



Conceptual Foundations of Citizen Competence

Author(s): James H. Kuklinski and Paul J. Quirk

Source: *Political Behavior*, Vol. 23, No. 3, Special Issue: Citizen Competence Revisited (Sep., 2001), pp. 285-311

Published by: [Springer](#)

Stable URL: <http://www.jstor.org/stable/1558386>

Accessed: 19/12/2014 09:15

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at <http://www.jstor.org/page/info/about/policies/terms.jsp>

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.



Springer is collaborating with JSTOR to digitize, preserve and extend access to *Political Behavior*.

<http://www.jstor.org>

CONCEPTUAL FOUNDATIONS OF CITIZEN COMPETENCE

James H. Kuklinski and Paul J. Quirk

This article considers some of the challenges that attend efforts to assess citizen performance. We begin by demonstrating the often- unarticulated complexity of evaluating performance in any domain. To do this, we identify four distinct conceptual elements that comprise an evaluation—identification of task, selection of criterion, choice of empirical indicator, and explication of standard—and illustrate with an example that is relatively free of ambiguity: performance in basketball. Using this framework, we then review research in three general areas of study: mass belief systems and issue consistency, political knowledge, and the use of political heuristics. We find that no study articulates all four elements (or adequate substitutes associated with an alternative framework). As a result, problems arise. Most significantly, any particular study is likely to use criteria that are unsatisfactory in important respects or to employ empirical indicators that do not validly measure the criteria. Across studies, conclusions often vary as a function of unarticulated differences in assumptions, definitions, and measures. We conclude by drawing a few lessons for future research, while also recognizing the impressive progress that the study of public opinion and citizen competence has made over the last 40 years.

Key words: citizens; performance; competence; public opinion.

The stuff of politics is contestable. There is no single right way to vote, no single right position on issues, no single right set of beliefs. From the standpoint of studying citizen performance, this observation is bad news. It means that scholars cannot evaluate the quality of decisions in a straightforward fashion. Assessing performance would be simple if liberal or conservative decisions were always the right decisions or if a select group of individuals who were known to “get it right” always agreed. For scholars who study such things, unfortunately, neither is the case.

James H. Kuklinski, Department of Political Science and Institute of Government and Public Affairs, University of Illinois at Urbana–Champaign (kuklinsk@uiuc.edu); Paul J. Quirk, Department of Political Science and Institute of Government and Public Affairs, University of Illinois at Urbana–Champaign (p-quirk@uiuc.edu).

The only available option, therefore, is to seek out alternative ways of assessing performance, which scholars have done with imagination. They have considered whether citizens hold consistent positions across issues; whether they hold stable positions across time; whether they know relevant facts from a policy debate; whether they maintain their positions when given different framings of the same issue; whether their preferences are correlated with their values; whether their preferences resemble those of others who are well informed; and whether they effectively take cues from parties, politicians, interest groups, and other citizens.

Most authors do not use words like “citizen competence” or “citizen performance” in reporting their empirical work. Rather, they focus on more specific ideas, reflected in terms such as “heuristics,” “issue constraint,” “issue framing,” “factual knowledge,” and “political sophistication.” Nevertheless, assessing performance is an obvious underlying motivation, sometimes explicitly stated, in all of these bodies of work. Indeed, it is this common motivation that integrates (as shown in cross-references) what otherwise might be viewed as entirely distinct inquiries.¹

Rather than offer another focused study, we step back and take a broad look at the conceptual foundations of research that assesses citizen performance. We begin by demonstrating the often-unarticulated complexity of evaluating competence in any domain. We do so by identifying several distinct conceptual elements that comprise an evaluation of competence. We illustrate these concepts, and some difficulties that can arise, with an example that is relatively free of intrinsic ambiguity: performance in basketball.

We then review how selected studies have approached the evaluation of performance. We find that they have often failed to spell out the corresponding conceptual elements. At a minimum, therefore, the logic of their evaluations is in some respects unclear. More important, however, the lack of explicit attention to conceptual issues has had substantive consequences. The next section looks in some detail at specific problems that emerge in these studies. We find performance criteria that do not stand up to close scrutiny and empirical indicators that do not validly measure the criteria. We then broaden our view and show that conclusions about citizen competence vary across studies as a function of conceptual conflicts that are largely unrecognized. In effect, researchers speak past each other without confronting their central disagreements. We conclude with several lessons for future research on citizen performance.

THE ELEMENTS OF EVALUATION

In many contexts outside of politics, an evaluation of performance can look quite simple and data driven. An empirical finding is reported (say, a test score, a monthly sales figure, or the like); and a conclusion is drawn. But it

only looks this way, of course, because important conceptual elements of the evaluative process are taken for granted. Beneath the surface, a substantial conceptual apparatus is quietly at work.

We propose that there are four principal conceptual elements in any evaluation of performance. First, one must identify the *task* that the actor is asked to undertake. Because some actors perform multiple tasks and tasks can be divided and combined in various ways, the choice of a task need not be straightforward. Second, one must state a *criterion* by which the performance is to be evaluated—that is, the property or attribute that is taken to constitute the quality of performance. Third, one must select at least one empirical *indicator* of that criterion. Finally, to categorize levels of performance, one must identify *standards* with respect to the indicator. Standards map levels of the indicator onto a set of evaluative categories: satisfactory or unsatisfactory; very good, good, fair, or poor; or the like. In some contexts, standards are entirely comparative; in others, they reflect independent notions of success in the task.

The Case of Basketball

To illustrate these concepts and some of the difficulties that can arise, we briefly consider two tasks that occur in playing basketball. One important task that basketball players perform is shooting free throws during games. The obvious performance criterion is making the free throws: the ball should go through the hoop. There are other possible criteria, such as using proper form or avoiding delay, but none is remotely comparable in importance. The main empirical indicator is the percentage of free throws made, typically over the duration of the season. The standards, which depend on the level of play, might be that a good free throw shooter makes more than 75% of his or her shots, an average one 60–75%, and a poor one less than 60%. Simply stating that a basketball player is a 75% free-throw shooter might be sufficient to convey an evaluation, but only for someone who knows the game (and thus the relevant concepts) and the level of play.

Another, quite different task in basketball is playing defense. A possible criterion of competence—as we will see, a problematic one—is blocking shots. One plausible measure of this criterion is the number of blocked shots per game. With respect to standards, the evaluator might place players into the top third, the middle third, and the bottom third.

Nothing we have said thus far suggests that evaluating basketball performance is particularly difficult. Most basketball fans would probably see no difficulty at all: compile the relevant statistics and let them speak for themselves. In fact, however, our discussion belies how treacherous assessing performance can be, even in the simple case of basketball. Having enumerated the steps, we can begin to see more clearly what some of the challenges are.

Assuming that task, criterion, empirical indicator, and standard are all ex-

plicitly stated, three additional conditions should be met. First, the criterion should be (at least generally) necessary for competence. Suppose it is not. Then an individual could easily fail with respect to the criterion and yet be entirely competent as a result of satisfying other criteria that the researcher overlooks. Second, the criterion should be (at least generally) sufficient for competence. If it is not, then an individual could easily meet the criterion and yet be incompetent in light of other criteria. The two basketball tasks discussed above differ sharply in these respects. Making free throws is both necessary and sufficient to being a good free-throw shooter. No other information is needed for an assessment. Blocking shots, in contrast, is neither necessary nor sufficient to playing good defense, nor even generally so. A poor shot blocker could easily play good defense by making steals or interfering with opposing teams' plays; a good shot blocker could easily play poor defense by reacting to fakes and giving up easy shots. The single indicator of shot blocking thus tells one rather little about competence in defensive play, even though it is obviously relevant to such competence. Evaluating defensive play is a more formidable challenge than evaluating free throw shooting.

