

UC Berkeley

UC Berkeley Electronic Theses and Dissertations

Title

Information Effects in Politics and Public Policy

Permalink

<https://escholarship.org/uc/item/80r7058z>

Author

Kelly, Andrew Patrick

Publication Date

2013

Peer reviewed|Thesis/dissertation

Information Effects in Politics and Public Policy

By

Andrew Patrick Kelly

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Political Science

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Paul Pierson, Chair
Professor Robert Van Houweling
Professor Anne Joseph O'Connell

Fall 2013

Information Effects in Politics and Public Policy

© 2013

By Andrew Patrick Kelly

Table of Contents

Chapter 1 Introduction	1
Chapter 2 Do Elections Inform Voters? Evidence from the 2006 Senate Elections	15
Chapter 3 Policy Performance and Public Support	38
Chapter 4 Increasing College Knowledge: Information Effects in Social Policy Markets	60
Chapter 5 Conclusion	79
References	93
Appendix	105

List of Figures

Figure 3.1: Mean Responses and Treatment Effects, By Ideology 57
Figure 5.1: Effect Sizes Across Studies 81
Figure 5.2: Effects By Educational Attainment 85

List of Tables

Table 2.1: Citizen Accuracy on Senator’s Party Identification and Roll Call Positions..... 23
Table 2.2: Citizen Accuracy on Party Identification by Electoral Context..... 24
Table 2.3: Regression Model of Citizen Accuracy on Party Identification 25
Table 2.4: Citizen Accuracy on Roll Call Issues by Electoral Context 27
Table 2.5: Regression Models of Citizen Accuracy on Roll Call Issues 28
Table 2.6: Citizen Accuracy across Electoral Conditions and Type of Vote..... 30
Table 2.7: Regression Models of Citizen Accuracy by Vote Type..... 31
Table 2.8: Within Subjects Comparisons (Pooling Across Issues)..... 33
Table 2.9: Projection on Defecting Votes by Electoral Condition..... 35
Table 3.1: Respondent’s Perceptions of Food Stamp Performance 46
Table 3.2: Respondent’s Perceptions of Head Start Performance 46
Table 3.3: Effect of Food Stamp Error Rate Information on Outcome Variables 48
Table 3.4: Effect of Food Stamp Error Rate Information on Outcomes, By Ideology 49
Table 3.5: Effect of Head Start Evaluation Information on Outcome Variables 51
Table 3.6: Effect of Head Start Evaluation Information on Outcomes, By Ideology 52
Table 3.7: Effect of Head Start Fraud Information on Outcome Variables 53
Table 3.8: Effect of Head Start Fraud Information on Outcomes, By Ideology 54
Table 3.9: Effect Sizes for Each Type of Information 55
Table 4.1: Descriptive Statistics and Balance 68
Table 4.2: Proportion of Parents Who Felt Very or Fairly Well Informed..... 69
Table 4.3: Proportion of Parents Rating Each Dimension as Important 70
Table 4.4: Information Effects (Differences in Means) 71
Table 4.5: Regression Models of Treatment Effects..... 72
Table 4.6: Treatment Effects By Level of Education..... 74
Table 4.7: Treatment Effects By Self-Reported Level of Information 75

This dissertation is dedicated to my wife Lindsay, who was rooting me on from start to finish.

Abstract

Information Effects in Politics and Public Policy

By

Andrew Patrick Kelly

Doctor of Philosophy in Political Science

University of California, Berkeley

Professor Paul Pierson, Chair

For more than fifty years, political scientists have questioned whether citizens are sufficiently informed about politics to ensure representation and democratic accountability. On the one hand, surveys suggest that most voters are poorly informed about politics and policy, have inconsistent policy preferences, and suffer from deep-seated biases that hinder rational political judgment. On the other, aggregate studies of representation and collective opinion have produced evidence that voters do hold officials accountable when they fail to represent their interests, and that public opinion tends to reflect national events.

