

Political Knowledge

Cheryl Boudreau
Department of Political Science
University of California, Davis
One Shields Ave
Davis CA 95616
E-mail: clboudreau@ucdavis.edu

Arthur Lupia
4252 Institute for Social Research
426 Thompson Street
University of Michigan
Ann Arbor, MI 48104
E-mail: lupia@umich.edu

FORTHCOMING IN

The Cambridge Handbook of Experimental Political Science

James N. Druckman, Donald P. Green, James H. Kuklinski, and Arthur Lupia (eds.)
Cambridge University Press, 2011

Word count: 7500

Version 2.0. Prepared for the *Cambridge Handbook of Experimental Political Science*.

August 10, 2009. We thank Adam Seth Levine, Yanna Krupnikov, and John Bullock for helpful comments.

In many political surveys, citizens fail to answer, or provide incorrect answers to, fact-based questions about political figures and institutions. Such findings raise concerns about citizen decision making. A common claim is that citizens' poor performance on surveys reflects their incompetence in democratically meaningful contexts such as voting booths.

The scholarly home for such findings and concerns is the academic literature on political knowledge. A common analytic definition of political knowledge is that it is a measure of a citizen's ability to provide correct answers to a specific set of fact-based questions. Typical questions include "What is the political office held by [name of current vice president, British prime minister, or Chief Justice of the United States]?" and "Which political party has the most seats in the U.S. House of Representatives?"

Many people have used such data to conclude that the public is generally ignorant about politics and makes incompetent decisions as a result. In recent years, however, scholars have asked increasingly penetrating questions about the validity and meaning of such claims (e.g., Graber 1984, Popkin 1994). Some raised questions about the practice of basing broad generalizations of citizen knowledge or voter competence on a relatively small set of idiosyncratic, fact-based survey questions (e.g., Lupia 2006). Others uncovered logical and factual errors in claims about the kinds of political knowledge that are necessary, sufficient, or even related to the ability to make important democratic choices competently (e.g., Lupia and McCubbins 1998, Gibson and Caldeira 2009).

A common theme in these critiques is that many claims about political knowledge are built on vague or erroneous conjectures about two central questions:

1. What is political knowledge?

2. What is the relationship between political knowledge, as defined above, and task-related competence, by which we mean the ability to make the choice one would have made if knowledgeable about a set of relevant facts?

Many scholars simply presumed that the answer to the first question was obvious: survey-based political knowledge measures are valid representations of citizens' general knowledge of politics. With respect to the second question, many writers also presumed that not offering correct responses to these survey questions was equivalent to incompetence. Given how often this literature criticized citizens for what they did not know, it is ironic that it included so little introspection about the conditions under which these presumptions were true.

Experimental political science has added clarity, precision, and new insight into such matters. Unlike non-experimental studies, experiments allow scholars to 1) randomly assign subjects to treatment and control groups, 2) systematically manipulate aspects of survey-interview and decision-making environments, and 3) directly observe the answers subjects give (and the choices they make) under different conditions. These features of experiments have enabled scholars to clarify how particular aspects of the survey production process contribute to citizens' poor performance on survey-based political knowledge questions. Experiments have also clarified the relationship between the ability to answer certain questions and the ability to make important decisions competently. To be sure, there are many things about politics that citizens do not know. But experiments show that what citizens actually know about politics, and how such knowledge affects their choices, is very different than the conventional wisdom alleged.

In this chapter, we report on experiments that address the two questions specified above. In section 2, we address "what is political knowledge" by describing experiments

that manipulate the survey context from which most political knowledge measures are derived. These experiments reveal that existing knowledge measures are significantly affected by question wording, variations in respondents' incentives to think before answering, whether respondents feel threatened by unusual aspects of survey interview contexts, and personality variations that make some respondents unwilling to give correct answers to survey interviewers even when they are knowledgeable about the subject matter.

In section 3, we describe experiments that clarify the relationship between what citizens know and their competence. These experiments compare the choices that people make, or the opinions that they state, given different kinds of information. They reveal conditions under which people can (and cannot) make competent decisions despite lacking answers to fact-based political knowledge questions. Collectively, these clarifications provide a different view of citizen competence than is found in most non-experimental work on political knowledge. There are many cases for which not knowing the answers to survey-based political knowledge questions reveals very little about citizens' competence.

Experiments on Properties of Survey-Based Knowledge Measures

There is no denying that many citizens do not correctly answer political knowledge questions on traditional surveys. But it is not clear that these survey questions are valid measures of what citizens actually know about politics. Non-experimental studies have attempted to validate survey-based political knowledge measures in one of two ways.

Some have tried to demonstrate their validity by correlating them with factors such as ideology, interest in politics, education, and turnout. This approach is problematic because these factors are not themselves credible measures of political knowledge. For example, a person can be an ideologue without being knowledgeable and can be knowledgeable without being ideological.