More formidable, that is, unless one divides the task more finely than playing defense. For example, the evaluator could define one task as blocking shots, another as making steals, and yet another as preventing easy shots. Although this division eliminates some of the problems noted in the preceding paragraph, it potentially raises another: defining tasks so narrowly as to diminish the importance of the evaluation.²

Finally, the empirical indicator must be valid as a measure of the criterion. If it is not, any conclusion will be suspect. Unfortunately for researchers, valid indicators are sometimes hard to find. In basketball, for example, no one even attempts to measure players' floor leadership or their ability to help teammates get open for shots.

Implications for the Study of Citizen Performance

In any case, politics is not basketball. Evaluating competence is more problematic in politics than in basketball for several reasons. First, in politics, there are no indisputably right decisions, comparable to free throws that go through the hoop. One cannot measure performance as the percentage of issues on which a citizen takes the right position. Second, standards of performance for any indicator are not given by the conditions of a competition. Whether an individual citizen performs well or poorly compared with other citizens is not necessarily even relevant. Rather, the central issue is whether the citizenry, collectively, performs adequately. Finally, and most fundamental, tasks and performance criteria for citizens are not determined by a well-defined object of the game. In basketball, these matters are decided by reference to the

object of scoring more points than the opposing team. In politics, the object of the game is subject to varying interpretation. That is because, except figuratively, it is not a game.

Taken together, the complications of evaluating competence in any context and the special problems of evaluation in politics have three important implications: First, the elements of evaluation are not given or generally understood; rather, they require important choices. Second, such choices will inevitably vary from one study, or one approach, to another; there can be no expectation of uniformity. And finally, the choices made in particular studies will often be subject to serious difficulties. Among other hazards, a study might inadvertently fuse task, criterion, and indicator; adopt criteria that fall far short of necessity, sufficiency, or both; use empirical indicators that lack a reasonable expectation of validity; or set standards that are unworkable or inadequate. In general, there will be a temptation to employ whatever conceptual elements are convenient from the standpoint of empirical research.

It is important, therefore, that scholars spell out the elements of their evaluative approach and, where the rationale is not self-evident, provide the reasoning for their choices. In the absence of such discussion, the significance of any conclusion about citizen competence, positive or negative, is unclear. Even serious problems in the logic of an evaluation are likely to go undetected. And different studies or approaches will reach conflicting conclusions for reasons that are neither debated nor even recognized. In an important sense, such conclusions are essentially arbitrary.

In the next section, we turn to a few selected studies of citizen performance and assess how scholars have handled the conceptual elements of their evaluative approach. Before we proceed, however, two caveats warrant mention. First, we do not argue that our conceptual framework—distinguishing tasks, criteria, indicators, and standards—is the only workable framework for evaluating citizen competence. Indeed, some other scheme might prove more useful, generally or in certain contexts. At the present time, however, it is the only such framework available in the literature. As far as we can see, any satisfactory framework will need to cover at least the same conceptual ground. In any case, we do not criticize authors for failing to use our terms if their relevant intentions are reasonably clear.

Second, in discussing necessary and sufficient conditions for competence, we have a specific and limited purpose in mind. Most of the works we review below rely on probabilistic methods and interpret the world in terms of conditional probabilities. We are not suggesting that scholars abandon these methods and use necessity and sufficiency for purposes of ascertaining causality. Conceptually, however, we believe that assessing criteria for competence in terms of necessary and sufficient conditions, somewhat loosely defined, helps bring some central issues into focus.³

STUDIES OF CITIZEN PERFORMANCE: EXPLICATING THE ELEMENTS OF EVALUATION

There is no dearth of studies or, for that matter, perspectives from which to draw. We have chosen to examine three topics that have been highly active research areas at one time or another over the last 40 years: issue consistency, factual knowledge, and the use of heuristics.⁴ Even this literature is far too vast to explore in its entirety. Therefore, where exemplary studies exist, we focus principally on them. Where they do not, we cast our nets more widely.⁵

Belief Systems and Issue Consistency

The study of issue consistency is synonymous with the name Converse (1964), and so the choice of exemplar in this case is really no choice at all. Nearly 40 years old, Converse's seminal work has been the target of various criticisms, most having to do with issues of validity and reliability (Achen, 1975; Erikson, 1979; Jackson, 1979; Judd and Milburn, 1980). Our purpose is not to revisit those criticisms here. Rather, we want ultimately to consider the logic underlying Converse's evaluation of citizen performance, and we begin here with his explication of the evaluative steps.

Converse does not directly identify the citizen's task. The most defensible interpretation of his analysis is that he defines the citizen's fundamental task simply as understanding politics. He takes for granted that this understanding, or the lack of it, will affect performance in a wide range of more specific tasks—voting, expressing policy preferences, participating in public life, and the like.⁶ In effect, he suggests plausibly that understanding politics is the key to competence in all of them.

Converse's criterion with respect to this task is that a citizen should have an organized belief system concerning politics. He defines a belief system as "a configuration of ideas and attitudes in which the elements are bound together by some form of constraint or functional interdependence" (1964, p. 207). Elaborating, he adds that "'constraint' may be taken to mean the success we would have in predicting, given initial knowledge that an individual holds a specified attitude, that he holds certain further ideas and attitudes" (p. 207). Converse acknowledges (perhaps thinking of Lane [1962]) that, in principle, different citizens might hold a variety of idiosyncratic belief systems that would all serve this purpose. But practically speaking, he argues, the only way for ordinary citizens to develop a meaningful belief system is to learn political ideology, as political elites define it and as it is diffused throughout society. Developing one from scratch is beyond their capability. And, because party politics revolves around the liberal-conservative continuum, American political ideology is also the most relevant belief system.

Citizens need an organized belief system to help them make sense of the political world. As Converse describes it, “there are many crucial consequences of such organization: With it, for example, new political events have more meaning, retention of political information from the past is far more adequate, and political behavior increasingly approximates that of sophisticated ‘rational’ models, which assume relatively full information” (1964, p. 227). In short, an organized belief system provides the foundation for understanding politics.

Converse’s theoretical discussion, taken literally, suggests that having an organized belief system requires citizens actually to *be* liberals or conservatives, who subscribe to one set of ideas and policies or the other. That is, he says that a belief system is needed, and proceeds to define a belief system in terms of constrained attitudes. The notion of attitudinal constraint evidently demands faithful adherence to a liberal or conservative ideology.⁷ On the other hand, some of his early empirical analysis, in which he attempts to determine whether people identify parties in ideological terms, suggests that it might be enough to *know about* liberals and conservatives and their opposing ideas and policies, that is, to *understand* ideological politics.

Generally, we will follow Converse’s most direct statements and assume that he subscribes to the more restrictive definition of a belief system, which requires citizens to hold consistently liberal or conservative positions. However, we will also consider the implications of the more lenient possible interpretation—citizens only need to know what the liberal and conservative positions are—for the choice of criteria and indicators.

With respect to standards, Converse does not say directly how high the across-issue correlations must be for citizens to possess an adequate belief system. Instead, he simply compares the correlations among citizens with those among a sample of congressional candidates. He finds that the average correlation within domestic issues is .53 among the candidates and .23 among citizens; within foreign issues it is .37 among candidates and .23 among citizens. On the basis of these and other, similar comparisons, Converse concludes that most citizens do not have satisfactory belief systems.⁸

Factual Knowledge

Another important stream of work examines citizens’ knowledge of politics and policy. Significantly, this body of research has a good deal of conceptual disorder even in its purely empirically oriented dimension. Different researchers use the same or similar survey items to measure differently labeled concepts. In some instances, researchers measure *political sophistication* by asking respondents who the vice president is, what a line-item veto is, and which party controls the Senate (for excellent overviews, see Luskin, 1987, 1990).

