

Transparency in Experimental Research*

Cheryl Boudreau

Abstract

In recent years, the credibility of social science has been tarnished by widely discussed replication failures and a lack of reporting about what exactly researchers did when conducting their studies. In response, scholars, policymakers, and the public have called for greater transparency in social science research. In this chapter, I emphasize that transparency is an important public good. However, because individual researchers lack incentives to contribute to this public good, institutional solutions are needed. I discuss three institutions that facilitate transparency in experimental research: (1) preregistration, (2) reporting guidelines, and (3) the Data Access and Research Transparency (DA-RT) initiative. I also offer recommendations for what kinds of information researchers should preregister and report in their published articles and appendices. I conclude with a discussion of how researchers might be incentivized to make greater use of these institutions when designing, conducting, and publishing their experiments.

In recent years, the credibility of social science has been tarnished by widely discussed replication failures and a general lack of reporting about what exactly researchers did when conducting their studies. For example, a recent effort to

replicate 100 experimental studies published in highly ranked psychology journals found considerable differences between the results reported in the original and replication studies (Open Science Collaboration 2015). While there is some evidence that political science experiments might fare better (Coppock 2018; Mullinix et al. 2015¹), a

* I thank James N. Druckman, Donald P. Green, Neil Malhotra, Natalia Candelo, Jason Seawright, and participants in the Advances in Experimental Political Science conference at Northwestern University for helpful feedback.

1 Mullinix et al. (2015) successfully replicated over 80% of the 20 Time-Sharing Experiments for the Social Sciences (TESS) experiments that they included in

similarly large-scale replication of published experimental research in political science has not yet occurred. There is also reason to think that similar replication failures might happen in political science, given the lack of sufficient reporting in published experimental research (Gerber et al. 2014). Indeed, Gerber et al.'s (2014) content analysis shows that key details necessary for replication (e.g., the number of subjects assigned to treatment and control groups, the exact experimental materials, the measurement of outcome variables) are often omitted from laboratory, survey, and field experiments published in our discipline's top general interest and subfield journals. This lack of reporting would make replication difficult, resulting in what may be many failed attempts.

In response to these replication failures and reporting omissions, scholars, policymakers, and the public have called for greater transparency in social science research (Ioannidis 2014; Lupia 2008; Lupia and Elman 2014; Miguel et al. 2014). That is, they have emphasized the importance of openness about all aspects of scholarly research. This involves making the following kinds of information public: the details of the procedures used (including the materials used in experimental research), the hypotheses to be tested and data analyses to be conducted, as well as the data and code used to generate the results. The hope is that greater transparency about what exactly researchers did in their studies will facilitate replication and, in

their study. They considered a replication to be successful when they obtained a statistically significant result in the same direction as the original study or a null result when the original study produced a null result. In contrast to the Open Science Collaboration (2015), the materials for all of these TESS experiments were publicly available. This facilitated replication by allowing Mullinix et al. to use materials that were identical to those used in the original studies. Mullinix et al.'s higher rate of successful replication might also stem from the fact that their replication effort included unpublished studies, thereby avoiding well-established publication biases that also undermine replication efforts (Franco et al. 2014; see also Chapter 19 in this volume). Mullinix et al. also replicated experiments that had passed an initial peer review, were relatively well powered, and used a high-quality sample. Thus, their effort might represent a most likely case for successful replication.

turn, bolster the credibility of science (for a complementary discussion, see Chapter 19 in this volume). Greater transparency will also aid in the evaluation of claims that researchers make based on their research and shed light on whether and when their results generalize to real-world policy problems.

Despite the benefits of greater transparency in social science research, individual researchers often lack incentives to provide it. Indeed, transparency is a classic example of a public good in that many researchers lack incentives to contribute to it, even though academic disciplines would enjoy the benefits of it (e.g., greater scientific credibility) (see Nosek et al. 2015).^{2,3} As a result, transparency (like other public goods) will be undersupplied unless institutions give researchers an incentive to contribute. Until recently, there were few requirements or norms in political science that encouraged open research practices. Further, the reward systems in place at universities, academic publishers, and granting agencies often create incentives to be *less* transparent about one's research (Ioannidis 2014; Miguel et al. 2014; Nosek et al. 2015). These institutions typically reward large quantities of publications, and the publication process itself tends to favor statistically significant, novel findings over null results and replication studies (Nosek, Spies, and Motyl 2012). These factors may discourage researchers from reporting mixed or null results and conducting replications of their own or others' research. They may also discourage researchers from taking the time to make their research transparent, given that this takes time and effort away from publishing per se.

However, scholarly incentives are beginning to change as political science and other

- 2 In this way, transparency (like other public goods) is non-excludable and non-rival. That is, it is not possible to exclude researchers who are not transparent from the benefits of greater transparency (i.e., scientific credibility). Further, one researcher's enjoyment of the scientific credibility that results from greater transparency does not reduce other researchers' ability to enjoy that same benefit.
- 3 If preregistration and pre-analysis plans increase the chances that a manuscript is accepted for publication, then researchers would have an incentive to be more transparent.

social science disciplines implement institutions that encourage greater transparency. In particular, the experimental research community in political science has embraced three institutions designed to facilitate transparency: (1) preregistration and pre-analysis plans, (2) reporting guidelines, and (3) the Data Access and Research Transparency (DA-RT) initiative.⁴ In this chapter, I discuss each of these institutions and describe how they promote transparency in experimental research. I also offer specific recommendations for what kinds of information researchers should preregister and ultimately report in their published articles and appendices. I conclude with a discussion of how researchers might be incentivized to make greater use of these institutions when designing, conducting, and publishing their experiments.