Others have based validity claims on factor analysis. The argument they make is that if the same kinds of people can answer various survey questions correctly, then the questions must be effectively tapping a more general knowledge domain. The error in this logic follows from previous critiques of factor analysis in intelligence estimation. As Gould describes (1996: 48) "the key error of factor analysis lies in reification, or the conversion of abstractions into putative real entities." In other words, factor analysis yields meaningful results only if the selection of questions themselves is derived from a credible theory of information and choice. Lupia (2006: 224) finds that the selection of specific questions is subjective and not easily quantified. Hence, such attempts are of little relevance to the question of whether survey-based political knowledge questions capture citizens' true knowledge about politics.

Gibson and Caldeira (2009) offer an experiment that both questions the validity of existing data and provides a more effective means of measuring what citizens know. They argue (2009: 429) that "much of what we know — or think we know— about public knowledge of law and courts is based upon flawed measures and procedures." Their experiment clarifies how an important aspect of question format affects survey-based political knowledge measures.

In particular, many of the most famous survey-based political knowledge measures come from open-ended recall questions. An example of such a question, from the American National Election Studies (ANES), is as follows:

“Now we have a set of questions concerning various public figures. We want to see how much information about them gets out to the public from television, newspapers and the like.... What about ... William Rehnquist – What job or political office does he NOW hold?”

Open-ended questions allow respondents to answer in their own words.

From the 1980s through 2004, the ANES hired coders to convert transcribed versions of respondents’ verbatim answers into a simple binary variable. All responses were coded as “correct” or “incorrect.” The codes, and not the original responses, were then included in the public ANES dataset. For decades, scholars treated the codes as valid measures of respondents’ knowledge. While many analysts used this data to proclaim voter ignorance and incompetence, an irony is that almost no one questioned whether the measure itself was valid.

Gibson and Caldeira raised important questions about which responses should be counted as correct. In 2004, for example, William Rehnquist was Chief Justice of the United States. Upon inspecting the transcribed versions of ANES responses, Gibson and Caldeira found errors in coding and in which responses were coded as "correct." The ANES counted as correct only responses that included “Chief Justice” *and* “Supreme Court.” A respondent who said that Rehnquist is on the Supreme Court without saying Chief Justice or a respondent who simply said that he was a federal judge were coded as "incorrect."

Many analysts interpreted such codes as meaning that the respondent knew nothing about the subject matter. Gibson and Caldeira examined the transcriptions and found that they often showed at least partial knowledge of the topic. Hence, they argued (2009: 429) that past practices likely produced "a serious and substantial underestimation of the extent to which ordinary people know about the nation's highest court."

To assess the extent of this underestimation, Gibson and Caldeira embedded an experiment in a nationally drawn telephone survey of 259 respondents. Respondents were asked to identify the current or most recent political office held by William Rehnquist, John G. Roberts, and Bill Frist. A control group was asked these questions in the traditional open-ended format. A treatment group was asked to identify the same individuals in a multiple choice format.

With respect to Chief Justice Rehnquist, and using the traditional ANES method of scoring open-ended responses as "correct" or "incorrect," 12% of respondents correctly identified Rehnquist as Chief Justice. Another 30% identified him as a Supreme Court justice, but since these responses did not explicitly refer to him as Chief Justice, the standard ANES measure would have counted these responses as incorrect. The treatment group, by contrast, was asked to state whether Rehnquist, Lewis F. Powell, or Byron R. White was Chief Justice (with the order of the response options randomized across respondents). When asked the question in this format, 71% correctly selected Rehnquist. Gibson and Caldeira observed comparable results for Bill Frist and John Roberts.

The substantive impact of Gibson and Caldeira's (2009: 430) findings is that "[T]he American people know orders of magnitude more about their Supreme court than most other studies have documented." The broader methodological implication is that the

combination of open-ended questions, ANES's coding scheme, and a lack of fact-checking by critics of citizen knowledge contributed to an overly negative image of public knowledge. Gibson and Caldeira's work subsequently caused the ANES to restructure how it solicits and codes political knowledge (Krosnick et al. 2008). Hence, in this case, an experimental design not only influenced our understanding of political knowledge, but also improved how the concept is now measured.

Other experiments on survey interview attributes suggest further trouble for conventional interpretations of traditional political knowledge measures. Prior and Lupia (2008: 169) ask whether “seemingly arbitrary features of survey interviews” affect the validity of knowledge measures. They contend that the typical survey-based political knowledge assessment occurs in an unusual circumstance. Interviewers have incentives to complete interviews quickly. Respondents often do not want to prolong the interview. Hence, questions are asked, and answers are expected, in quick succession. Moreover, political knowledge questions typically appear in the survey with no advance notice. This “pop quiz” atmosphere is very different than circumstances in which having particular kinds of political knowledge matter most, such as elections. Election dates are typically known in advance. Hence, people who wish to become informed have an opportunity to do so before they cast a vote. The typical survey also provides no incentive for respondents to answer the questions correctly.