Those who answer these questions correctly are said to be more politically sophisticated than those who do not. Yet, others have construed similar questionnaire items as indicators of what is apparently a different concept: *political awareness*. Zaller (1990, p. 132; also see Zaller, 1992), for example, measures this concept with a “series of direct information tests (e.g., which party controls Congress, the term of office of a U.S. senator), ability to evaluate a variety of somewhat obscure political figures (e.g., Henry Jackson), ability to recall the names of Congressional candidates, and ability to locate accurately the policy positions of prominent individuals and groups.”

In fact, the concept-indicator nexus is even more varied. To quote Zaller (1990, p. 126) again, “variables purporting to measure ‘political awareness,’ ‘political expertise,’ ‘political sophistication,’ ‘cognitive sophistication,’ ‘political information,’ ‘political involvement,’ ‘media exposure,’ and ‘political interest’ appear regularly in the public opinion literature and are used (along with education) more or less interchangeably to explain the same family of dependent variables.” This proliferation of related, yet significantly different concepts indicates that there is little consensus about the central processes in which political information plays a role. This confusion on the empirical issues would not necessarily preclude articulating a coherent approach to assessing competence, but it certainly does not make it easier. If a lot of these concepts are interchangeable, then is there another, more encompassing concept that scholars have not yet identified and that everyone should be using?⁹ To put it more bluntly, if the concept is the same, why do different labels abound? Alternatively, if scholars offer the different labels for purposes of conceptual discrimination, then why do they use many of the same survey items to measure them?

The most comprehensive study of factual knowledge to date is Delli Carpini and Keeter’s *What Americans Know about Politics and Why It Matters* (1996), a study that was published more than 30 years after Converse’s. We focus on this study both to simplify our task and to show that some of the conceptual ambiguities that we found in Converse’s study also appear in the most recent and fully developed work on factual knowledge.

By comparison with Converse, Delli Carpini and Keeter devote considerably greater space to the basic elements of their approach to assessing competence, and they face many of the central issues in assessing citizen competence head-on. Their exposition is careful and often penetrating. Nevertheless, perhaps largely because their inquiry is exceptionally wide ranging, the discussion falls short of laying out a coherent perspective. In some respects, it is unclear exactly how to interpret their effort. The difficulties show how far scholars still need to go to put such assessment on a sounder footing.

To begin with, Delli Carpini and Keeter are notably thorough in identifying

a range of specific tasks that citizens are asked to undertake. Setting forth a litany of prescribed activities, they note that citizens must:

select qualified representatives (both within parties and in general elections) for local, state, and national offices; serve as the pool from which representatives are selected; reward and punish officeholders for their past performances; vote directly on policy issues through initiative and referenda; fill the thousands of voluntary, appointed, and bureaucratic civic roles required for the machinery of campaigns, elections, and government to work effectively; help shape local, state, and national political agendas through numerous outlets from public opinion polls to public demonstrations to direct contact with public officials; support and cooperate with the implementation of public policies; navigate government bureaucracies for information, goods, and services; attend local government and civic meetings; and more. (1996, p. 4)

These are the kinds of specific tasks that Converse left unstated. Of course, Delli Carpini and Keeter do not examine citizen performance on every one of these tasks. Rather, when they turn to the central empirical chapter for assessing competence (chapter 6), they identify four tasks that are more general: holding democratic values, participating in civic and political life, holding high-quality opinions, and acting in one's enlightened self-interest.

The relationships between and within the two lists of tasks, however, are complex and murky. Whereas the specific tasks all entail physically doing something, two of the general tasks—holding democratic values and holding high-quality opinions—do not. Moreover, three of the more general tasks—holding democratic values, holding high-quality opinions, and acting in one's enlightened self-interest—are likely requirements for competently performing some of the specific tasks, as well the fourth general task, participation. And one of these three, holding high-quality opinions, also seems to precede one of the others, acting in one's enlightened self-interest. In short, some tasks are seemingly subsumed by others, which in turn are subsumed by yet others. But there is no formula for choosing one level of task over another, nor a hierarchical structure for relating them to each other. Finally, one of the general tasks, acting on enlightened self-interest, looks more like a criterion for evaluating performance in other tasks. As we will see later, many students of political heuristics have made this notion—or at least enlightened action (with or without self-interest)—the central criterion for judging performance: the task is expressing preferences and the criterion is that the preferences should be enlightened.

In any event, Delli Carpini and Keeter also propose a considerable variety of performance criteria. For the task of holding democratic values, the criterion is being politically tolerant of others. For the task of civic and political

participation, the criteria are voting and taking part in campaigns. For holding high-quality opinions, they are holding opinions of any kind and holding stable and consistent opinions. Finally, for acting in one's enlightened self-interest, the criterion is expressing issue preferences that comport with one's ascriptive group memberships—male or female, black or non-black, and so on.

These criteria do not make a tidy package. Two of them, voting and campaign activity, are familiar from the list of specific tasks. More important, the whole set of criteria does not have a consistent theme. Rather, they seem to incorporate three separate and even conflicting ideas: enlightened action, along with its antecedents or consequences; self-interest, in some matters or in some degree; and democratic values, evidently regardless of self-interest.

Delli Carpini and Keeter do an admirable job of identifying and providing a rationale for one or more empirical indicators for each of the criteria. The indicators include: for political tolerance, expressed attitudes toward the courts and civil liberties; for participation, voting and working in campaigns; for stable and consistent opinions, stability on NES 7-point issue items over two waves and attitude constraint, respectively; and for enlightened self-interest, the relationships between demographic attributes and issue preferences. All of the indicators have clear *prima facie* connections with the criteria they are supposed to address and are supported by explicit discussion. Although we find problems with the demographic attributes-issue preferences indicator, the rationale for using it is carefully stated.

Like Converse, Delli Carpini and Keeter do not offer explicit standards to identify satisfactory levels with respect to the indicators. In some of their analyses, however, there is a clear, even though unspoken standard: the adequate level is the actual level of those who do “best”—that is, those who are most willing to let their most disliked group give a speech, those who give the most money to campaigns, and so on. In their many graphs, they show that citizens with more information are consistently more likely to reach that top level. The authors do not claim that these standards reflect an analysis of the consequences of different performance levels. They appear to reflect a mere stipulation that the highest level achieved by the best performing group, or the highest level discriminated by a measure, must be satisfactory.

Heuristics

Introduced into the study of public opinion in the 1980s, the political heuristics literature centers on two interrelated ideas: (1) neither an organized belief system nor much factual knowledge is necessary to adequate performance; rather, (2) citizens can compensate for their absence by relying on heuristics, or mental shortcuts, to make their decisions. This research put the

political environment and the cues it provides at the center of public opinion research. For many scholars, it also elevated the ordinary citizen from a hopeless incompetent to a reasonably capable participant in democratic politics (Brady and Sniderman, 1985; Carmines and Kuklinski, 1990; Lupia, 1994; Lupia and McCubbins, 1998; Mondak, 1994; Mutz, 1998; Popkin, 1991; Sniderman, Brody, and Tetlock, 1991; but see Bartels, 1996; Kuklinski and Quirk, 2000, and Luskin, 2001, for a different view).

Despite the common ideas, research on heuristics is highly diverse. For one thing, researchers have proposed a wide range of potential cues—including but not limited to parties, interest groups, politicians, election results, other citizens whom one knows, other citizens whom one does not know, those perceived to be knowledgeable, and the level of consensus among visible political elites. For another, authors appear to differ in their assumptions about the limits of heuristics, with some (e.g., Popkin, 1991) expressing more enthusiasm than others (e.g., Sniderman et al., 1991) about their effectiveness in promoting competent decisions. Almost everyone, however, reaches far more optimistic conclusions than psychologists do (examples include Kahneman, Slovic, and Tversky 1982, and Nisbett and Ross, 1980; a lively overview is Piattelli-Palmarini, 1994). Finally, some authors associate the use of heuristics primarily with less-informed citizens (Sniderman et al., 1991), while others find heuristics essential to competence for all citizens (Popkin, 1991; and Delli Carpini and Keeter, 1996, who generally criticize the heuristics literature). On the whole, this literature is explicitly concerned with evaluating competence and is quite transparent in its evaluative approach.