18.1 Institutions That Promote Transparency in Experimental Research

18.1.1 Preregistration and Pre-Analysis Plans

One institution that promotes transparency in experimental research is preregistration, which may also include a pre-analysis plan. Preregistration is the practice of developing one's research questions, hypotheses, research design, and analyses *before* observing the data and making that information public on an independent registry. As part of the preregistration process, researchers may also create and submit pre-analysis plans that describe in detail the procedures they will use when collecting and analyzing the data (e.g., planned data analyses and statistical tests). For example, a researcher who wants to conduct an experiment to study how campaign advertisements

affect citizens' perceptions of candidates' traits (e.g., competence, viability) would begin by developing his or her research question, hypotheses, experimental design, and planned analyses. Before conducting the experiment, the researcher would submit a written description of these aspects of the study to an independent registry. The researcher may also choose to submit a pre-analysis plan that provides greater detail about the planned data analyses and statistical tests. In general, preregistration and pre-analysis plans encourage researchers to think carefully about their planned experiments before implementing them, which is good research practice.

In addition to being valuable tools for researchers, preregistration and pre-analysis plans are beneficial to the scientific community because they facilitate transparency. One way that preregistration does so is by ensuring that there is a record of unpublished research on a particular topic. That is, even if researchers do not submit a pre-analysis plan, they can still create a record of their research by preregistering the study itself. This makes their data collection effort public and provides information about the number of studies contained in the "file drawer" (i.e., studies that were conducted but never published). This is important because of biases in the publication process that favor significant findings over null results (Franco et al. 2014). By providing information about the existence of unpublished studies, preregistration can shed light on whether other data collection efforts on a particular topic were undertaken. And, if researchers update their registrations to include a description of the results they ultimately obtained, this can also reveal whether published findings on a particular topic are counterbalanced by unpublished null or contradictory results (for further discussion, see Chapter 19 in this volume). It is similarly important for researchers conducting meta-analyses to know the universe of data that was collected on a particular topic.

Preregistration together with a pre-analysis plan also facilitates transparency by revealing what findings were predicted

4 Coffman and Niederle (2015) and Coffman et al. (2017) propose other institutions that encourage independent replications of published research. These include the creation of a journal dedicated to publishing replications and a requirement that replications be cited together with the citation of the original study. Nonetheless, transparency is necessary for such independent replications to take place.

ex ante (i.e., before observing the data) as part of confirmatory research versus discovered *ex post* as part of exploratory research. Both types of research are valuable, and which approach researchers take will depend upon their particular topic and research program. However, it is important for researchers to be clear about this aspect of their research because it provides information about whether particular theories were tested directly versus developed during the data analysis process (Ledgerwood 2019). It also enables scholars to determine their level of confidence in a study's findings and statistical tests (Ledgerwood 2019; Nosek et al. 2018). For example, if a researcher specified the analyses used to test his or her predictions prior to observing the data, then we have greater confidence that he or she did not obtain significant results simply by chance as a result of repeated tests. Conversely, if a researcher blurs the distinction between *ex ante* and *ex post* analyses, this can result in an unknown degree of type-1 error inflation (i.e., the probability that an observed significant result is merely due to chance) (Ledgerwood 2019; Nosek et al. 2018; Simmons et al. 2011).

Practically speaking, there are several independent registries that researchers can use to preregister their studies. These registries were developed by scholars in different disciplines and feature varying levels of flexibility. One of the most basic registries is [Aspredicted.org](https://aspredicted.org), which was developed by psychologists. This registry asks all researchers to answer the same eight questions about their planned study. These questions ask researchers to report their main question or hypothesis, as well as the intended sample size, number of experimental conditions, and measurement of key dependent variables. The questions also ask researchers to describe the analyses they will use to test the main question or hypothesis, as well as any secondary analyses they plan to conduct.⁵ Researchers

are also asked whether the data have yet to be collected or whether some or all of the data have been collected. They are also asked whether there is anything else that they would like to preregister. Researchers can choose to make their answers to these questions public or keep them private. In this way, [Aspredicted.org](https://aspredicted.org) provides a standardized preregistration process that reveals what researchers predicted and planned to analyze *ex ante* versus what they explored *ex post*. An advantage of this registry is that it is easy to use and unambiguous about what aspects of their studies researchers should preregister. The disadvantage is that it is less flexible than other registries and does not allow researchers to provide additional details about their planned studies (such as the exact materials used in the experiment or a separate pre-analysis plan). In fact, this registry permits researchers to provide few details about their planned data analyses if they so choose.

An independent registry that provides greater flexibility is the Evidence in Governance and Politics (EGAP) registry. This registry was designed by political scientists for the preregistration of experimental and observational studies of governance and politics. Like [Aspredicted.org](https://aspredicted.org), the EGAP registry asks researchers to answer a standardized list of questions about their planned study. Some of these questions are similar to the [Aspredicted.org](https://aspredicted.org) questions in that they ask researchers to report their main hypotheses, planned analyses, and sample size. However, the EGAP registry also asks researchers to provide additional information, including whether and when their study received institutional review board (IRB) approval, whether they conducted a power analysis before the data were collected, whether the researchers themselves or a third party will implement the planned study, and whether the researchers have permission to publish the study regardless of what the results show. In contrast to [Aspredicted.org](https://aspredicted.org), the EGAP registry also encourages researchers to submit a more detailed pre-analysis plan, as well as the exact materials that will be used in the study, the code that will be used to analyze

5 Information about the number of analyses that researchers conduct is necessary for correcting *p*-values for multiple comparisons.

the data, and the actual data once they are collected. The EGAP registry is also different from [Aspredicted.org](https://aspredicted.org) in that it requires registrations to be made public within 18 months. This promotes transparency because it reveals what researchers planned to do in their study and allows others to compare that to what the researchers actually did in their published research. This also facilitates replication because it makes public the materials on which particular studies are based. It also ensures that any unpublished studies in the “file drawer” are made public.