To determine the extent to which odd survey interview attributes cause poor performance on political knowledge quizzes, Prior and Lupia assigned over 1200 randomly selected members of Knowledge Networks' national internet panel to one of four experimental groups. The control group was asked fourteen political knowledge

questions in a typical survey interview environment, with little time to answer the questions (60 seconds from the time that the question first appeared on screen) and no motivation to answer correctly. Treatment groups received greater opportunity and/or incentive to engage the questions. One treatment group was offered one dollar for every question answered correctly. Another group was offered 24 hours to respond to the fourteen questions. The third treatment group was offered time and money.

Each of the treatments produced a significant increase in questions answered correctly. Simply offering a dollar for correct answers increased the average number of correct answers by 11%. Offering extra time produced an 18% increase. Time and money together increased the average by 24%.

The effect of money alone is noteworthy as the only difference between the control and treatment groups is that the latter is paid a small amount for each correct answer. Treatment group respondents did not have time to look up correct answers. Hence, the treatment group's performance gain indicates that respondent motivation is a determinant of existing political knowledge measures.

Looking at experimental effects across population groups reinforces the conclusion. The largest effects are on respondents who report that they follow politics "some of the time" (rather than "most of the time" or "not at all"). For them, simply paying a dollar yields a 32% increase in correct answers. Hence, for people whose attention to politics is infrequent, the typical survey interview context provides insufficient motivation for searching the true content of their memories in response to conventional political knowledge questions. This finding (2008: 169) implies that "conventional knowledge measures confound respondents' recall of political facts with

variation in their motivation to exert effort during survey interviews [and] provide unreliable assessments of what many citizens know when they make political decisions.”

Other experiments examine how social roles and survey contexts interact to affect respondent performance. McClone, Aronson, and Kobrynowicz (2006) noted that men tend to score better than women on survey-based political knowledge tests. Conventional explanations for this asymmetry included the notion that men are more interested in politics. McClone et al. saw "stereotype threat" (Steele and Aronson 1995) as an alternative explanation. In a typical stereotype threat experiment, members of stigmatized and non-stigmatized groups are randomly assigned to experimental groups. A treatment group is given a test along with a cue suggesting that members of their stigmatized group have not performed well on previous evaluations of this kind. The control group receives the test without the cue. While most experimental work demonstrated threat effects on African-American subjects, there was growing evidence that women were also affected by stereotype threat.

McClone et al.'s phone-based experiment was conducted on 141 undergraduates, 70 men and 71 women. A ten-question index measured political knowledge. Questions were drawn from sources such as the National Election Studies. Stereotype threat was manipulated in two ways. The first manipulation was interviewer gender. Respondents were randomly assigned to be interviewed by men or women. The second manipulation pertained to the description of the political knowledge questions. Treatment groups were told that "the survey you are participating in this evening has been shown to produce gender differences in previous research." Control groups heard that "the survey you are

participating in this evening has not been shown to produce any gender differences in previous research whatsoever."

Men scored higher than women overall, which is consistent with non-experimental findings. However, this asymmetry was reliably moderated by the experimental factors. When the political knowledge questions were portrayed as non-diagnostic of women's abilities, or when women were interviewed by other women, *there was no significant difference* between men's and women's scores. When diagnosticity or male interviewers were introduced, the asymmetry emerged but neither factor affected men's scores. *This effect was confined to women*. Hence, gender-related factors affected political knowledge scores. This experiment suggests important limits on the extent to which survey-based political knowledge tests can be interpreted as valid measures of what the population as a whole knows about politics.

Other experiments show that whether survey questions allow or encourage "don't know" (henceforth, DK) responses affects political knowledge measures. These experiments suggest that DK options may cause scholars to underestimate what the public knows about politics because some people are less likely than others to offer answers when they are uncertain. For example, Jeffrey Mondak and his colleagues designed several split-ballot experiments (i.e., random assignment of survey respondents to experimental conditions) in two surveys, the 1998 National Election Studies Pilot Study and a Tallahassee-based survey. In each survey, the control group received knowledge questions that began with a DK-inducing prompt, "Do you happen to know..." and interviewers were instructed not to probe further after a DK response. Treatment groups received questions with identical substantive content, but different

implementation. In both surveys, each question included a guess-inducing phrase such as "even if you're not sure I'd like you to tell me your best guess." In the National Election Studies (NES) version, moreover, the interviewer first recorded any DK responses and then probed further for substantive answers to determine whether respondents who initially responded DK actually knew about the concept in question.

In each experiment, respondents were significantly less likely to choose DK when they were encouraged to guess than when DK was encouraged. Interviewer probing decreased DKs even further (Mondak 2001). Another analysis of the effect of discouraging DKs showed women significantly more likely than men to respond DK even when encouraged to guess (Mondak and Anderson 2004). In this analysis, discouraging DKs reduces the extent to which men outperform women (in terms of questions answered correctly) by about half.