The heuristics literature is characterized by a focus on clear, well-defined tasks. In a few cases, authors examine a judgment task, such as identifying the ideological tendency of a political group (Brady and Sniderman, 1985). But, for the most part, the citizen's task in this work is to make decisions about policies and candidates, that is, to express preferences. In one of the most important studies, for example, Lupia (1994) analyzes how California voters reached decisions on several insurance referenda put before them in the 1988 election. One reason for this direct focus on concrete decision tasks is that the literature proceeds from skepticism about the need to perform difficult mental tasks.

This literature is also direct and generally coherent with respect to criteria. In some cases, scholars identify the criterion somewhat casually, almost in passing. They ask whether citizens can make “good” or “reasonable” decisions in some context, without worrying about general definitions of those terms. In other cases, they provide a more elaborate formulation of what is probably the same criterion: that citizens should make the same decisions that they would make if they were well informed.

On the other hand, heuristics scholars are generally silent about the deep

ambiguities that lurk in the pregnant phrase “as if well informed.” What information is necessary to be well informed? Does being well informed mean merely having information, or is it also necessary to understand it? Is the assumption that the well informed always make the right decision, or are other hypothetical conditions—careful deliberation, absence of extreme emotion, or the like—also intended? Rather than pause to consider such difficulties—or indeed, even to notice them—the researchers usually skip quickly to the indicators.

Heuristics scholars have selected indicators in one of two ways, corresponding to the looser and more precise statements of the criterion. As an indicator for “good” or “reasonable” decisions, scholars sometimes look at whether citizens who are exposed to a particular item of information behave in ways that, on the face of it, reflect sensible use of that item of information. In his article on presidential positions as cues, for example, Mondak (1993) reports that citizens respond more favorably to a cue indicating the president’s sponsorship of a policy when the president has high approval ratings than when he has low ratings. Such a response makes sense, in that if one approves of the president, one should usually expect to favor his policies.

As the indicator for the more elaborate, as-if-well-informed criterion, scholars have compared decisions by the focal group, the one whose competence is at issue, to the decisions of another group that is assumed to have similar values and to be well informed. We will call such a comparison group a well-informed-proxy group and this method a well-informed-proxy comparison. The more the focal group acts like the well-informed-proxy group, of course, the more competent it is. Here again, the researchers typically do not pause for much discussion about the proxy group and its qualifications to stand for the well informed, whatever that criterion is taken to mean. As we will see, the conceptual weaknesses in this approach have mostly to do with the qualifications of the proxy group.

Appealing to criteria that are essentially open-ended in their demands—good, reasonable, or well informed decision—the political-heuristics literature not surprisingly is vague and inconsistent with respect to standards: How closely should citizens approximate a well-informed decision, however defined, to be regarded competent? No one has come forth to propose an answer. In effect standards are set by the most straightforward analysis of a given indicator. Thus, if the indicator is citizens’ response to a particular item of information, the standard is that people who receive it should move in the correct direction, rather than the opposite. If people approve of the president’s performance, for example, they should support his policy proposals more, not less. If the indicator is a well-informed-proxy comparison, the standard is a function of the proxy group. That is, whatever the nature of the proxy group, the focal group should roughly approximate its decisions.

THE CRITERION-INDICATOR NEXUS

We have found that authors do not always clearly identify the central elements in the evaluation of citizen performance—in our terms, the task, criterion, indicator, and standard. Trying to explicate them required us to attribute a variety of decisions and assumptions, not always with certainty. Even when the authors lay out some of these elements, they often devote little effort to identifying their strengths, weaknesses, and limitations. These omissions matter because they can hide difficulties in a particular study or produce essentially arbitrary, yet consequential differences across studies. Ultimately, they can undermine the value of any conclusions about citizen competence.

Next, we document some of these problems. In this section, we deal principally with the selection of criteria and empirical indicators in each of the three areas of research. In the following section, we will consider two examples of inconsistency across studies. In one case, the inconsistency arises from a change in the definition of the task; in the other, from the use of widely varying standards.

Ideological Belief System as Performance Criterion

By our principal interpretation, Converse proposes a task for citizens of understanding politics and introduces holding an ideological belief system as the performance criterion with respect to the task. As we have noted, the rationale for this criterion is that an organized belief system is what enables people to deal with the complexity of politics.

For citizens to understand American politics, must they hold ideological belief systems (in Converse's strict sense of having clear-cut ideological positions)? It is hard to see why. A citizen can know what liberals and conservatives stand for without joining either side and even without merely compromising between them.¹⁰ Consider, for example, a discriminating moderate who has reasons for taking liberal positions on some issues and conservative positions on others. Indeed, an important political group consists of well-educated suburbanites who are conservative on economic issues but liberal on moral or cultural issues. Few would suppose that such individuals are less able to understand politics than dedicated ideologues are.

If holding an ideological belief system is not necessary for understanding politics, is it at least sufficient? Again, it seems not. Ideological concepts relate to very general dimensions and effects of public policy—benefits for lower-income groups, government activism, economic freedom, traditional morality, and so on. If nothing else mattered to citizens' values and interests, then reacting appropriately to the ideological content of policy debate would amount to fully understanding it.¹¹ All that would matter about a policy is

how liberal or conservative it was. But such a conception overlooks numerous domain-, situation-, and policy-specific effects. Health is one thing, education another. Policies can be reckless, irresponsible, prudent, or efficient. Some policies are more likely to work than others. All these distinctions would be lost (or conflated with ideological categories) for citizens who tried to understand politics solely on the basis of an ideological belief system. To deal with such distinctions, and recognize evidence with respect to them, requires more than an ideology.

Finally, it is not even clear that having an ideological belief system only helps people understand politics and never hinders them. The experimental work of Lodge and colleagues (Lodge and Hamill, 1986; also see Lodge and Taber, 2000, on motivated reasoning) suggests that those with strong partisan identifications, presumably also the possessors of the most complete belief systems, are indeed able to remember more facts and arguments about politics than others. But what they remember is also more selective and more often distorted; thus it is potentially more misleading. Those lacking an ideological belief system retain a smaller sample of the available information, but also a more representative one. In a similar vein, our work on misinformation (Kuklinski, Quirk, Jerit, Schweider, and Rich, 2000) has found that highly partisan citizens hold more biased factual beliefs about welfare than others. Such ideologically selective perception undoubtedly leads to easier decision-making and a greater feeling of certitude; but those effects do not represent superior understanding in any straightforward sense.

As we have noted, there are alternative interpretations of Converse's approach. But the alternatives do not eliminate the difficulties; they merely shift them from one step to another. For example, a defensible construction of Converse would identify his criterion as merely recognizing ideological concepts and positions rather than actually holding an ideological belief system. That modification would make the criterion more plausibly necessary for understanding American politics. But it would render his indicator, issue constraint, unnecessary—arguably, invalid—for the criterion. People need not be consistent themselves to understand what liberal and conservative positions are.

Enlightened Preferences and Group Self-Interest

As we have noted, Delli Carpini and Keeter propose a concept of enlightened self-interest as one of their main criteria of competent performance. The term, however, contains an important ambiguity. If self-interest is defined broadly, so that it includes public-regarding values and preferences, the criterion is essentially equivalent to acting rationally or "getting it right." But, if self-interest refers specifically to self-regarding interests and excludes or omits

other-regarding interests, then the criterion is much narrower. To the extent that people act on broader interests, this narrower criterion is neither necessary nor sufficient for competent performance.

