used particular cases to study a broader question. In short, that article reported the results of a study in which I invited myself into conversations among groups of people meeting in gathering places like gas stations and cafés in communities throughout Wisconsin, especially small towns and rural places, so that I could better understand how group consciousness might lead people to support limited government

Katherine Cramer is Professor of Political Science and Director of the Morgridge Center for Public Service at the University of Wisconsin-Madison. She is online at kathy.cramer@wisc.edu and https://faculty.polisci.wisc.edu/kwalsh2/.

[1] Cramer Walsh 2012.