For our purposes, the most interesting finding from this work is that many respondents chose DK for reasons other than ignorance. Mondak and Davis (2001) analyzed the responses to interviewer probes that respondents offered after they initially chose DK. These responses were significantly more likely to be correct than responses that would have emerged from "blind guessing." Taken together, Mondak and his colleagues show that whether surveys encourage or discourage DK responses substantially affects what traditional political knowledge questions actually measure. This work reveals that many previous measures confound what respondents know with how willing they are to answer questions when they are uncertain.

Building on these experiments, Miller and Orr (2008) designed an experiment where the option of responding DK was not just discouraged, it was eliminated altogether.

It was run on 965 undergraduates via internet surveys. Each respondent received eight multiple-choice political knowledge questions and each question contained three substantive response options. What differed across their experimental groups was question format. The first group's questions encouraged DKs. The second group's questions discouraged DKs. For the third group, DKs were unavailable. Miller and Orr found that discouraging DKs (rather than encouraging them) led to a substantial drop in the use of the DK option. They also found that discouraging DKs (rather than encouraging them) corresponded to an increase in the average percentage of correct answers given per respondent.

When comparing the DK-omitted group's responses to those of the DK-encouraged and the DK-discouraged, interesting patterns emerged. The "average percent correct" was highest for the DK-omitted group, as was the "average percent incorrect." The most interesting thing about the comparison between the DK-encouraged and DK-omitted groups is that the increase in percent correct (from 61% to 70%) was higher than the increase in percent incorrect (from 21% to 29%). This is interesting because each question had *three* response options. Hence, if the DK-encouraged and DK-omitted groups were equivalent, and if all that omitting DK options does is cause respondents to guess haphazardly, then respondents who would have otherwise chosen DK should have only a *one-in-three* chance of answering correctly. Hence, if respondents were simply guessing, then the increase in "average percent incorrect" should roughly double the increase in "average percent correct." Instead, the increase in corrects was larger than the increase in incorrects. Miller and Orr's experiment shows that the DK option attracts not just respondents who lack knowledge about the questions' substance but also people who

possess relevant information but are reticent to respond for other reasons (such as lack of confidence or risk aversion).

While Miller and Orr's work suggests that DK responses hide partial knowledge, research by Sturgis, Allum, and Smith (2008) suggests a different conclusion—that DK responses conceal little knowledge. Sturgis et al. integrated a split ballot experiment into a British telephone survey. Each respondent was asked three knowledge-related questions. Each question contained a statement, and the respondent was asked to say whether it was true or false. 1006 respondents were randomly assigned to one of three conditions. In one condition, the question's preamble included the DK-encouraging phrase, "If you don't know, just say so and we will skip to the next one." In a second condition, that phrase was substituted with the DK-discouraging phrase, "If you don't know, please give me your best guess." In the third condition, the original DK-encouraging statement was included but the response options were changed. Instead of simply saying "true" or "false," respondents in this group could say whether the statement was "probably true," "definitely true," "probably false," or "definitely false." Moreover, in conditions 1 and 3, respondents who initially said DK were later asked to provide their "best guess."

Sturgis et al. find that discouraging DKs significantly decreased their frequency (from 33% to 9%). Providing the "definitely" and "probably" response options also reduced DKs (to 23%). Turning their attention to partial knowledge, they then analyze the answers given by respondents who initially responded DK and then chose true or false after interviewer probing. For two of the three questions, probing elicited correct answers at a rate no better than chance. For the other question, two-thirds of the new responses elicited were correct. From these results, they conclude that "when people who

initially select a DK alternative are subsequently asked to provide a ‘best guess,’ they fare statistically no better than chance.” But their results suggest a different interpretation – that there are types of people and question content for which encouraging DKs represses partial knowledge and that there is still a great deal that we do not know about the types of questions and people that make such outcomes more or less likely. Further, with only three true/false questions, their experiment does not provide a sufficient basis for privileging their conclusions over those of Miller and Orr. As Miller and Orr (2008:779) note, “The availability of three options from which to choose [the Miller-Orr method] may motivate respondents to draw on their partial knowledge, whereas the true/false format might not.”

For decades, surveys have been used to produce measures of political knowledge. The modal finding is that many members of the public do not provide correct answers to the kinds of questions that political scientists place on surveys. Some respondents offer incorrect answers. Others respond DK.

Many scholars and analysts have used this data to draw broad conclusions about public ignorance. At this point, few disagree that there is a lot that the public does not know. But what these findings mean is another matter altogether. Experiments on survey-based political knowledge measures have shown that many underappreciated attributes of survey interview contexts (including question wording, respondent incentives, and personality variations) are significant determinants of past outcomes. Hence, as a general matter, survey-based political knowledge measures are much less valid indicators of what citizens know about politics than many critics previously claimed.

When Do Citizens Need to Know the Facts?