In our view, Delli Carpini and Keeter equivocate on the point. They begin their discussion of political information and enlightened self-interest by invoking a concept of “enlightened preferences,” in which a person’s interest is whatever he or she would choose “with fullest attainable understanding of the experiences resulting from that choice” (p. 238; quoting Dahl, 1989, pp. 180–181). So defined, the concept is completely general with respect to motivation.

In the search for an indicator, however, they head for narrower ground. They note that interests or preferences are in part “socially constructed” (1996, p. 239) and often reflect group influences. To facilitate “an empirical study of interests,” therefore, they adopt the following logic:

If more-informed citizens are better able to discern their interests, and if material interests differ across groups in the population, it should be possible to detect the influence of information by comparing the opinions of better- and lesser-informed members of different groups (p. 239).

Specifically, they propose to look at whether knowledge: (1) “sharpens the differences between groups, moving their members closer to positions that are arguably consistent with their group norms and material circumstances”; (2) “encourage[s] consensus building,” moving citizens to positions that reflect a greater understanding of the circumstances of groups to which they do not belong; or (3) “move[s] mean opinion a significant amount to the left or right” (p. 239). The apparent implication is that any of these changes would indicate an enhancement of enlightened self-interest.

Delli Carpini and Keeter focus mainly, however, on the sharpening of differences. They select several sets of contrasting demographic groups (economically disadvantaged versus advantaged, black versus non-black, young versus old, and women versus men). They note that these groups have “important material and cultural differences . . . that should be and often are reflected in their expressed opinions about certain issues” (1996, p. 240). They then compare opinions of the contrasting demographic groups, conditioned on level of knowledge, using issues that are relevant for the group differences.

They find, for example, that more politically knowledgeable women support abortion rights more than less knowledgeable women do. Moreover, in a regression-based simulation of preferences, the predicted difference between women at the top of the knowledge scale and those at the bottom is markedly greater than that between such men. To Delli Carpini and Keeter, the findings suggest that knowledgeable women understand their self-interest better than

less knowledgeable women; and that in this case, at least, political knowledge promotes enlightened self-interest. In the end, therefore, the indicators of enlightened self-interest are very simple: support for abortion rights among women, and the difference in such support between women and men.

Despite the authors' fairly lengthy explanation, the rationale for this indicator is not entirely clear. On one interpretation, suggested by some of the language we have quoted,¹² their argument is circular: That is, they assume that more-informed citizens are better able to discern their interests. They then observe that more-informed women support abortion rights. And they infer that support for abortion rights is the enlightened self-interest position for women. So far, this works. But then they use the same observation about informed women supporting abortion rights to conclude that information helps people achieve enlightened self-interest. At this point, the single observation is doing inadmissible double duty: defining enlightened self-interest and explaining how people arrive at it. The difficulty arises from Delli Carpini and Keeter's objectives in this study. That is, they could plausibly use the positions of more-informed women to identify the enlightened self-interest position for women—provided they did not also intend to determine the effect of information on the ability to achieve enlightened self-interest. But that is indeed what they intend.

For their purpose the only coherent basis for the indicator—and undoubtedly their main intention—is simply to stipulate the content of enlightened self-interest on the basis of common knowledge about the effects of policies: for example, that women generally have a self-interest in abortion rights; that the economically disadvantaged have a self-interest in liberal domestic policies; and so on. It boils down to the researcher proposing a judgment about what policies benefit a group.

Such judgments are often fairly easy to make. Nevertheless, we have three reservations about such a stipulated group self-interest indicator. The most obvious is that the indicator is only as good as the stipulation of group interests. In the case of women and abortion rights, it is more tenuous than might appear. It is true that many women, at some time, have an abortion, want an abortion, or face a significant possibility of wanting one. All such women have a clear and direct interest in abortion rights. But many other women perceive no possibility of wanting an abortion, and oppose abortion rights on moral or religious grounds. To say that supporting abortion rights is in their self-interest appears unwarranted.

In addition, employing the indicator will often require questionable other-things-being-equal assumptions. Suppose, for example, that women have different moral beliefs about abortion than men, or that they attach more weight to those beliefs. Such differences could outweigh the differences in material interests that are used to stipulate a group-based self-interest.

Finally, and most fundamentally: The stipulated-group-interest approach presumes that citizens *ought* to decide on the basis of self-interest. As a Tom Wolfe character said, “Greed is good.” To liberal academics, that sounds fine when it suggests that women should favor abortion rights, that blacks should favor affirmative action, or that the economically disadvantaged should favor active government. But it is much less attractive when it suggests that white males should oppose affirmative action, that the economically secure should oppose social programs for the poor, and so on.¹³ In short, the group-self interest indicator is potentially useful but needs to be employed with considerable caution.

Indicators of “Getting It Right”

The political heuristics school generally has employed unadulterated versions of the broad, getting-it-right criteria of performance—making good decisions, deciding as if well informed, and so on. They have not sought to simplify analysis through such expedients as imposing restrictions on citizens’ motivation. However, they have dealt rather casually with the conceptual difficulties.

As we have noted, the getting-it-right criteria are infested with difficult conceptual issues. What is a good or a reasonable decision? What does it mean to be well informed, in the case of ordinary citizens? Do people who have the right information necessarily make good decisions? And if not, what additional conditions—concerning deliberative effort, thought processes, or other matters—are needed for a full statement of the criterion? These issues are theoretically difficult and partly normative. There is no prospect of reaching complete and specific consensus on such issues. But only a few scholars have even mentioned them directly, and then only briefly (Mondak, this issue; the most thorough discussion is Mutz, 1998, ch. 7).

Nevertheless, the main shortcoming of the political heuristics school’s approach to assessing competence has not been the lack of a fully elaborated conceptual definition of the getting-it-right criterion.¹⁴ It has been the failure to develop convincing empirical indicators for any respectable version of it. Despite the obvious challenges of developing such indicators, researchers have not invested major efforts in doing so. Instead, they have made sometimes-questionable inferences from readily available data and indicators.

In the simplest case, political heuristics studies have merely shown that citizens change preferences, in what is presumed to be the right direction, in response to a particular item of information—such as an endorsement by a citizens’ group or a popular president. The inference is that citizens use appropriate cues to decide competently. Such findings, however, do not demonstrate that they make reasonable decisions overall. Citizens may change preferences too much or too little for the significance of the information. Or they

may begin at an unreasonable starting point. To get decisions right, citizens must respond reasonably to the entirety of the information about a decision.

Political heuristics scholars are by no means alone in drawing broad inferences about competence from citizens' responses to particular information or stimuli (most of the research on political information does the same; see also Cobb and Kuklinski, 1997; Kuklinski and Quirk, 2000). In many circumstances, this approach offers advantages of convenience and analytic simplicity that scholars cannot easily overlook. Nevertheless, one cannot make a convincing assessment of competence—whether positive or negative—by analyzing responses to a single aspect of a complex decision environment. It requires assessing, on some basis, actual decisions—in effect, the summation of all the influences on them. The question is on what basis to assess them.

Political heuristics scholars, as well as others, have addressed this problem. In their most ambitious efforts to assess getting-it-right competence, they have employed what we have called well-informed-proxy group comparisons. In principle, the logic is appealing. The approach considers entire decisions, and thus takes into account all the influences that go into them. In practice, however, it has serious problems of its own.

The key to the approach is to find an appropriate well-informed-proxy group. Such a group should have two qualifications: First, the members should have values and interests similar to those of the focal group—ideally, the identical values and interests. Second, they should make, in some serious sense, well-informed and capable decisions on the issue at hand. The more informed and capable the proxy group is, the more rigorous is the resulting criterion. Unfortunately, however, it is exceedingly difficult to identify proxy groups that meet both these tests. What is crucial, therefore, is to think carefully about how well such a group measures up in each respect.