These micro and macro patterns present something of a puzzle: how can political accountability exist when overall levels of citizen competence are low? More recent research has argued that while citizens are not encyclopedias of political knowledge, they can often learn what they need to know to make rational choices. These studies focus on the dynamics of political learning and the role of the information environment in providing voters with useful cues and policy information. When examined through this lens, citizens often appear more capable than was previously thought.

But important questions about information effects and political learning remain. Some studies find that citizens will update their beliefs in predictable ways when they learn something new. Others find that political misperceptions are hard to erase, even in the face of efforts to correct them. Under what conditions will citizens learn from new political information? Are some messages more potent than others? And do simple cues always help voters navigate the political world?

This dissertation contributes to these debates by examining information effects across three different venues: electoral politics, attitudes toward government programs, and social policy markets. Using three different data sources, I explore whether the political environment helps citizens to become more informed and whether new information prompts them to revise their preferences and perceptions in rational ways. I begin with intuitions from a rational learning model that help to explain the divergent findings in the literature and serve as a guide for the empirical chapters that follow.

I find considerable evidence of information effects across all three venues. Citizens in the context of an election campaign are more informed about the party identification of their senator, and this party cue helps them to make inferences about the senator's issue positions. Importantly, however, this increase in accuracy is only evident on votes where the senator voted the party

line. When senators defected to vote with the other party, citizens' reliance on party cues led them astray. Similarly, I show how providing citizens with information about the performance of social programs like Food Stamps and Head Start has a significant effect on their perceptions and preferences. Negative performance information appears to be more potent than positive, but the findings show no evidence of ideological bias. Finally, I find that citizen-consumers in the market for public colleges use additional information about student outcomes to update their preferences.

The dissertation concludes by returning to the rational learning model in an effort to reconcile these findings with those in the existing literature. I argue that the divergent findings are largely due to different assumptions about the precision of new political information, the strength of prior beliefs, and the probability that citizens receive political messages.

Chapter 1 Introduction

American politics and government are more transparent than ever before. Anyone with a smart phone can look up the day's roll call votes to see how their member of Congress voted on a controversial issue. In the middle of a heated campaign, voters can rely on fact-checking websites to see whether the President's latest claim was mostly true, half-true, or "pants-on-fire." Deficit hawks can access federal evaluations of social programs and cry foul when officials continue to fund those found to be ineffective or inefficient. And the beneficiaries of public services can use government-issued data to compare different providers before deciding where to invest their public benefits.

In theory, the trend toward a more transparent, information-rich politics should empower citizens to exert better control over their government. As President Obama (2009) explained to the heads of federal agencies shortly after taking office, "transparency promotes accountability and provides information for citizens about what their government is doing" (1). Such "openness," Obama (2009) wrote, "will strengthen our democracy and promote efficiency and effectiveness in government" (1). The President is the latest champion of a longer-term trend toward openness that began in the 1970s and has picked up speed in the digital age. The recorded teller vote, the Freedom of Information Act, campaign finance reform, randomized evaluations of social programs, and the proliferation of third-party fact-checkers, data clearinghouses, and rankings all spring from this push for transparency.

The link between transparency and democratic accountability fits with classic principal-agent models of political representation and service provision. In these models, the principal's key challenge is monitoring the agent's behavior. Making the behavior of agents more transparent—reporting on the votes they cast, the performance of the programs they oversee, or the quality of the public service they provide—makes it less costly to monitor, thereby enhancing constituents' ability to hold agents accountable for their performance. Less costly monitoring should, in turn, create incentive for agents to be more responsive to constituent preferences. Don't like what you learned about your senator's roll call votes on taxes? Vote for the challenger. Does your child's school rate poorly on math and reading tests when compared to other schools in the district? Raise a stink at the next school board meeting. If that fails and a choice exists, send your child elsewhere. Even in the absence of information, the threat of moving from one agent's column to that of a competitor is present so long as constituents have a choice. But choice plus information about agent behavior enhances citizens' ability to monitor, punish, and reward. The more informed the principals, the riskier it is for agents to ignore them.

This is how transparency should work in theory, anyway. In reality, this tidy narrative depends on the ability and motivation of citizens to fulfill their monitorial duties. They must then process the additional information they receive, combining it with prior beliefs and arrive at a new position that reflects their preferences. They must then act on that revised belief, choosing to stick with the status quo or change course.