In addition to shedding light on the validity of survey-based political knowledge measures, experiments also clarify the conditions under which knowing particular facts about politics is necessary, sufficient, or relevant to a citizen's ability to make competent choices. Non-experimental research on this topic has often tried to characterize the relationship between political knowledge and politically relevant behaviors or opinions with anecdotes or regression coefficients that presume a simple linear and unconditional relationship between the factors. These characterizations are never derived from direct evidence or rigorous theory. Instead, many critics have merely presumed that any fact they deem worth knowing must be a necessary condition for competent political decision making and proceeded with their claims about who is and is not competent.

With experiments, scholars can evaluate hypotheses about the relationship between specific kinds of political knowledge and competence at various politically relevant tasks. To illustrate how and why experiments are well suited for this purpose, we discuss several examples of such experiments. These examples reveal conditions under which people who lack certain kinds of information make competent choices nevertheless.

Lupia and McCubbins (1998), for example, acknowledge that many citizens lack factual knowledge about politics, but they emphasize that uninformed citizens may be able to learn what they need to know from political parties, interest groups, and the like (i.e., speakers). Lupia and McCubbins use a formal model to identify conditions under which citizens can make competent choices as a result of such interactions.

They then use experiments to evaluate some of their theoretical conclusions. In their laboratory experiments, subjects guess the outcomes of unseen coin tosses. Coin

tosses are used because they, like many elections, confront people with binary choices. They are also easy to describe to experimental subjects and, hence, permit other aspects of the informational environment to be manipulated.

Before subjects make predictions, another subject (acting as “the speaker”) makes a statement about whether the coin landed on heads or tails. Subjects are told that the speaker is under no obligation to reveal what he or she knows about the coin truthfully. After receiving the speaker’s statement, subjects predict the coin toss outcome.

To evaluate key theoretical hypotheses, Lupia and McCubbins systematically manipulate multiple aspects of the informational context, including whether subjects receive statements from a speaker who benefits when they make correct predictions, or whether they receive statements from a speaker whose interests conflict with their own. They also introduce penalties for lying and probabilistic threats that false statements will be revealed (i.e., verification). In other words, to evaluate conditions under which subjects can learn enough to make correct predictions, Lupia and McCubbins vary the financial incentives and beliefs of subjects and the speaker.

Lupia and McCubbins’s experiments show that the conditions identified by their theory do in fact lead subjects to trust the speaker’s statements and make correct decisions. Specifically, when subjects perceive a speaker as being knowledgeable and having common interests, subjects trust the speaker’s statements. When they are in conditions under which such perceptions are likely to be correct, these patterns of trusting lead subjects to make correct predictions at a very high rate -- one that is substantially greater than chance and often indistinguishable from the predictions they would have made if fully informed. Similarly, when a sufficiently large penalty for lying or

probability of verification is imposed upon the speaker, subjects trust the speaker's statements and make correct predictions at very high rates (even when it is known that the speaker would otherwise benefit if subjects make incorrect predictions about the coins).

These experiments highlight conditions under which uninformed citizens can increase their competence by learning from others. But the experiments evaluate only a few of Lupia and McCubbins's theoretical implications. One question that their experiments leave open is whether a speaker's statements are equally helpful to more and less knowledgeable citizens. On one hand, it is possible that a speaker's statements will be more helpful to citizens who have preexisting knowledge about the choices they face. On the other hand, a speaker's statements may be more helpful to less knowledgeable citizens, who might be more open to advice. Lupia and McCubbins' experiments do not address these questions because all subjects know that there is 50% chance that a fair coin will land on heads. Hence, prior to interacting with the speaker, subjects cannot be more or less knowledgeable about how to accomplish the task of correctly predicting the coin toss.

Boudreau (2009) replicates Lupia and McCubbins' experiments but substitutes math problems for coin tosses. An advantage of using math problems is that subjects vary in their levels of preexisting knowledge. Some subjects know a lot about how to solve math problems. Others do not. A second advantage is that there exists a valid, reliable, and agreed upon measure of how knowledgeable subjects are about this type of decision— SAT math scores. Thus, Boudreau collects subjects' SAT math scores prior to the experiments. She uses the experiments to clarify conditions under which a

speaker's statements about the answers to the math problems help low-SAT subjects to behave as though they are high-SAT subjects.

Boudreau finds that when the speaker is paid for subject success, is subject to a sufficiently large penalty for lying, or faces a sufficiently high probability of verification, both low-SAT and high-SAT subjects achieve large improvements in their decisions (relative to counterparts in the control group, who do the problems without a speaker). Low-SAT subjects improve so much that it reduces the achievement gap between them and high-SAT subjects. This gap closes even when the size of the penalty for lying or probability of verification is reduced (and is, thus, made more realistic because the speaker may have an incentive to lie). This result occurs because high-SAT subjects do not improve their decisions (and apparently ignore the speaker's statements), but low-SAT subjects typically improve their decisions enough to make them comparable to those of high-SAT subjects. By using an experimental task for which subjects vary in their levels of knowledge, Boudreau further clarifies the conditions under which less informed citizens can make competent choices.