The most flexible, but also the most involved, independent registry is the Center for Open Science’s Open Science Framework (OSF) registry. Distinct features of this registry are its provision of separate forms for different kinds of studies, as well as the comprehensiveness of those forms. Like the [Aspredicted.org](https://aspredicted.org) and EGAP registries, the OSF registry provides a standard general purpose preregistration form. However, the OSF form is considerably more detailed about what it asks researchers to report (i.e., 10 pages of questions compared to one and four pages for the [Aspredicted.org](https://aspredicted.org) and EGAP registries, respectively). Like the EGAP registry, the OSF registry also encourages researchers to submit a separate pre-analysis plan and additional materials that will be used in the study. In contrast to the other registries, however, it provides separate forms for registering different types of studies, such as replication studies, pre-accepted registered reports, and studies for which the data have already been collected. It also allows researchers to keep their preregistration materials private for up to four years, at which point they will be made public. These aspects of the OSF registry give researchers considerable flexibility when preregistering their studies and encourage them to be transparent about all aspects of their research. That said, the greater demands of the OSF registry’s forms and the number of different choices to be made also potentially present greater barriers to preregistration.⁶

6 Though not discussed here, economists have also developed independent registries for experimental

18.1.2 Reporting Guidelines

A second institution that facilitates transparency in experimental research is the adoption of reporting guidelines. Reporting guidelines are checklists that guide researchers on what information to report about their experiments in their published articles and appendices. They promote transparency by encouraging researchers to publish the details of what exactly they did when designing, conducting, and analyzing data from their experiments.⁷ By providing a specific list of items that should be reported, they also help prevent information that is necessary for evaluating a study and its conclusions from being omitted from published research.

Transparency in published research is important for several reasons. First, the scholarly community will be unable to make sense of and evaluate a particular study if relevant details are not provided. For example, if a researcher does not publish his or her experimental materials (either in the article or an appendix), it is impossible to tell what participants experienced during the experiment and whether the treatment was administered as described in the published article. It is also difficult to know what inferences should be drawn from a study that is not described in adequate detail. Second, policymakers and the public will be unable to determine whether a study’s results are relevant to real-world policy problems. For instance, if a researcher does

and other types of research (e.g., the American Economic Association’s Randomized Controlled Trials Registry, as well as the Registry for International Development Impact Evaluations (RIDIE)). The American Economic Association’s registry allows researchers to register their studies either *ex ante* or *ex post*. It also assigns a digital object identifier to each study that is registered so that any changes made to it can be tracked over time.

7 Importantly, scholars have developed transparency checklists that provide information about the research process that produced certain results. These checklists reveal whether a study was preregistered, whether the preregistration included certain kinds of information (e.g., intended statistical analyses), and whether the resulting manuscript provides details necessary for evaluating the study (e.g., the operationalization of measures, the rationale for the sample size) (Aczel et al. 2020).

not report information about the nature of the sample (e.g., a population-based sample versus a convenience sample; see Chapter 9 in this volume), it is difficult to know whether the study's conclusions should apply to the population at large or only to a particular segment of it (see Chapters 15 and 21 in this volume). Finally, a lack of clarity about what exactly researchers did in their experiments or how they measured their outcome variables hinders scholarly efforts to replicate and extend the research. Indeed, it is unlikely that researchers will be inspired by and motivated to build upon prior experiments that were not explained in sufficient detail. This, in turn, undermines scientific progress and the credibility of social science research.

Given the importance of adequate reporting in published experimental research, scholars in a variety of academic disciplines have developed official reporting guidelines. In many cases, the reporting guidelines have been adopted or endorsed by academic journals as a statement of what information researchers should include in the articles and appendices that they publish there. For example, scholars in medicine created the Consolidated Standards of Reporting Trials (CONSORT) guidelines to improve the transparency of randomized controlled trials in medical research. The CONSORT guidelines include a 25-item checklist, as well as a diagram that helps researchers to convey how many participants were assessed for eligibility in a study, excluded for particular reasons, randomized into treatment and control groups, and ultimately analyzed. Over 500 medical journals (including leading journals like the *New England Journal of Medicine* and the *Journal of the American Medical Association*) have endorsed the CONSORT guidelines, with many of them requiring researchers to complete the CONSORT checklist and/or diagram before submitting their manuscripts for review.

Social scientists have engaged in similar efforts. In particular, economists have developed guidelines for what types of information should be reported in published experimental economics research (Palfrey and Porter

1991). These guidelines include details that are of particular relevance in economics experiments (e.g., details about payments to subjects, practice trials, matching procedures, and tests of subjects' comprehension of the experiment). Prominent economics journals (e.g., *Econometrica*, *American Economic Review*, *Journal of Political Economy*) have embraced aspects of these guidelines and list on their websites the information that should be reported about experimental research that is published there. In psychology, the American Psychological Association produced and continues to update its Journal Article Reporting Standards (JARS). A unique feature of these reporting guidelines is that they provide separate lists of items to be reported for quantitative, qualitative, and mixed-methods research (Appelbaum et al. 2018; Levitt et al. 2018). The JARS also specify additional information that should be reported in experimental research (e.g., evidence of treatment fidelity, use of incentives, and length of experimental sessions) and in replication studies (e.g., comparisons of the recruitment procedures, demographic characteristics of the subjects, and administrations of the experiment in the original versus replication study) (Appelbaum et al. 2018).