Research from political science and behavioral economics provides reason to doubt the narrative at each juncture. Though information is plentiful, Americans are often poorly informed about politics, policy, and government services (Miller and Stokes 1963; Converse 1964; Delli Carpini and Keeter 1996). Similarly, when faced with a choice of schools, health plans, or retirement accounts, citizens often fail to capitalize on available information that would help them find the best option (Thaler and Sunstein 2009). And even when they do encounter data

that could help them make a decision, citizens' prior beliefs, party loyalties, and other biases can lead them to marginalize that information and adhere to existing misperceptions (Bartels 2002).

Optimistic forecasts about the power of transparency to improve public accountability and government responsiveness must contend with reams of research on the frailties of political judgment. Indeed, the growing emphasis on informational interventions in public policy—the idea that we can “nudge” people to make better decisions—is built on this realization (Thaler and Sunstein 2009). However, if proponents of open government and market-based policies have been Pollyannaish, the critique of citizen competence has often gone too far in the opposite direction. Critics look at low levels of political knowledge and assume that most citizens are not capable of the kind of reasoning healthy democracies and healthy markets require. If that were true, though, why do we see aggregate evidence that voters often oust legislators who are too extreme? That public opinion tracks with national and global events? That families flock to the best public schools and school districts?

Both the optimistic and pessimistic storylines are also under-specified. In the indictment of citizen competence, for instance, are we to believe that information processing and political judgment are equally suspect across all settings and decision-making contexts? Or are there conditions under which new information can affect citizen beliefs and preferences in predictable and rational ways? Transparency hawks make similar errors of oversimplification and generalization, assuming that better data alone will cure all manner of information problems. They make this argument despite copious evidence that simply providing information hardly ensures that citizens will use it. These arguments also ignore the likelihood that information effects are likely heterogeneous. Which citizens will benefit most from transparency? Will it be those who are already well informed, or the ones who know the least about politics and policy? Will the knowledgeable get even more sophisticated, potentially exacerbating existing inequalities in political voice and access? And are there domains in which information effects are more or less likely?

In an effort to shed some light on these unanswered questions, this dissertation takes an empirical look at information effects across three domains: electoral politics, attitudes toward social programs, and social policy markets. First, I use the staggered terms of the U.S. Senate to assess whether voters are better informed about elected officials who are up for reelection. I also ask whether improvements in accuracy reflect simple cue taking rather than deeper learning and explore the limits of cognitive shortcuts. Second, I use data from an experimental survey to ask whether messages about the performance of two federal social programs—Food Stamps and Head Start—affect individuals' perceptions and policy preferences. I then document how those effects vary across different political groups and different types of information. And third, I examine how additional information can affect preferences in one government-funded social policy market: public colleges and universities. I then look at how those information effects vary by socioeconomic status. The concluding chapter summarizes the findings and attempts to reconcile them with the existing research.

The goal here is not to provide an up or down verdict on whether enhanced transparency will or will not strengthen democratic governance and public accountability. Instead, the hope is that by looking at information effects across different domains, respondents, and political judgments, this study can help to clarify both the opportunities to create a better-informed citizenry and the likely limits of those efforts. These findings can then inform both future political science research as well as contemporary policy debates.

In this introductory chapter, I begin by providing a brief snapshot of the changing informational environment in each of the three domains studied. I then describe some of the prior research on information effects in politics and public policy and lay out a stylized rational learning model that guides the chapters that follow. I conclude with a preview of what is to come.

A More Transparent State

Persistently disappointing levels of citizen competence are particularly curious in light of long-term trends in both educational attainment and the information environment. The worlds in which election campaigns are waged, public policies are made and evaluated, and social services are delivered have changed dramatically since the first studies of political awareness were published. For one thing, Americans have become far more educated since the 1958 congressional election, when Miller and Stokes' (1963) famously found “not more than a chemical trace” that constituents knew the first thing about their elected representatives (54). Between 1959 and 2012, the percentage of Americans who had finished four years of college increased from about 8 percent to 31 percent; high school completion doubled from well under half (44 percent) to 88 percent over the same period (U.S. Census 2012).