One of the strengths of Boudreau's and Lupia and McCubbins's experiments is that subjects make decisions for which there is an objectively correct or incorrect choice under different conditions. This approach is advantageous because it allows them to measure precisely whether and when a speaker's statements help subjects make a greater number of correct decisions than they would have made on their own. However, what does it mean for citizens to make "correct" decisions in electoral contexts? Lau and Redlawsk (1997, 2001) address this question by conducting experiments in which subjects learn about and vote for candidates in mock primary and general elections.

Subjects in Lau and Redlawsk's experiments are provided with different types of information about fictional candidates. Subjects access this information by clicking on labels (such as "Walker's stand on defense spending") that appear on computer screens. Subjects can also learn about a candidate's partisanship, ideology, and appearance, as well as endorsements and polls. After subjects gather information about the primary election candidates (which may be two candidates or four, depending on the experiment), they vote for one of these candidates. Subjects repeat this process for the general election candidates. At the completion of the experiments, subjects receive *all* of the information that was available for two candidates from the primary election (not just the information they clicked on during the experiment). Subjects are then asked whether they would have voted for the same candidate if they had all of this information when they made their decisions.

Lau and Redlawsk (1997) find that subjects, in the aggregate, are adept at voting correctly. According to one of their measures, approximately 70% of subjects voted correctly. However, Lau and Redlawsk (2001) identified conditions that hinder subjects' ability to vote correctly. For example, they find that although heuristics significantly increase the ability to vote correctly among subjects who score high on their political knowledge and political interest index, they *decrease* less knowledgeable and less interested subjects' ability to vote correctly.¹ Lau and Redlawsk also find that characteristics of the information environment limit subjects' ability to vote correctly. Specifically, subjects are less likely to vote correctly when the number of primary candidates increases from two to four and when the choice between the candidates is

¹ The index combines subjects' levels of political knowledge, political behavior, political interest, political discussion, and media use.

more difficult (i.e., the candidates are more similar). Thus, although Lau and Redlawsk observe high levels of correct voting in the aggregate, they also show that there are conditions under which aspects of the information environment have detrimental effects on subjects' ability to make correct decisions.

Continuing Lau and Redlawsk's emphasis on the political environment, Kuklinski et al. (2001) use experiments to assess the effects that other aspects of the environment have on citizens' decisions. In contrast to Lau and Redlawsk's focus on correct voting, Kuklinski et al. contend that the ability to make tradeoffs is fundamental to being a competent citizen. Thus, they conduct survey experiments in which they measure subjects' ability to make tradeoffs among competing goals for health-care reform under different conditions.

Subjects in Kuklinski et al.'s experiments view seven different health-care goals (e.g., universal coverage, no increase in taxes, uniform quality of care), and they rate on a scale of one to ten how much of each goal a health-care plan must achieve for them to consider the plan acceptable. The key to this experiment is that the health-care goals conflict with one another; that is, no health-care plan can realistically achieve all of the goals. In various treatment groups, Kuklinski et al. manipulate the conditions under which subjects rate the seven health-care goals. In one group, subjects are given general information about the need for tradeoffs when designing any program. In a second group, subjects are given motivational instructions that encourage them to take their decisions seriously. In a third group, subjects are given both general information and motivational instructions. In a fourth group, subjects are given diagnostic information about the exact tradeoffs involved in health-care reform (e.g., that we cannot provide health coverage for

everyone and simultaneously keep taxes low). In a fifth group, this diagnostic information is provided along with motivational instructions. Kuklinski et al. then observe whether and under what conditions information and/or motivation improves subjects' ability to make tradeoffs (measured as the extent to which subjects reduce their demands for conflicting goals), relative to subjects in the control group who do not receive any information or motivational instructions.

Kuklinski et al.'s experiment reveals conditions under which new information improves subjects' ability to make tradeoffs and eliminates differences between more and less knowledgeable subjects. Specifically, where control group subjects tended not to make tradeoffs, treatment group subjects were much more likely to do so when both general information and motivational instructions were provided. Kuklinski et al. also show that diagnostic information about health-care tradeoffs induces subjects to make these tradeoffs, regardless of whether they are motivated to do so and regardless of their knowledge level. Indeed, when diagnostic information is provided, less knowledgeable subjects are just as capable of making tradeoffs as more knowledgeable subjects. In this way, Kuklinski et al. demonstrate that when information in the environment is sufficiently diagnostic, it can substitute for preexisting knowledge about politics and eliminate differences between more and less knowledgeable subjects.

Other experiments clarify the effect of policy-specific knowledge on citizens' abilities to express their opinions. Using survey experiments, Gilens (2001) suggests that policy-specific knowledge may be more relevant to such expressions than more general conceptions of political knowledge. Gilens randomly determines whether subjects receive specific information about two policy issues (crime and foreign aid). Subjects in the

treatment group receive information about two news stories, one showing that the crime rate in America has decreased for the seventh year in a row and one showing that the amount of money spent on foreign aid has decreased and is now less than one cent of every dollar that the American government spends. Control group subjects simply learn that two news stories have been released, one pertaining to a government report about the crime rate and one pertaining to a report about American foreign aid. Gilens then asks all subjects about their level of support for government spending on prison construction and federal spending on foreign aid.