In fact, the proxy groups that scholars have used are problematic on both scores. The most common strategy has been to select the *better informed* group in a representative sample of citizens and define them, just on that basis, as well informed. Such a group has some weaknesses with respect to representativeness. More informed citizens might differ from others in a variety of relevant ways. Some of the differences—in income, education, and political ideology, for example—are likely to be measured and thus can be taken into account statistically at the individual level. But other differences, for example, in cultural values or cognitive styles, usually will not be measured. If these unmeasured differences have important effects on preferences, they would tend to confound comparisons between the two groups.

But the more serious shortcomings of these readily available proxy groups concern their information and capability for competent decision. The problem is obvious: Few scholars have ever suggested that any large fraction of ordinary citizens is, by any meaningful standard, well informed on a typical important

policy question. Even on generic NES-type issue questions (do more for the poor or do less, and the like), high-performers are those who have any opinion at all, prior to being asked a question (Zaller, 1992). On specific real-world proposals (the Bush administration's airport security proposal, or congressional Democrats' economic stimulus package), it is unlikely that more than a small fraction of citizens could summarize the principal facts and arguments presented in the media.

The use of such groups thus can easily produce misleading results. In perhaps the most celebrated study using a well-informed-proxy group in an analysis of political heuristics, Lupia's (1994) elegant analysis of the California referendum on automobile insurance, the proxy group was also essentially a group of relatively informed ordinary citizens—the top-scoring 30% of the sample on a quiz about the provisions of the propositions.¹⁵ He found that respondents who scored lower on the provisions but knew the industry position voted similarly to the well-informed group. Specifically, both groups massively supported the proposition that imposed a 20% rollback in insurance rates. He concludes that knowing the industry position enabled otherwise uninformed citizens to act as if they were well informed.

But the supposition that this proxy group was well informed (as opposed to merely relatively informed) seems to us very doubtful. It included all respondents who correctly answered more than half of the essentially true-false questions. More important, there is no evidence that these respondents had followed and understood the media debate about the effects of the propositions. Indeed most of this group arguably overlooked a crucial consideration: that a large mandatory price rollback for a competitive industry such as automobile insurance was predictably unworkable, and strongly opposed by the preponderance of reputable experts (Lascher and Powers, 1997).

The general point is that the choosing a well-informed-proxy group for comparison requires serious attention to the requirements for informed, deliberative decision. There is room for varying levels of rigor in defining them. But if competence implies that citizens make decisions that advance their basic interests or preferences, then a proxy group should have, in some respectable sense, the necessary information for such a decision. In particular, it should be reasonably informed about the effects of policies. Scholars have rarely addressed whether a proxy group actually meets such requirements.

DIFFERENT CHOICES, CONFLICTING CONCLUSIONS

There is no reason, of course, why scholars should necessarily agree on the citizen's task, the criteria by which to judge performance, the empirical indicators by which to measure the criterion, or the standards by which to place

people into one or another evaluative category. We should expect a range of choices across studies—exactly what we have found. However, if we also find that conflicting conclusions about citizen competence largely reflect these conceptual differences, one would have to wonder what to make of the collective product. If, in addition, the grounds for alternative choices are not clearly spelled out and debated, the uncertainty would be even greater. In this section, we identify two examples of such conflicts: one arising from different tasks, the other from different standards.

Converse Versus the Heuristics School

The rationale for Converse's ideological belief-system criterion is closely tied to his choice of the broad task of understanding politics, rather than voting, expressing opinions, or the like. He undoubtedly saw analytic advantages in selecting a task that affects performance in a variety of specific activities. And at the time of his writing, it might have seemed unexceptionable.

In recent years, however, the task of understanding politics has become controversial. The heuristics school argues that citizens can perform competently without retaining much information or understanding much about politics. In short, Converse's task is unnecessary. Ironically, these scholars appeal to ideology not as a mental apparatus that enables people to retain and integrate extensive information but, rather, as a labeling device that substitutes for such capability.

We are not suggesting that Converse and the heuristics school have fundamentally incompatible or even incommensurable approaches. For one thing, they might each have part of the truth: citizens might need some degree of understanding to use heuristics effectively (Delli Carpini and Keeter, 1996). For another, to the extent that their assumptions are in conflict, the issue may be subject to adjudication on the basis of evidence. For the time being, however, we are left with two contradictory conclusions: most citizens do not perform well at all and most citizens perform adequately.

Standards of Convenience

Scholars have rarely given direct attention to the matter of standards—that is, of what findings on an indicator show competent performance. When they have identified a standard, or implicitly used one to reach conclusions, their choices understandably have been driven in part by convenience. In using the issue-constraint indicator, for example, Converse compares ordinary citizens with candidates for Congress. The motivation for this seemingly very demand-

ing standard, as best we can surmise, is that measures of constraint were available for both citizens and candidates. Since it would be hard to decide whether citizens should have, say, two-thirds as much constraint as candidates, one-half, or some other fraction, the simplest position, and apparently the one that Converse adopts, is to demand a roughly equivalent level.

In rather dramatic contrast, Delli Carpini and Keeter offer an implicit standard, which they apply to more informed citizens, that they perform detectably better than the less informed. In their group self-interest analysis, for example, they find that more informed women show greater support for abortion than less informed women. More precisely, the finding is that this difference is statistically significant. (At least they mention no other threshold for the magnitude of the difference.) They count such findings as part of their evidence for the competence of the more informed. This seemingly very easy standard is again apparently a matter of convenience. It enables Delli Carpini and Keeter to assess competence by making comparisons within a single data set. It also helps them to support one of their central arguments: those with relatively greater information among American citizens are sufficiently informed to perform competently.

In Lupia's insurance referendum study (1994), the standard, which he applies to the less informed (with and without available cues), is that they make the same vote choices as the better informed. In this case, the absence of statistically significant difference is the apparent threshold. Although strict in one respect, it ends up an easy standard. To pass the test, cue takers must match the criterion group, as with Converse; but as we have seen, it is a criterion group that probably knew very little about the effects of the policy. Here again, the apparent motivation—again, understandably—is convenience. The standard permits an assessment of competence through comparisons within a single data set. And it helps support Lupia's argument that cue takers can perform competently.

In short, therefore, Converse demands that citizens match the most sophisticated elites; Delli Carpini and Keeter ask only that they do better than the uninformed; and Lupia asks that they do as well as the better informed. In no case does the standard rest on any direct consideration of what citizens must accomplish to advance their interests in the political system. Indeed, taken together, the latter two standards draw different conclusions from the same observation: For Delli Carpini and Keeter, a difference between the more informed and less informed would show that the more informed do well. For Lupia, it would show that the less informed do poorly. One study assumes that the less informed fail; the other that the more informed pass. In principle, however, it seems clear that both could fail or both pass.

In fairness, we see no easy answer to the question of standards. The authors' avoidance of direct discussion of the issue has been understandable. But it

has led them to use dramatically different standards, and draw correspondingly different conclusions, with no particular rationale.

LESSONS LEARNED

The authors whose works we have reviewed are among the best scholars in American politics; the problems we have identified cannot be attributed to poor scholarship. That we have found conceptual and empirical difficulties in three distinct literatures reflects, rather, the challenges that attend the study of citizen competence.

In this last section, we step back and reflect on some lessons that we have learned in the process of reviewing the various studies. Where we find it appropriate, we also offer recommendations for future research.