While Americans were getting more highly educated, changes within and outside of government have made information about politics and public policy more widely available than ever before. Three examples stand out.

Roll Call Voting and Position Taking: Since the 1970s, it has become increasingly easy for constituents to learn about the voting behavior of individual members of congress. The Legislative Reorganization Act of 1970 required that roll calls in the House of Representatives' Committee of the Whole be recorded and made publicly available. Before the recorded teller reform, members could vote one way on penultimate votes in the Committee of the Whole and the opposite way on final passage. In the mid-1990s, the new Republican majority expanded these recording requirements to cover roll call votes taken by committees. According to Ferejohn (1999), these efforts to open up the policymaking process had clear consequences in congress: roll call participation doubled, and voting on committee roll calls rose from 40 percent to 90 percent.

More detailed roll-call voting records have created opportunities for new intermediaries to collect, simplify, and publicize position-taking information in ways that are easy to understand. Interest groups were the earliest adopters, publishing “report cards” that graded individual members of congress according to the votes they cast on important issues. What started with a handful of groups in the 1970s has ballooned over time; Project Vote Smart lists 365 interest groups that have issued ratings since 2008. And media groups have taken on political fact-checking duties to keep candidates and elected officials honest. Each day, sites like PolitiFact, FactCheck.org, and countless others assess the factual accuracy of statements from elected officials, challengers and policy advocates.

The Effectiveness of Government Programs: In the early 1970s, Rivlin (1971) famously argued that the growing federal role in social policy was hamstrung by a lack of rigorous research on what programs were working, which were not, and how to improve existing programs. “Little is known about how to produce more effective health, education, and other social services,” Rivlin (1971, 7) wrote, “and neither social service systems nor federal programs are organized to find out.”

That would change. Advances in social science—the adoption of randomized controlled trials from medicine in particular—allowed researchers to measure the effects of large social

programs directly (Baron and Haskins 2011). Congress increasingly built evaluation requirements into legislation, and the results of those studies helped to guide future rounds of policymaking. Though many large programs have escaped scientific evaluation up to now, evidence of effectiveness (or the lack thereof) has become a common feature of social policy debates. Over the past thirty years, a dozen large-scale federal programs have been subjected to rigorous evaluation, including Head Start, Job Corps, Upward Bound, Welfare-to-Work, and Children's Health Insurance. Program outcomes have generally disappointed, discouraging advocates and providing rhetorical ammunition to fiscal conservatives intent on cutting the budget.

Like them or not, program evaluations exert increasing influence on budgetary politics in an era of federal deficits. In the welfare reform movement of the 1990s, for instance, it was randomized studies of state-level innovations that laid the groundwork for the federal efforts that came later (Rogers-Dillon 2004). They have also become a favorite presidential management tool. President George W. Bush's Program Assessment Rating Tool (PART) sought to steer money toward programs with evidence of success. President Obama's Office of Management and Budget (OMB) officially told agencies that their budget requests should be evidence-based.

Informing Consumers of Government Services: Public accountability is not only exerted at the ballot box; increasingly, citizens are empowered to hold the providers of government services accountable by "voting with their feet." Market-based approaches to social policy—which allow eligible citizens to choose a service provider from among a set of public and private organizations—have taken root in education, healthcare, job training, housing, and other areas. Political scientists have closely examined the behavior of citizens in these markets for local public goods, casting decisions to "exit" or "enter" as acts of political participation in their own right (Hirschman 1970; Lowery & Lyons 1989; Teske et al. 1993). By choosing to move from one town or school to another, citizens signal their dissatisfaction with the existing level of service. In the aggregate, these exits put pressure on public officials to be more responsive to the needs of their constituents. In this way, mobility and choice provide citizens with another mechanism by which to hold governments and public organizations accountable.