In political science, the Experimental Research Section of the American Political Science Association (APSA) formed a committee to develop reporting guidelines for political science experiments.⁸ The committee consisted of six members with expertise in the different types of experiments typically conducted in political science: laboratory, survey, and field experiments. The committee began by creating separate lists of items to be reported for each type of experiment and also reviewing the reporting guidelines developed in medicine, economics, and psychology. In light of the overlap in the separate lists of reporting items for

⁸ In a related effort, EGAP developed a set of principles that apply to its membership. These principles involve human subjects' protection, transparency, rights surrounding the review and publication of findings, the publication of data, and the disclosure of any remuneration the researcher received (see McDermott 2014 for further discussion).

laboratory, survey, and field experiments, the committee created a single set of reporting guidelines (Gerber et al. 2014). Like the guidelines in other disciplines, the political science guidelines recommend that researchers report the information necessary for evaluating and replicating their studies. This information includes key details about the selection and recruitment of subjects, the randomization procedures, the administration of the treatment(s), the measurement of outcome variables and covariates, the nature of any attrition or missing data, and the role of any funding agencies in the analysis of the experimental data. These guidelines are described in detail in Gerber et al. (2014) and in Section 18.2.2 of this chapter (where the checklist of reporting items that Gerber et al. (2015) subsequently developed is also provided). They were also approved by a vote of the Experimental Research Section⁹ of APSA and are recommended by its journal, the *Journal of Experimental Political Science*.

While the political science reporting guidelines share common elements with those developed in other disciplines, they are unique in that they explicitly consider the different types of experiments that political scientists conduct. Thus, within the single set of guidelines are specific instructions for certain types of experiments. For example, the guidelines suggest that researchers conducting survey experiments provide information about the survey firm that was used, how it recruits respondents, the response rate (if possible), and details about the weighting procedures used (if any). The guidelines also specify that laboratory experiments (and other types of experiments,

when relevant) report information such as the total number and length of experimental sessions, whether incentives were provided, the order of the treatments in within-subject designs, and whether subjects were quizzed on the experimental instructions.

18.1.3 The DA-RT Initiative

A third institution that promotes transparency in experimental research is the DA-RT initiative in political science. This initiative has been led by prominent quantitative and qualitative researchers in political science with the goal of fostering openness about the data used to produce particular conclusions (Lupia and Elman 2014). To this end, the DA-RT initiative has focused on creating incentives for researchers to share information about the evidence on which their claims are based so that others can better interpret and evaluate them. Unlike reporting guidelines, the DA-RT initiative is not specific to a particular type of research, nor does it impose a single set of standards on all political science research.¹⁰ Rather, it endorses general principles about data access and research transparency and facilitates the application of those principles within quantitative and qualitative research traditions (Lupia and Elman 2014).¹¹ As is true of the reporting guidelines described above, the experimental research community in political science and its journal (the *Journal of Experimental Political Science*) have embraced the DA-RT principles.

9 While the guidelines enjoy broad support, there have been objections to certain aspects of the guidelines and proposals for extensions to them. For example, Mutz and Pemantle (2015) argue that the guidelines should recommend manipulation checks (they currently do not), should only recommend the reporting of response rates if the researcher claims that he or she has a random probability sample, and should not require the reporting of randomization checks (see Gerber et al. 2015 for a response). Franco et al. (2017) build on the guidelines by proposing standards for the use of weights in survey experiments and the reporting of how they were constructed.

10 While there have been debates about this initiative (particularly as it applies to qualitative research; see Elman and Kapiszewski 2014), most are not relevant to experimental research. Connors et al. (2019), however, raise a consideration that is relevant to some types of experimental research: that survey respondents' knowledge about transparency (i.e., that their responses will be made public, even if anonymous) affects their responses.

11 In particular, it appointed committees that were tasked with developing more specific guidelines for particular research communities. The end result was the creation of documents entitled "Guidelines for Data Access and Research Transparency in the Quantitative Tradition" and "Guidelines for Data Access and Research Transparency in the Qualitative Tradition" (Lupia and Elman 2014).

One way in which the DA-RT initiative facilitates transparency in experimental and other kinds of research is through its efforts to revise APSA's policies governing data access and research transparency. In particular, it helped revise APSA's *Guide to Professional Ethics in Political Science* so that it now states that "researchers have an ethical obligation to facilitate the evaluation of their evidence-based knowledge claims through data access, production transparency, and analytic transparency so that their work can be tested or replicated." The ethics guide also now specifies that this involves sharing the data used to make particular claims, the details of the procedures used when collecting and analyzing the data, as well as explaining the connection between one's data and conclusions. This is an accomplishment because the previous ethics guide encouraged data access only when findings had been challenged. In contrast, the revised ethics guide makes data access and research transparency the default, rather than the exception (Lupia and Elman 2014).

The DA-RT initiative has also promoted transparency through its efforts to get academic journals to implement its data access and research transparency principles. To this end, the DA-RT initiative wrote the Journal Editors Transparency Statement (JETS) that outlines procedures that a journal will adopt to ensure that the data, code, and/or analyses used in the research it publishes will be made public. The JETS specifies that a journal will require authors to make their data publically available on a trusted digital repository (such as Dataverse) at the time of publication, though editors may grant exceptions for confidential data.¹² It also states

that a journal will require authors to describe the analytical procedures used to generate their results clearly and provide access to all materials used when analyzing the data. In addition, it specifies that a journal will maintain a consistent data citation policy to ensure that credit is given to authors who collect and generate new data. It also imposes a deadline for journals to implement these changes.

The DA-RT initiative then approached political science journal editors about signing on to the JETS, and 27 of them did so as of 2015 (including the editors of the *Journal of Experimental Political Science*). Several of these journals (*American Journal of Political Science*, *Quarterly Journal of Political Science*, and *International Organization*) went one step further by verifying that the data and code that authors submit actually reproduce the results they report in their articles before publishing them. As such, these 27 journals as of 2015 require authors to make public the data, code, and analyses that were used to produce their results as a condition of publication. This increases transparency by allowing others to observe and inspect the data and analyses that a researcher used to produce his or her results. It also facilitates reproducibility and replication, which are important aspects of scientific credibility, and creates powerful incentives for researchers to be more open about the data, code, and analyses on which their claims are based.