First, *the conceptual elements need explicit definition*. The framework we employed throughout this article—distinguishing task, criterion, indicator, and standard—is not the only one that scholars might offer. Our more important message is that the field would benefit from researchers choosing an appropriate set of elements and stating clearly how they propose to specify them for a particular study. What is the task, what is the criterion, and so on? Of course, acknowledging and taking responsibility for such decisions subjects a researcher to critical evaluation that might otherwise be avoided. But that is precisely our point: if we wish to move the study of citizen competence forward *conceptually*, such critical evaluations must occur. To date, scholars have focused most of their attention on empirical issues; and the exchanges have benefited everyone. We now need to do the same at the conceptual level.

Second, *thinking in terms of necessary and sufficient conditions helps bring conceptual problems to the fore*. In empirical matters, students of public opinion use probabilistic methods, as of course they must. However, thinking about concepts in terms of necessary and sufficient conditions can help researchers determine the limits of their criteria and measures. We have pointed out several instances where authors unwittingly consider only necessary or only sufficient conditions. This alone is not a serious problem; it becomes one when the researcher fails to see the implications of his or her choices for the conclusions that validly can be reached.

As we hinted earlier, we see no reason to apply these two conditions so restrictively that no study can get over the threshold. As general guidelines, however, they can serve a valuable purpose.

Third, *any choice of tasks has consequences*. We have found conceptions of tasks to vary not only across studies, but, in some cases, even within a single work (Delli Carpini and Keeter, 1996). In rather crude terms, scholars have pursued two approaches. Much of the scholarship has converged on discrete, specific tasks that citizens routinely perform and that have consequences for

the political system—such as voting for candidates, voting on referenda, and expressing opinions in polls, among others. Others define the task more broadly. Thus, Converse defines the task as understanding politics in a way that permits citizens to give meaning to events, retain information, and approximate rational behavior (or so we have argued).

Each form of task, specific and broad, has an advantage and a disadvantage. Specific tasks leave normative and conceptual issues to another stage (especially the choice of criteria); and their relevance is undeniable. On the other hand, choosing specific tasks arguably misses the forest for the trees. The strength of the broader conception is that it focuses on “the forest”: fundamental prerequisites to adequately performing any of the specific tasks. Yet, we also saw a weakness of this conception: a chosen criterion can embody potentially controversial theories about how people achieve competence in the specific tasks. We discussed, for example, how the heuristics school rejected the very task—understanding politics—that underlies all of Converse’s discussion.

The choice of task itself implicitly makes assumptions about how democracies do and should work. The more that scholars succeed in making these assumptions explicit, the easier it will be to understand why conclusions differ.

Fourth, *choosing criteria and indicators presents a dilemma*. In reviewing the three literatures, we identified a range of criteria—holding an ideological belief system, being tolerant of others, making good or reasonable decisions, acting in one’s self-interest, and choosing as if one were fully informed, among others. Unfortunately, there is a tradeoff between the conceptual adequacy of a criterion—its apparent necessity and sufficiency for competence—and its potential for valid measurement.

Narrower criteria, such as possessing an ideological belief system or tolerating others, are easier to measure empirically. But they also fall short with respect to necessity or sufficiency for competence, if not both. On the other hand, broad criteria—deciding reasonably, acting in one’s enlightened interest (not restricted to selfish interest), or deciding as if one is well informed—will almost always pass muster as necessary and sufficient. For many scholars, they are the very definition of competence.¹⁶ Unfortunately, however, any such criterion is extremely difficult to measure.

Scholars can take and have taken different attitudes toward this dilemma. In our view, however, the balance of advantages generally favors the broader, even if less measurable, getting-it-right criteria.¹⁷ In the first place, such criteria seem capable of eliciting general agreement among empirically oriented public-opinion scholars. To our knowledge, such scholars have offered virtually no criticism of the criterion as a matter of principle. Moreover, the problems with indicators for these criteria are not entirely intractable. Scholars can probably find or create more suitable (especially better informed) well informed-proxy groups. And they can bring to bear additional methods. Snider-

man and Theriault (1999; also see Druckman, this issue), for example, have measured correlations between individuals' values and their policy preferences. In fact, the multiple approaches suggest the possibility of strengthening findings through the use of multiple indicators within and across studies.

Finally, *setting standards is the biggest challenge of all*. When it comes to defining standards, there is an inevitable problem of half-full, half-empty. How competent is competent enough? The question becomes especially crucial if the criterion is a potentially very rigorous one, such as acting as if one is well informed. How closely should we expect ordinary citizens to match the decisions of a suitable, well-informed proxy group? At the limit, how much should their decisions resemble, say, those of experts or candidates for Congress? To be more specific, what magnitudes of correlation among citizen issue positions would have—and should have—led Converse to conclude that citizens perform adequately?

The choice of standards poses special difficulties. It depends on all of the other conceptual decisions; one cannot even talk about standards until tasks, criteria, and indicators are set. Moreover, it turns in part on questions that concern feasibility, that are largely normative, or that go beyond the public-opinion field: How much competence is plausibly attainable? How much of what is attainable should we insist on? And what difference do various levels of competence make for political outcomes? For example, how much incompetence among citizens does it take to do real damage in a representative democracy? Unlike the case with criteria, we cannot expect anything approaching consensus with respect to standards anytime soon. At most, we can expect that scholars reveal their presumptions about standards, and not pretend that conclusions about competence are somehow independent of any standard.

A FINAL LESSON

Converse, Delli-Carpini and Keeter, Lupia and McCubbins, Mutz, Popkin, Sniderman—these are just some of the authors whose work we have discussed in the preceding pages. The citations of these names alone provide our final lesson: we know far more about citizen competence today than we did 40 years ago. Most important, we think, is the introduction of ideas that did not exist before publication of the various works. Terms like heuristic, belief system, impersonal influence, political knowledge, and political sophistication are now commonplace. Each represents a new and alternative way to think about citizen performance. Moreover, scholars have introduced creative and highly rigorous methodologies that have given the field insights it did not have before. As a field, we should take pride in our accomplishments; and we have every reason to hold high expectations for the next generation of studies.

Acknowledgments. This article benefited greatly from the comments that Brian Gaines, David Darmofal, and Diana Mutz made on an earlier version.

NOTES

1. The other three articles in this volume are responses to a call for a special issue on citizen competence. One deals with knowledge, another with framing, and a third with a myriad of topics on public opinion.
2. At the other extreme, an owner or a committee making the most valuable player award might be interested in overall performance on the court. The most likely approach would be to identify all of the relevant criteria and appropriate empirical indicators and then come up with a composite measure.
3. In his recent work, Ragin (2000) presents methods to test for necessary and sufficient conditions in terms of probabilities.
4. All three topics focus entirely and only on citizen performance with respect to public opinion. As Weissberg (this issue) observes, this is only a part of what falls under the so-called citizen competence literature.
5. This means that we inevitably will not cite studies that are just as deserving as those we do cite. Moreover, we are only examining political science research; economists and sociologists have addressed similar questions (see, for example, Camerer, 1995).
6. One conceivably could argue that these are the tasks Converse has in mind. However, citizen understanding clearly is more prominent in his overall discussion. To the extent that Converse refers to voting and the like at all, it is in passing.
7. He does not say whether consistent moderation would indicate an appropriate belief system or rather the lack of one.
8. However, this leaves the question, what correlations among citizens would represent minimum competence? Would a correlation of, say, .33 suffice for domestic issues? Setting aside specific numbers, must citizens have roughly the same degree of constraint as congressional candidates to be judged competent, or would some significantly lower figure suffice? It is tempting to say that explicit standards are not necessary as long as the correlations are reported. The crux of the matter, though, is that these correlations must be interpreted to reach a conclusion. And the interpretation requires a standard, even if it is not acknowledged.
9. Obviously, some of these concepts are more interchangeable, on their face, than others.
10. Of course, Converse spends considerable time earlier in his essay demonstrating that people do not understand basic American political ideology. Our critique is not directed at Converse's overall work, which is impressively thorough, but at the particular criterion of holding an ideological belief system. This is undoubtedly the criterion to which scholars have given the most attention over the years.
11. And there would be little if any change in public preferences. Citizens would draw on their belief systems, as Converse defines them, at the outset and stay there. Where public debate would come into play is not clear.
12. See the indented quote above, which seems to propose both an assumption that information promotes enlightened self-interest and a plan for testing the effect of information.
13. One might attempt to modify the criterion by saying that only the less privileged should act on self-interest. But that would merely amount to adopting an overt political position and stipulating that everyone should work for social equality.
14. Using a creative scrolling technique, Lau and Redlawsk (2001) find evidence that people use heuristics, but often with bad results. Their finding is in line with the dominant view in social psychology.