As these quasi-markets have expanded, so too has the recognition that citizens are often poorly informed about their options. As is the case in politics, uninformed citizens can be vulnerable to manipulation and make choices that distort competitive pressures. In response, governments have begun to publish more data on the quality and cost of different public service options. Every state in the country now reports the percentage of students that are proficient in math and reading in each of their public schools and disaggregates those data by demographic groups. At the federal level, policymakers in the executive branch have built report cards like Medicare Hospital Compare and the College Scorecard to report data directly to citizens in a consumer-friendly way.

Falling on Deaf Ears? Information and Political Behavior

These efforts to increase government transparency are often popular and bipartisan. Being anti-transparency is not a political winner, particularly when trust in government is at historic lows. And these efforts are almost certain to help some citizens arrive at better decisions in some domains. But is all this additional information going to improve the health of our democracy? That depends on whether citizens actually receive the new information and can use it in rational ways. On this question, a sizable body of research provides little reason for optimism.

Political science's most enduring finding is arguably its most discouraging: the average citizen is poorly informed on questions of politics and public policy. The list of shortcomings (and studies) is long: citizens struggle to come up with the names and issue positions of their elected representatives; they know even less about policy issues and public problems and often harbor stubborn misperceptions; and they even have trouble expressing coherent and consistent policy preferences of their own (Campbell et al. 1960; Miller and Stokes 1963; Converse 1964; Luskin 1987; Delli Carpini and Keeter 1996; Bartels 1996; Kuklinski et al. 2000; Gilens 2001). In their sweeping study of 2,000 survey items measuring political knowledge over fifty years, Delli Carpini and Keeter (1996) found that just forty percent of those items were answered correctly by more than half of respondents.

More telling, perhaps, is that these patterns have been stable for the last half-century, despite significant gains in educational attainment, information availability, and communications technology (Delli Carpini 2005). Indeed, generation after generation of political scientists has come to similar conclusions: "where political information is concerned, the mean is very low but the variance is very high" (Converse 2000, 331); "The political ignorance of the American voter is one of the best-documented features of contemporary politics" (Bartels 1996, 194); "The failure of voters to meet the demands of even relatively modest theories of political participation highlights the severity of the challenge of political ignorance for democratic theory" (Somini 2013, 13). According to Converse (2000), these patterns are so firmly established that "we hardly need argue low information levels anymore" (331).

Citizens are not only under-informed; they are often misinformed and unwilling to update their beliefs. What citizens think they know about politics, policies and trends in economic and social conditions is often systematically incorrect, colored by partisan and ideological biases, racial attitudes, and faulty heuristics (Bartels 2002; Gilens 1996; Gaines et al. 2007). For instance, partisans often believe the economy improved under a president from their party, regardless of the actual trend (and vice versa for presidents from the opposing party) (Gerber and Huber 2010). Citizens often view the beneficiaries of welfare programs as "undeserving" and overestimate the proportion of welfare recipients that are African American (Gilens 1999; Kuklinski 2000). And though Americans assign low ratings to the public schools as a whole, they usually give their local public school high marks (Loveless 1997).

When it comes to political learning, these misperceptions and information deficits are not easily overcome, even when researchers try to erase them. For one thing, the least politically sophisticated are often the least likely to receive new information, while the most aware may also be the best equipped to reject it (Zaller 1992; Lodge and Taber 2000). Taber and Lodge (2006) find evidence that citizens encounter new information with "motivated skepticism" (755). Survey respondents evaluate arguments they agree with as being stronger than those they disagree with. And, when given a choice, citizens often choose sources with which they are likely to agree (Taber and Lodge 2006). As a result, new information often polarizes opinions of those who disagree at the outset rather than leading them to converge. Experimental research has shown that prior beliefs and party cues tend to swamp the effect of policy information, leading respondents to choose candidates on the basis of partisanship rather than policy representation (Rahn 1993; Cohen 2003). Even when experiments explicitly set out to correct common misperceptions, partisans hang onto their inaccurate beliefs, and some become even more convinced of their accuracy (Bullock 2007; Nyhan and Reifler 2010).

These findings pose a fundamental challenge to the rosy claims of the open government movement, and to political accountability more generally. As Bullock (2011) remarked recently,

“if people ignore facts about policy even when exposed to such facts, there is little reason to expect that facts will help them to make better decisions or protect them from manipulation by elites” (496).