18.2 Recommendations for What to Preregister and Report in Experimental Research

The three institutions described above are designed to facilitate transparency in social science research in general and experimental research in particular. In principle, these institutions create incentives for individual researchers to be more open about all aspects of their research, from its design and implementation to its analysis and ultimate publication. In doing so, they help ensure the provision of a public good that yields an important benefit for academic disciplines (namely, greater scientific credibility).

¹² The importance of making scientific data publicly available was recently illustrated in political debates over environmental regulations that govern clean air and water. In particular, the Environmental Protection Agency (EPA) under the Trump administration proposed a policy whereby the EPA would no longer consider scientific research unless the data on which that research is based can be made public. Such a policy would prevent regulators from considering a large body of scientific research on the health-related effects of exposure to air pollution, pesticides, and other chemicals when crafting rules (Friedman 2018).

In practice, however, it is not always clear what exactly researchers should preregister and report in their published articles. Achieving greater clarity on these fronts is important because the institutions will not have their intended effects if individual researchers do not know how to respond to them. Stated differently, the institutions are necessary but not sufficient for increasing transparency. Also needed is an understanding of what they mean for individual researchers. Thus, in what follows, I offer specific recommendations about the kinds of information researchers should provide when preregistering their experiments and reporting the results in published research.

18.2.1 Preregistration Recommendations

Given the relatively recent focus on preregistration in political science, as well as the variety of independent registries available, researchers may be unsure about what aspects of their studies should be preregistered, whether a pre-analysis plan should be submitted, and what information to include in a pre-analysis plan. In general, preregistration and pre-analysis plans should be thought of as planning documents for a prospective experiment (see Druckman et al. 2018). Ideally, the process of designing an experiment to answer a particular research question or to test a given hypothesis will provide a roadmap for the kinds of information to include during the preregistration process. Table 18.1 lists the types of information that researchers should consider when planning their experiments and therefore provide during the preregistration process. Depending on the specific registry used, this information can either be provided on the preregistration form itself or included in a supplemental pre-analysis plan.

As the recommendations in Table 18.1 indicate, the preregistration process need not be a daunting or overly time-consuming process. Indeed, the items listed in Table 18.1 are things that researchers should have already considered before administering their experiment. For example, when planning and designing an experiment, researchers must

consider the kinds of questions they will ask participants or the types of behaviors they will observe in order to measure their outcomes of interest. Presumably, researchers choose particular questions or behaviors because they enable them to test a hypothesis of interest. Although some of the registries discussed above do not ask for all of the items listed in Table 18.1, the preregistration of such information allows others to observe how researchers planned to measure their key outcomes of interest, as well as the connection between their measures and their theories or hypotheses. This, in turn, sheds light on what analyses and findings were specified *ex ante* versus discovered *ex post* as part of exploratory research.

Other items in Table 18.1 are things that, absent preregistration, researchers may have overlooked when planning their experiment. For example, researchers may have previously made decisions about how to handle missing data and/or how to define and detect noncompliance after they collected and began analyzing their data. They may also have neglected to develop and incorporate standard operating procedures that list default practices for handling certain types of issues, like how to handle attrition (Lin and Green 2016). By encouraging researchers to think and make decisions about these issues in advance, the preregistration process facilitates not only transparency, but also good research practices.

Another source of uncertainty in the preregistration process is how much detail to provide about particular aspects of a study. Table 18.1 seeks to clarify this by providing examples of the kinds of details that researchers should provide for each recommended item. For example, when describing the measurement of the constructs used in their study, researchers should not only describe how they measured their dependent variable(s), but also any moderators, mediators, or other covariates they plan to use. When describing their data source, researchers should include specific information about the nature of the sample (e.g., college students versus adults; a national sample versus a state or local

Table 18.1 Recommended information for preregistration and pre-analysis plans in experimental research.

<i>Type of information</i>	<i>Recommended details</i>
Research question	Convey the main question the study is designed to address
Existing literature/theories	Brief summary of existing literature/theories that relate to the research question
Hypotheses	State predictions that are clearly linked to/derived from existing literature/theories
Data source	Provide information about the nature of the sample, how participants will be recruited, the intended sample size, the stopping rule for data collection, and justifications for the intended sample size (e.g., power analyses)
Experimental design	Summarize the number and nature of the treatment/control groups; also provide the exact materials and instructions used in the experiment
Measurement	Describe how the constructs stated in the hypotheses will be operationalized; specify not only the measurement of the dependent variable(s), but also any moderators, mediators, or other covariates
Data analysis plan	Specify how exactly the data will be used to test the hypotheses, including planned statistical tests, how missing values will be treated, and whether covariates will be included (and if so, which covariates); describe the set of results that would support (or refute) the stated hypotheses (see also Blair et al. 2019)
Plan for handling noncompliance	Describe how noncompliance is defined and detected if relevant (e.g., via manipulation checks, attention screeners, reaction times); specify how noncompliance and/or attrition will be handled in the data analyses
Standard operating procedures	Refer to standard operating procedures that state default practices for handling certain types of issues, like how to handle attrition, how to define noncompliance, and whether to exclude subjects who state that they discerned the purpose of the experiment (Lin and Green 2016)

sample), how participants will be recruited (e.g., email, flyers, telephone), the intended sample size, the stopping rule governing when data collection will cease, as well as any power analyses or considerations that were taken into account. In general, the information that researchers provide during the preregistration process should be detailed enough to allow others to compare what the researcher planned to do with what the researcher actually did in his or her published research (Ledgerwood 2019).

18.2.2 Reporting Recommendations

In contrast to preregistration, there are well-developed guidelines and norms in political science for what to report in published experimental research. For example, researchers are now expected to explain and often include the exact materials

and instructions used in their experiments when submitting a manuscript to an academic journal. They are also expected to publish those experimental materials and instructions in an appendix (typically online) upon publication of the article. As discussed above, Gerber et al. (2014) present official guidelines for what should be reported in published experimental research in political science. The 18-item checklist that Gerber et al. (2015) subsequently developed based on the official guidelines is particularly useful for helping researchers determine what to report about their laboratory, survey, or field experiments in their published articles and appendices. As such, this checklist is provided in Table 18.2.