Gilens's results demonstrate that the provision of policy-specific information significantly influences subjects' opinions, even among subjects who possess high levels of general political knowledge. Specifically, treatment group subjects (who learn that the crime rate and foreign aid spending have decreased) are much less likely than the control group to support increasing government spending on prison construction and decreasing American spending on foreign aid. Interestingly, Gilens shows that policy-specific information has a stronger influence on subjects who possess high levels of general political knowledge. Gilens suggests that citizens' ignorance of policy-specific facts (and not a lack of general political knowledge) is what hinders them from expressing their opinions effectively on certain policy issues.

Kuklinski et al. (2000) also assess whether policy-specific facts are relevant to citizen opinions. In contrast to Gilens, Kuklinski et al. distinguish between citizens who are uninformed (i.e., who lack information about particular policies) and citizens who are misinformed (i.e., who hold incorrect beliefs about particular policies). Indeed, Kuklinski et al. suggest that the problem facing our democracy is not that citizens are

uninformed, but rather that citizens confidently hold incorrect beliefs and base their opinions upon them.

Kuklinski et al. assess experimentally whether and when the provision of correct policy-specific information induces citizens to abandon incorrect beliefs and express different opinions. In one treatment group, subjects receive six facts about welfare (e.g., the percentage of families on welfare, the percentage of the federal budget devoted to welfare) before they express their opinions. In another treatment group, subjects first take a multiple-choice quiz on these six facts about welfare. After answering each quiz question, subjects in this treatment group are asked how confident they are of their answer. The purpose of the quiz and the follow-up confidence questions is to gauge subjects' beliefs about welfare and how confidently they hold them. In the control group, subjects simply express their opinions about welfare policy.

Kuklinski et al. also conduct follow-up experiments in which the information about welfare is made more salient and meaningful than the six facts provided in the initial experiments. To this end, Kuklinski et al. ask subjects what percentage of the federal budget they think is spent on welfare and what percentage of the budget they think should be spent on welfare. Immediately after answering these two questions, subjects in the treatment group are told the correct fact, which for most subjects is that actual welfare spending is lower than either their estimate or their stated preference. Control group subjects answer these two questions but do not receive the correct fact about actual welfare spending. At the completion of these experiments, subjects in both groups express their level of support for welfare spending.

In both experiments, Kuklinski et al. find that subjects are grossly misinformed about welfare policy. Indeed, the percentage of subjects who answer particular multiple-choice quiz questions incorrectly ranges from 67% to 90%. Even more troubling is their finding that subjects who have the least accurate beliefs are the most confident in them. Kuklinski et al. also show that the opinions of subjects who receive the six facts about welfare are no different than the opinions of subjects in the control group, which indicates that treatment group subjects either did not absorb these facts or failed to change their opinions in light of them. However, when a fact about welfare is presented in a way that explicitly exposes and corrects subjects' incorrect beliefs (as in the follow-up experiments), subjects adjust their opinions about welfare spending accordingly. In this way, Kuklinski et al. demonstrate that policy-specific information can induce misinformed citizens to abandon their incorrect beliefs and express informed opinions, but only when it is presented in a way that is meaningful and relevant to them.

Taken together, experiments focusing on choice contexts have shown conditions under which knowledge of the kinds of facts that have been the basis of previous political knowledge tests are neither necessary nor sufficient for competent choices. Some of these conditions pertain to the kind of information available. Other conditions pertain to the context in which the information is delivered. Collectively, these clarifications provide a different view of citizen competence than is found in most non-experimental work on political knowledge. While there are many cases in which lacking information reduces citizens' competence, experiments clarify important conditions under which things are different. In particular, if there are relatively few options from which to choose (as is true in many elections), if people can be motivated to pay attention to new information, if the

information is highly relevant to making a competent decision, and if people's prior knowledge of the topic, the speaker, or even the context leads them to make effective decisions about whom and what to believe, then even people who appear to lack political knowledge as conventionally defined can vote with the same level of competence as they would if better informed.

Conclusions

Our argument in this chapter has been as follows. First, there are many ways in which the survey questions that scholars use to measure political knowledge are not valid indicators of what citizens know about politics. Second, for circumstances in which such measures are valid, it is not clear that this knowledge is necessary, sufficient, or relevant to a citizen's ability to perform important democratic tasks, such as voting, competently (i.e., as they would if they knew all relevant facts). Third, experiments have shed important new light on both of these issues.

Collectively, the kinds of experiments we have described above are not only helping scholars more effectively interpret existing data, but they are also beginning to help scholars create better measures of political knowledge. Experiments aid interpretation of existing data by helping scholars better reconcile possible interpretations of political knowledge data with increasingly rigorous and transparent logic and evidence about how citizens think. Already, experiments are changing how political knowledge is measured by clarifying relationships between the content of people's thoughts, what people say to survey interviewers, and how both concepts relate to the choices that they

make. In a very short period of time experiments have transformed the study of political knowledge.