Alternative Takes on Information and Citizen Competence

Naturally, researchers have asked how representation and basic governance could exist in this morass. After all, ideal models of democracy rely on citizens to monitor the behavior of government officials and then hold them accountable at the ballot box. If they lack the ability or motivation to access and process that information, the ideal model breaks down—quickly.

But aggregate evidence suggests that the electorate *does* hold officials accountable for their policy positions (Canes-Wrone, Brady and Cogan 2002; Ansolabehere, Snyder and Stewart 2001), the performance of the economy (Fiorina 1978), and the quality of the public services they provide (Berry and Howell 2007). Recent elections have produced predictable but extreme swings in the ideological complexion of Congress (Bafumi and Herron 2010). And studies of collective opinion have shown that public attitudes and preferences respond to national events and trends in rational, predictable ways (Stimson, MacKuen and Erikson 1995; Page and Shapiro 1992).

These patterns present a puzzle: how can poorly informed voters exercise meaningful political accountability? A growing body of research suggests that under the right circumstances, citizens can learn what they need to know to make reasoned political choices. Rather than asking “are citizens informed” across different issues, candidates, and levels of government, this strand of work has asked “under what conditions can citizens learn what they need to know?”

Heuristics and cue-taking provide one potential answer. Drawing on insights from psychology and behavioral economics, this literature argues that by using simple cues and rules of thumb, citizens can arrive at reasoned political choices in spite of their overall lack of political knowledge (Sniderman, Brody, and Tetlock 1991; Popkin 1991; Lupia 1994; Lupia and McCubbins 1998). The argument is not that citizens are more informed, per se, but that cues like party labels, endorsements from interest groups, or the race and gender of candidates can provide voters with a shortcut to the rational choice. In the world of political heuristics, party cues tend to reign supreme. Research has documented the important influence these cues have on citizens’ inferences about the policy positions of candidates (Conover and Feldman 1989) and the likelihood that policy information will affect preferences for candidates (Rahn 1993). That citizens regularly rely on heuristics to navigate the political world is clear. But questions remain as to whether these shortcuts invariably lead to reasoned political choices (Lupia 1994) or to misperceptions and faulty inferences (Kuklinski and Quirk 2000), and the conditions under which either outcome is more likely.

Other studies have taken a fresh look at how voters learn about policy issues and candidates using longitudinal and experimental research designs. In his study of campaign effects across four different cases, Lenz (2009) argues that what may look like the short-term “priming” effect of campaigns actually reflects a longer-term process of learning and opinion change. In the midst of a campaign, Lenz (2009) argues, some constituents learn about the position of their party on a policy issue and then adopt that position as their own. Though Lenz (2009) describes this process of party-driven attitude change as being less “normatively appealing” than attitude-based party change, the pattern constitutes political learning just the same (834).

But other evidence suggests policy information can have a direct, independent effect on policy attitudes as well. Indeed, a handful of studies have argued that “policy-specific

information” belongs at the center of research on political knowledge and learning (Gilens 2001). Gilens (2001) makes a convincing case that facts about policy design, cost, and outcomes have more influence over policy preferences than general political knowledge. Nor are the two sets of facts always highly correlated. Political sophisticates may lack policy-specific information, while members of so-called “issue publics” can be quite well informed on particular policy issues even in the absence of general political knowledge (Krosnick 1990; Hutchings 2003).

The emphasis on policy information has broadened notions of what citizens know and can learn about politics. In contrast to earlier studies of representation, for instance, Ansolabehere and Jones (2010) find that most constituents are able to place their senator on specific roll call votes. Macro studies of citizen beliefs and media coverage have found that knowledge of particular policy issues reflects the “information environment:” the more the media covers a particular policy issue the more likely citizens are to know relevant facts about that issue (Jerit, Barabas, and Bolsen 2006; Jerit and Barabas 2006; Barabas and Jerit 2009). Meanwhile, a series of survey experiments have shown that policy information can affect preferences under particular circumstances (Gilens 2001; Bullock 2011).