15. The authors calculated the size of the group from information provided in Table 4 and footnote 14.
16. There are two important exceptions. These criteria implicitly bracket out citizens' choices of values. Some scholars want to include certain particular values as part of the definition of competence. Others want to consider competence in the choice of values. In either case, deciding as if fully informed would not be a sufficient condition for competence.
17. The political heuristics school has played an important role in shifting attention toward the broadest criteria of citizen competence, deciding as if well informed. It has taken this tack, among other reasons, to leave citizens' methods of accomplishing that result unspecified.

REFERENCES

- Achen, C. H. (1975). Mass political attitudes and the survey response. *American Political Science Review* 69: 1218–1231.
- Bartels, L. M. (1996). Uniformed voters: information effects in presidential elections. *American Journal of Political Science* 40: 194–230.
- Brady, H. E., and Sniderman, P. M. (1985). Attitude attribution: a group basis for political reasoning. *American Political Science Review* 79: 1061–1078.
- Camerer, C. (1995). Individual decision making. In J. H. Hagel and A. E. Roth (eds.), *The Handbook of Experimental Economics*, pp. 587–616. Princeton NJ: Princeton University Press.
- Carmines, E. G., and Kuklinski, J. H. (1990). Incentives, opportunities, and the logic of public opinion in American political representation. In J. A. Ferejohn and J. H. Kuklinski (eds.), *Information and Democratic Processes*, pp. 240–268. Urbana: University of Illinois Press.
- Cobb, M. D., and Kuklinski, J. H. (1997). Changing minds: political arguments and political persuasion. *American Journal of Political Science* 41: 88–121.
- Converse, P. (1964). The nature of belief systems in mass publics. In D. Apter (ed.), *Ideology and Discontent*, pp. 206–261. New York: Free Press.
- Dahl, R. A. (1989). *Who Governs? Democracy and Power in the American City*. New Haven: Yale University Press.
- Delli Carpini, M. X., and Keeter, S. (1996). *What Americans Know About Politics and Why It Matters*. New Haven, CT: Yale University Press.
- Druckman, J. N. (2002). The implications of framing effects for citizen competence. *Political Behavior* 23: 225–256.
- Erikson, R. (1979). The SRC panel data and mass political attitudes. *British Journal of Political Science* 9: 89–114.
- Jackson, J. E. (1979). Statistical estimation of possible response bias in close-ended issue questions. *Political Methodology* 6: 393–424.
- Judd, C., and Milburn, M. (1980). The structure of attitude systems in the general public: comparison of a structural equation model. *American Sociological Review* 45: 627–643.
- Kahneman, D., Slovic, P., and Tversky, A., eds. (1982). *Judgment Under Uncertainty: Heuristics and Biases*. New York: Cambridge University Press.
- Kuklinski, J. K., and Quirk, P. J. (2000). Reconsidering the rational public: Cognition, heuristics, and public opinion. In A. Lupia, M. D. McCubbins, and S. L. Popkin (eds.), *Elements of Reason: Cognition, Choice, and the Bounds of Rationality*, pp. 153–182. New York: Cambridge University Press.
- Kuklinski, J. K., Quirk, P. J., Jerit, J., Schwieder, D., and Rich, R. F. (2000). Misinfor-

- mation and the currency of democratic citizenship. *The Journal of Politics* 62: 790–816.
- Lane, R. (1962). *Political Ideology*. New York: Basic.
- Lascher, E. L., and Powers, M. R. (1997). Expert opinion and automobile insurance reform: an empirical assessment. *Journal of Insurance Regulation* 16(2): 197–222.
- Lau, R. R., and Redlawsk, D. P. (2001). Advantages and disadvantages of using cognitive heuristics in political decision making. *American Journal of Political Science* 45: 951–971.
- Lodge, M. G., and Hamill, R. (1986). A partisan schema for political information processing. *American Political Science Review* 83: 399–419.
- Lodge, M., and Taber, C. (2000). Three steps toward a theory of motivated political reasoning. In A. Lupia, M. D. McCubbins, and S. L. Popkin (eds.), *Elements of Reason*, pp. 183–213. New York: Cambridge University Press.
- Lupia, A. (1994). Shortcuts versus encyclopedias: information and voting behavior in California insurance reform elections. *American Political Science Review* 88: 63–76.
- Lupia, A., and McCubbins, M. D. (1998). *The Democratic Dilemma: Can Citizens Learn What They Need to Know?* Cambridge: Cambridge University Press.
- Luskin, R. (1987). “Measuring political sophistication. *American Journal of Political Science* 31: 856–99.
- Luskin, R. (1990). Explaining political sophistication. *Political Behavior* 12: 331–361.
- Luskin, R. (2001). Political psychology, political behavior, and politics: questions of aggregation, causal distance, and taste. In J. H. Kuklinski (ed.), *Thinking About Political Psychology*, pp. 217–252. New York: Cambridge University Press.
- Mondak, J. J. (1993). Source cues and policy approval: the cognitive dynamics of public support for the Reagan agenda. *American Journal of Political Science* 37: 186–212.
- Mondak, J. J. (1994). Cognitive heuristics, heuristic processing, and efficiency in political decision-making. In M. Delli Carpini, L. Huddy, and R. Y. Shapiro (eds.), *Research in MicroPolitics*, vol 4, pp. 84–107. Greenwich, CT: JAI Press.
- Mondak, J. J. (2002). Asked and answered: knowledge levels when we will not take “Don’t Know” for an answer. *Political Behavior* 23: 199–224.
- Mutz, D. C. (1998). *Impersonal Influence: How Perceptions of Mass Collectives Affect Political Attitudes*. New York: Cambridge University Press.
- Nisbett, R., and Ross, L. (1980). *Human Inference: Strategies and Shortcomings of Social Judgment*. New York: Prentice-Hall.
- Piattelli-Palmarini, M. (1994). *Inevitable Illusions: How Mistakes of Reason Rule Our Minds*. New York: John Wiley.
- Popkin, S. L. (1991). *The Reasoning Voter*. Chicago: University of Chicago Press.
- Ragin, C. (2000). *Fuzzy Set Social Science*. Chicago: University of Chicago Press.
- Sniderman, P.M., Brody, R. A., and Tetlock, P. E. (1991). *Reasoning and Choice: Explorations in Political Psychology*. New York: Cambridge University Press.
- Sniderman, P. M., and Theriault, S. M. (1999). The dynamics of political argument and the logic of issue framing. Paper Presented at the Annual Meeting of the Midwest Political Science Association, Chicago, IL, April 15–17.
- Weissberg, R. (2002). Democratic political competence: clearing the underbrush and a controversial proposal. *Political Behavior* 23: 257–284.
- Zaller, J. (1990). Political awareness, elite opinion leadership, and the mass survey response. *Social Cognition* 8: 125–153.
- Zaller, J. (1992). *The Nature and Origins of Mass Opinion*. New York: Cambridge University Press.