Yet, experiments have just begun to scratch the surface of the multifaceted ways in which thought and choice interact. They have examined only a few of the many attributes of survey interviews that can affect responses. They have also examined only a few of the many ways that particular kinds of political knowledge can affect politically relevant decisions and opinions. Their different results have also raised questions about whether and when particular types of information eliminate differences between more and less knowledgeable citizens (compare Lau and Redlawsk 2001 to Kuklinski et al. 2001 and Boudreau 2009). As research in fields of study such as political cognition, the psychology of the survey response, and political communication evolve, more questions will be raised about the validity of extant knowledge measures and whether particular kinds of knowledge are relevant to democratic outcomes. While such questions can be studied in many ways, experiments should take center stage in future research. Hence, the experiments we have described represent the beginning, rather than the end, of a new attempt to better understand what people know about politics and why it matters.

References

Boudreau, Cheryl. 2009. "Closing the Gap: When Do Cues Eliminate Differences between Sophisticated and Unsophisticated Citizens?" *The Journal of Politics* 71(3): 964-976.

Gibson, James L., and Gregory A. Caldeira. 2009. "Knowing the Supreme Court? A Reconsideration of Public Ignorance of the High Court." *The Journal of Politics* 71(2): 429-441.

Gilens, Martin. 2001. "Political Ignorance and Collective Policy Preferences." *American Political Science Review* 95(2): 379-396.

Gould, Stephen J. 1996. *The Mismeasure of Man. Revised Edition*. New York: W.W. Norton and Co.

Graber, Doris. 1984. *Processing the News: How People Tame the Information Tide*. New York: Longman.

Krosnick, Jon A., Arthur Lupia, Matthew DeBell, and Darrell Donakowski. 2008. "Problems with ANES Questions Measuring Political Knowledge." Posted at www.electionstudies.org/announce/newsltr/20080324PoliticalKnowledgeMemo.pdf

Kuklinski, James H., Paul J. Quirk, Jennifer Jerit, and Robert F. Rich. 2001. "The Political Environment and Citizen Competence." *American Journal of Political Science* 45(2): 410-424.

Kuklinski, James H., Paul J. Quirk, Jennifer Jerit, David Schwieder, and Robert F. Rich. 2000. "Misinformation and the Currency of Democratic Citizenship." *The Journal of Politics* 62(3): 790-816.

Lau, Richard R. and David P. Redlawsk. 1997. "Voting Correctly." *American Political Science Review* 91(3): 585-598.

Lau, Richard R. and David P. Redlawsk. 2001. "Advantages and Disadvantages of Cognitive Heuristics in Political Decision Making." *American Journal of Political Science* 45(4): 951-971.

Lupia, Arthur, and Mathew D. McCubbins. 1998. *The Democratic Dilemma: Can Citizens Learn What They Need to Know?* New York: Cambridge University Press.

Lupia, Arthur. 2006. "How Elitism Undermines the Study of Voter Competence." *Critical Review* 18 (1-3): 217-232.

McClone, Matthew S., Joshua Aronson, and Diane Kobrynowicz. 2006. "Stereotype Threat and the Gender Gap in Political Knowledge." *Psychology of Women Quarterly* 30: 392-398.

Miller, Melissa, and Shannon K. Orr. 2008. "Experimenting with a "Third Way" in Political Knowledge Estimation." *Public Opinion Quarterly* 72(4): 768-780.

Mondak, Jeffery J. 2001. "Developing Valid Knowledge Scales." *American Journal of Political Science* 45(1): 224-238.

Mondak, Jeffery J., and Belinda Creel Davis. 2001. "Asked and Answered: Knowledge Levels When We Will Not Take "Don't Know" for an Answer." *Political Behavior* 23(3): 199-224

Mondak, Jeffery J., and Mary R. Anderson. 2004. "The Knowledge Gap: A Reexamination of Gender-based Differences in Political Knowledge." *The Journal of Politics* 66(2): 492-513.

Popkin, Samuel L. 1994. *The Reasoning Voter: Communication and Persuasion in Presidential Campaigns*. Chicago: University of Chicago Press.

Prior, Markus, and Arthur Lupia. 2008. "Money, Time, and Political Knowledge: Distinguishing Quick Recall from Political Learning Skills." *American Journal of Political Science* 52(1): 168-182.

Steele, Claude M., and Joshua Aronson. 1995. "Stereotype Threat and the Intellectual Test Performance of African-Americans." *Journal of Personality and Social Psychology* 69: 787-811.

Sturgis, Patrick, Nick Allum, and Patten Smith. 2008. "An Experiment on the Measurement of Political Knowledge in Surveys." *Public Opinion Quarterly* 72(1): 90-102.