An inspection of Gerber et al.'s (2015) checklist reveals that a good place to start when reporting the details of experimental research is the information provided during

Table 18.2 Gerber et al.'s (2015) checklist of reporting items for experimental research.*Items to report*

Eligibility and exclusion criteria for participants
Details of recruitment and selection of participants, including incentives and any firms used
Type of experiment (lab, survey, field), mode, location, and dates conducted
Response rate or other participation metric (and how calculated), when possible
Details of randomization procedure
Baseline means and standard deviations for demographics and other pretreatment measures by experimental group
Whether blinding took place and how it was accomplished
Description of the treatment(s), as well as description of the control group
Details of experiment: its duration, number of participants, within- versus between-subject design, piggybacking/ordering/repetition of treatments, use of deception, use of incentives
Evidence treatment was delivered as intended, if available
Definitions of outcome measures and covariates, as well as noting whether the level of analysis differs from the level of randomization
Identification of analyses specified <i>ex ante</i> versus <i>ex post</i> exploratory analyses
Information in CONSORT participant flow diagram
Sample means and standard deviations for outcome variables using intent-to-treat analysis
Patterns of missing data, attrition, and methods of addressing these issues if missing data and/or attrition are present
Description of weighting procedures, if used
Institutional review board approval, preregistration, source of funding, conflicts of interest
Availability of replication materials and data set

the preregistration process. Indeed, many of the items in the checklist are also listed in Table 18.1 as information that researchers should preregister before conducting their experiments. For example, Tables 18.1 and 18.2 both recommend that researchers provide information about the nature of their sample, how participants were recruited, and the details of their experiment. Both Tables 18.1 and 18.2 also recommend providing information about the measurement of outcome variables and covariates. In addition, the item in Table 18.2 that recommends reporting whether particular analyses were specified *ex ante* or conducted *ex post* is something that stems directly from the preregistration process. In this way, the preregistration process actually facilitates the reporting process.

Other items in Gerber et al.'s (2015) checklist will not have been provided during the preregistration process because they are only knowable after the data have been collected. Examples include the dates on which the experiment was conducted,

response rate, patterns of missing data and attrition, evidence that the treatment was delivered as intended, as well as sample means and standard deviations for outcome variables using intent-to-treat analysis. These items are important to report in published experimental research because they allow others to assess the internal validity of an experiment and the generalizability of its results. For instance, if a researcher included a manipulation check, reporting the results of it can shed light on whether participants actually took up the treatment (see Chapter 12 in this volume, as well as Kane and Barabas 2019). This helps others to evaluate whether the treatment was delivered as intended and hence the researcher's claims about cause and effect. Further, reporting information about the response rate, the extent of missing data, and attrition allows others to evaluate how broadly the results of an experiment should be generalized. If the response rate is quite low or if a large percentage of participants dropped out before the completion of a study, this may indicate that the sample is

not representative of the broader population to which the researcher seeks to generalize.

In a similar manner, reporting the exact location and dates of an experiment provides information about the context in which the experiment took place. Such contextual information is important for assessing the generalizability of the results and evaluating subsequent replication efforts. For example, the researcher conducting an experiment on the effects of campaign advertisements on citizens' perceptions of candidates' traits might obtain different results if the experiment is conducted during an actual election campaign versus at some other time. If the experiment is conducted during an actual campaign, the amount of time between the implementation of the experiment and Election Day might also affect the results. If the researcher provides the location and dates of the experiment (and ideally a description of the political environment in which the experiment took place) in his or her published research, others can take this information into account when assessing the extent to which the results might generalize to other settings. Further, any failures to replicate the researcher's results might stem from the necessarily different context in which the replication study took place (i.e., it is not possible to recreate the particular election in which the original study took place). In this way, information about the location and dates of an experiment, as well as the political context at the time of the experiment, can also aid in the evaluation of subsequent replication efforts.

In addition to the items listed in Gerber et al.'s (2015) checklist, it is helpful if researchers can provide information about how the demographics of their sample compare to the broader population to which they seek to generalize (see Mutz and Pemantle 2015). This can shed light on whether a low response rate or substantial attrition is problematic for generalizability purposes (see Chapters 9 and 21 in this volume). Consider the following example: a researcher administers his or her experiment to a sample of registered voters in a particular city. The researcher obtains his or her

sample by contacting a random sample of voters listed in the city's master voter file. Only a small percentage (say 5%) responds to the researcher's invitation to participate in the experiment and actually does so. Needless to say, this low response rate raises questions about whether and to what extent the researcher's findings can be generalized to the city's voting population. Thus, an examination of how the demographic characteristics of the researcher's sample compare to those of registered voters in (or even residents of) the city is informative. Information about the latter can be obtained from census data or other sources. While not perfect and limited to the observable characteristics included in both the sample and population data sets, this sort of comparison at least sheds some light on the generalizability of the researcher's findings. It can also help researchers decide whether to weight the data from their study (see Chapter 21 in this volume and Miratrix et al. 2018). Reporting this information, along with the other recommended reporting items, is also important for meta-analyses of research on particular topics.

18.3 Conclusion

At a time when the credibility of social science research has been tarnished by widely discussed replication failures, insufficient reporting, and inadequate data sharing, the importance of transparency in experimental research has never been greater. However, because transparency is a public good, institutions that create incentives for individual researchers to be open about the details of their experiments are necessary. In this chapter, I discussed three institutions designed to do so: (1) preregistration and pre-analysis plans, (2) reporting guidelines, and (3) the DA-RT initiative. Together, these institutions promote transparency in all aspects of experimental research, from its design and implementation to its analysis and ultimate publication in academic journals.