But policy information is no panacea. A subset of this research has found that citizens are as biased and stubborn as we thought. For instance, Gaines et al. (2007) find that while members of different partisan groups learned many of the same facts about the War in Iraq, their interpretations of those facts differed according to their partisanship. Republicans saw mounting casualties as “small” while Democrats saw things very differently. In one experiment on welfare attitudes, Kuklinski et al. (2000) found that correcting misperceptions about welfare policy had no effect on respondents’ policy preferences. Even more perversely, Nyhan and Reifler (2010) found that attempting to correct inaccurate beliefs about Bush-era policies like Iraq, tax cuts, or stem cell research caused some partisans to double-down on inaccurate priors.

However, the literature is far from conclusive on this subject. Bullock’s (2011) experimental study of attitudes toward health policy provides the most direct test of policy information versus partisanship. He found that information about the direction of policy change—a move in a more liberal or conservative direction—had a larger substantive effect on respondents’ beliefs than did cues about party positions on the issue. According to Bullock (2011), the results show that when partisans are “exposed to more than meager descriptions of policy . . . they can arrive at policy views that are independent of and even contrary to the views of their party leaders” (512).

In sum, recent literature on information effects defies a simple verdict on citizen competence. What citizens learn about politics and policy issues has much to do with the information environment, how the information is presented, and citizens’ predispositions. At the very least, these alternate takes have encouraged researchers to push beyond a focus on levels of political knowledge (what citizens have learned prior to being surveyed) to an examination of the dynamics of political learning (when is information likely to affect preferences and behavior).

Political Science and Informational Nudges

By asking the question in this way, political behavior scholars have actually been designing and studying informational “nudges,” placing them in the middle of a broader, cross-disciplinary conversation about how targeted information can help individuals make informed choices (Thaler and Sunstein 2009). Like political scientists, social policy scholars have recognized that individuals rarely make decisions like ideal rational actors, but suffer from limited cognitive capacity and built-in biases. In a litany that echoes the canon on political ignorance, Kahneman and Kreuger (2006, 3) explain:

a large literature in behavioral economics and psychology finds that people often make inconsistent choices, fail to learn from experience, exhibit reluctance to trade, . . . and depart from the model of the rational economic agent in other ways. If people display bounded rationality when it comes to maximizing utility, then their choices do not necessarily reflect their “true” preferences.

In response, scholars have suggested that informational nudges can help citizens navigate complex decisions about health and well-being. Nudges are attractive because they are inexpensive to administer and balance the desire to both empower citizens to choose and minimize the chances that they choose poorly (Thaler and Sunstein 2009). By correcting information asymmetries, the argument goes, these interventions can potentially invigorate all manner of competitive markets.

Examples abound. Because they do not recognize the benefits of compounding interest, employees may not opt to contribute to a generous retirement plan. Sending them a brief flyer that explains the benefits can encourage them to do so (Clark, Maki, and Morrill 2013). Because choosing a prescription drug plan under Medicare Part D can be confusing, most seniors opt into a default plan. Providing personalized information about plan costs encourages participants to switch to one that saves them money (Kling et al. 2012). And because gifted, low-income students may not have a sense of the types of colleges and financial aid they may qualify for, many attend institutions that are less selective. Mailing these students a personalized college and financial aid guide boosts application rates (Hoxby and Turner 2012). Survey and field experiments in political science fit into this tradition, they just rarely recognize it (Gerber, Green and Larimer, 2008).

But for all of the enthusiasm around informational nudges in social policy, their effectiveness depends on many of the same obstacles that hinder citizens in politics: limited cognitive capacity, selective perception, motivated skepticism, and so on. Sometimes informational interventions work, sometimes they do not. For instance, political scientists have found that information about the performance of local public schools can affect citizen perceptions in some settings (Chingos, Henderson, and West 2012) but not others (Buckley and Schneider 2007; Howell 2006). And even when interventions provide information to citizens directly, sometimes the information alone is not enough to change behavior, and other times it is. When a group of economists randomly assigned families to receive an estimate of their teenager’s financial aid package from a tax professional, the information alone had no effect (Bettinger, Long, Oreopolous and Sanbonmatsu, 2012). But a fifteen-minute video about college opportunities had lasting effects on high school students’ aspirations (Oreopolous and Dunn 2013).