Though necessary, these institutions are not sufficient for increasing transparency

in experimental and other types of social science research. Researchers also need an understanding of what exactly to preregister and report about their experimental research. Thus, I offered specific recommendations about the kinds of information researchers should preregister and ultimately report in their published articles and appendices. While there are other kinds of information that researchers may want to preregister and report, my hope is that these recommendations provide researchers who wish to be more open about their research with a template or starting point for doing so.

Incentives that encourage researchers to make greater use of these institutions (preregistration, in particular) are also important. Unlike reporting guidelines and data access requirements, preregistration is almost entirely optional in political science.¹³ Given that it takes time and effort to preregister a study, what might induce researchers to pay these costs? In recent years, academic journals and scholars have developed different kinds of incentives that might encourage researchers to preregister their studies. For example, journals like the *Journal of Experimental Political Science* and *Comparative Political Studies* have allowed for the preacceptance of preregistered reports. This allows researchers who preregistered their studies to submit them for results-blind review and to gain acceptance based on the quality of the theory and experimental design, regardless of whether null results are ultimately obtained. Other journals print badges on articles that confirm that they were produced in an open manner (e.g., preregistered, data shared).¹⁴ Scholarly organizations and foundations have also tried to create incentives by sponsoring competitions

that award cash prizes to excellent studies that were preregistered (e.g., the Center for Open Science's Preregistration Challenge and the Election Research Preacceptance Competition sponsored by the Laura and John Arnold Foundation). Learning to follow these procedures in graduate school so that they become ingrained research habits is also important (see Druckman et al. 2018).¹⁵ So, too, is recognizing that an additional benefit of preregistration is that it enables researchers to lay claim to their ideas, theories, and experimental tests at the outset of their research.

Even with these incentives, scholarly misgivings about preregistration remain a barrier to the widespread use of this institution. In contrast to the relatively broad agreement about reporting and data-sharing principles and practices (at least within the experimental research community in political science), there have been a number of arguments made against the preregistration process. One argument is that fully prespecifying social science research is close to impossible, given the inherent complexity of the social world. In this view, an attempt to anticipate and prespecify all aspects of such research may lead to unwieldy pre-analysis plans at best and misplaced focus on the part of researchers at worst (Olken 2015). With respect to the latter, the concern is that preregistration may induce researchers to focus on simple, straightforward hypotheses at the expense of conducting more nuanced exploratory analyses that allow researchers to learn from the data they collected (Coffman and Niederle 2015; Olken 2015). A separate concern is that funding limitations in the social sciences often prevent initial exploratory results from later being tested in a confirmatory (i.e., *ex ante*) way. Thus, the exploratory results are often all that the scholarly community has to go on, and there is a risk that they will be downgraded in

13 An exception is development research produced as part of EGAP, where preregistration appears to be the norm.

14 Research suggests that "open data" and "open materials" badges effectively promote transparency. Specifically, Kidwell et al. (2016) find that after the journal *Psychological Science* began offering these badges, there was a considerable increase in the percentage of articles that reported open data. They also find that the data that researchers made available were more likely to be usable and complete when badges were earned relative to when badges were not earned.

15 To the extent that academic departments and universities begin to take research transparency into account in hiring and/or tenure decisions, this creates additional incentives for junior scholars to make use of these institutions.

importance or not produced in the first place (Olken 2015).

While these are valid concerns, it is important to note that they pertain to the practice, not principle, of preregistration. That is, these arguments do not deny the value of greater transparency or the role of preregistration in facilitating it. Rather, they highlight potential negative consequences that might arise from the actual practice of preregistration. Some of these consequences (e.g., unwieldy pre-analysis plans) can be avoided with the development of clear guidelines about what exactly to include during the preregistration process. Others, such as the concern that exploratory research might be viewed as less important than confirmatory research, are more challenging to address because they involve perceptions of the relative value of these two types of research. Ultimately, journal policies and scholarly norms that encourage researchers to report not just their confirmatory analyses but also their exploratory ones in published research may help put these two types of research on more equal footing.

References

- Aczel, Balazs, Barnabas Szasz, Alexandra Sarafoglou, Zoltan Kekecs, Šimon Kucharský, Daniel Benjamin et al. 2020. "A Consensus-Based Transparency Checklist." *Nature Human Behavior* 4(1): 4–6.
- Appelbaum, Mark, Harris Cooper, Rex B. Kline, Evan Mayo-Wilson, Arthur M. Nezu, and Stephen M. Rao. 2018. "Journal Article Reporting Standards for Quantitative Research in Psychology: The APA Publications and Communications Board Task Force Report." *American Psychologist* 73(1): 3–25.
- Blair, Graeme, Jasper Cooper, Alexander Coppock, and Macartan Humphreys. 2019. "Declaring and Diagnosing Research Designs." *American Political Science Review* 113(3): 838–859.
- Coffman, Lucas C., and Muriel Niederle. 2015. "Pre-analysis Plans Have Limited Upside, Especially Where Replications are Feasible." *Journal of Economic Perspectives* 29(3): 81–98.
- Coffman, Lucas C., Muriel Niederle, and Alis-tair J. Wilson. 2017. "A Proposal to Organize and Promote Replications." *American Economic Review* 107(5): 41–45.
- Connors, Elizabeth C., Yanna Krupnikov, and John Barry Ryan. 2019. "How Transparency Affects Survey Responses." *Public Opinion Quarterly* 83(S1): 185–209.
- Coppock, Alexander. 2019. "Generalizing from Survey Experiments Conducted on Mechanical Turk: A Replication Approach." *Political Science Research and Methods* 7(3): 613–628.
- Druckman, James N., Adam J. Howat, and Kevin J. Mullinix. 2018. "Graduate Advising in Experimental Research Groups." *PS: Political Science and Politics* 51(3): 620–624.
- Elman, Colin, and Diana Kapiszewski. 2014. "Data Access and Research Transparency in the Qualitative Tradition." *PS: Political Science and Politics* 47(1): 43–47.
- Franco, Annie, Neil Malhotra, and Gabor Simonovits. 2014. "Publication Bias in the Social Sciences: Unlocking the File Drawer." *Science* 345(6203): 1502–1505.
- Franco, Annie, Neil Malhotra, Gabor Simonovits, and L. J. Zigerell. 2017. "Developing Standards for Post-Hoc Weighting in Population-Based Survey Experiments." *Journal of Experimental Political Science* 4: 161–172.
- Friedman, Lisa. 2018. "The E.P.A. Says It Wants Research Transparency. Scientists See an Attack on Science." *The New York Times*, March 26.
- Gerber, Alan, Kevin Arceneaux, Cheryl Boudreau, Conor Dowling, Sunshine Hillygus, Thomas Palfrey et al. 2014. "Reporting Guidelines for Experimental Research: A Report from the Experimental Research Section Standards Committee." *Journal of Experimental Political Science* 1(1): 81–98.
- Gerber, Alan S., Kevin Arceneaux, Cheryl Boudreau, Conor Dowling, and D. Sunshine Hillygus. 2015. "Reporting Balance Tables, Response Rates and Manipulation Checks in Experimental Research: A Reply from the Committee that Prepared the Reporting Guidelines." *Journal of Experimental Political Science* 2(2): 216–229.
- Ioannidis, John P. A. 2014. "How to Make More Published Research True." *PLoS Medicine* 11(10): e1001747.
- Kane, John V., and Jason Barabas. 2019. "No Harm in Checking: Using Factual Manipulation Checks to Assess Attentiveness in Experiments." *American Journal of Political Science* 63(1): 234–249.
- Kidwell, Mallory C., Ljiljana B. Lazarevic, Erica Baranski, Tom E. Hardwicke, Sarah Piechowski, Lina-Sophia Falkenberg et al. 2016. "Badges to Acknowledge Open

- Practices: A Simple, Low-Cost, Effective Method for Increasing Transparency.” *PLoS Biology* 14(5): e1002456.
- Ledgerwood, Alison. 2019. “New Developments in Research Methods.” In *Advanced Social Psychology: The State of the Science*, 2nd ed., eds. Eli J. Finkel and Roy F. Baumeister. Oxford: Oxford University Press, pp. 39–61.
- Levitt, Heidi M., Michael Bamberg, John W. Creswell, David M. Frost, Ruthellen Josselson, and Carola Suarez-Orozco. 2018. “Journal Article Reporting Standards for Qualitative Primary, Qualitative Meta-Analytic, and Mixed Methods Research in Psychology: The APA Publications and Communications Board Task Force Report.” *American Psychologist* 73(1): 26–46.
- Lin, Winston, and Donald P. Green. 2016. “Standard Operating Procedures: A Safety Net for Pre Analysis Plans.” *PS: Political Science and Politics* 49(3): 495–500.
- Lupia, Arthur. 2008. “Procedural Transparency and the Credibility of Election Surveys.” *Electoral Studies* 27: 732–739.
- Lupia, Arthur, and Colin Elman. 2014. “Openness in Political Science: Data Access and Research Transparency.” *PS: Political Science and Politics* 47(1): 19–42.
- McDermott, Rose. 2014. “Research Transparency and Data Archiving for Experiments.” *PS: Political Science and Politics* 47(1): 67–71.
- Miguel, E., C. Camerer, K. Casey, J. Cohen, K. M. Esterling, A. Gerber et al. 2014. “Promoting Transparency in Social Science Research.” *Science* 343(6166): 30–31.
- Miratrix, Luke W., Jasjeet S. Sekhon, Alexander G. Theodoridis, and Luis F. Campos. 2018. “Worth Weighting? How to Think About and Use Weights in Survey Experiments.” *Political Analysis* 26(3): 275–291.
- Mullinix, Kevin J., Thomas J. Leeper, James N. Druckman, and Jeremy Freese. 2015. “The Generalizability of Survey Experiments.” *Journal of Experimental Political Science* 2: 109–138.
- Mutz, Diana C., and Robin Pemantle. 2015. “Standards for Experimental Research: Encouraging a Better Understanding of Experimental Methods.” *Journal of Experimental Political Science* 2(2): 192–215.
- Nosek, Brian A., Jeffrey R. Spies, and Matt Motyl. 2012. “Scientific Utopia: II. Restructuring Incentives and Practices to Promote Truth Over Publishability.” *Perspectives on Psychological Science* 7(6): 615–631.
- Nosek, B. A., G. Alter, G. C. Banks, D. Borsboom, S. D. Bowman, S. J. Breckler et al. 2015. “Promoting an Open Research Culture.” *Science* 348: 1422–1425.
- Nosek, Brian A., Charles R. Ebersole, Alexander C. DeHaven, and David T. Mellor. 2018. “The Preregistration Revolution.” *Proceedings of the National Academy of Sciences of the United States of America* 115(11): 2600–2606.
- Olken, Benjamin A. 2015. “Promises and Perils of Pre-Analysis Plans.” *Journal of Economic Perspectives* 29(3): 61–80.
- Open Science Collaboration. 2015. “Estimating the Reproducibility of Psychological Science.” *Science* 349: aac4716.
- Palfrey, Thomas, and Robert Porter. 1991. “Guidelines for Submission of Manuscripts on Experimental Economics.” *Econometrica* 59: 1197–1198.
- Simmons, Joseph. P., Leif D. Nelson, and Uri Simonsohn. 2011. “False-positive Psychology: Undisclosed Flexibility in Data Collection and Analysis Allows Presenting Anything as Significant.” *Psychological Science* 22(11): 1359–1366.