Toward a Better Understanding of Information Effects

We are left, then, with a set of divergent findings in both political science and behavioral economics. Oftentimes citizens know little about politics and policy, but sometimes they know more than we would expect. Sometimes they learn when provided with an opportunity to do so, and sometimes they do not. How should we make sense of this body of work? Both types of studies have tended to accumulate, citing one another and noting how their findings differ but only occasionally exploring the roots of why they differ (see earlier book-length treatments in Lupia and McCubbins 1998; Sniderman, Brody, and Tetlock 1991).

This divergence is due, at least in part, to the fact that scholars in search of information effects have employed very different assumptions about where to look for political learning, how

best to elicit it, and why we might see it in some settings and not others. Assumptions vary at each stage of the learning process, from the factors that influence whether citizens receive political messages to the types of messages that are likely to affect attitudes to the prior beliefs on which those messages are supposed to act. Even among those with a more dynamic view of political learning, the question of whether information has an effect on beliefs has received more attention than equally important questions about why information would cause citizens to update in some contexts but not others. The implicit assumption is that in a rational world, all types of information should lead all citizens to update their beliefs.

But we need to be more specific about how variation in the components of the learning process is likely to affect these predictions. In an effort to do so, this section lays out a simple, stylized model of rational learning that guides the chapters that follow. I focus on three components of the learning process: the probability that citizens receive particular pieces of information, the relevance of the messages themselves, and the prior beliefs that those messages are supposed to affect.

To derive some expectations, I employ the basic intuitions from Bayesian models of updating but include a prior step: whether or not citizens actually receive information in the first place. Conditional on receiving the message, citizens must then combine that new information with their prior beliefs to arrive at an updated posterior belief. Bayesian updating models describe the arithmetic process by which this combination takes place and serve as a useful framework for assessing political learning (for the most comprehensive and useful review of Bayesian updating in politics, see Bullock 2009). The message and the prior beliefs are each characterized by their location and their precision. Location refers to an individual's best estimate of where a political object falls on a given dimension, such as a policy's outcome, a candidate's trustworthiness, or a school's quality. The precision is the inverse of the variance of that estimate. More simply, the precision is a subjective measure of the uncertainty surrounding the estimated location of the prior or the message. The posterior belief is the weighted average of the two locations, each weighted by their precision.

Note that there are two important quantities that influence the likelihood and magnitude of attitudinal change. First is the difference between the location of the prior and the location of the new information. Holding precision constant, the greater the difference, the larger the expected revision in beliefs. Intuitively, messages that run counter to prior beliefs are likely to surprise citizens. Second, and perhaps more importantly, the ratio of the two precisions determines the strength of new information compared to the prior belief (Bullock 2009). The more precise the prior relative to the new information, the smaller the attitude change. If the precision of the new message rivals the precision of the prior belief—perhaps because the prior is under-informed or the message is particularly informative—it will have a strong influence on the resulting attitude.

Political scientists have obviously studied each of these components at length, but research on information effects has not always been sufficiently precise in laying out its expectations on each.

The Probability that Citizens Receive Information

Before citizens can learn from new information, they have to receive it, and the probability of reception varies across individuals and contexts (Zaller 1992). The traditional approach to measuring political knowledge has focused on whether citizens can answer a battery of factual items about government, politics, and policy (Delli Carpini and Keeter 1996). Because these facts are public knowledge, the implicit assumption is that each fact is equally available for

public consumption and equally relevant to the tasks of citizenship. So even though senators are elected in staggered terms, citizens' ability to identify and place one who is up for reelection is considered as important as their ability to do so for the other. In this literature, variation in the probability that citizens receive and retain new information is typically cast as a function of individual-level attributes like education, interest, and prior levels of political sophistication (Price and Zaller 1993). When citizens are not able to answer these factual questions correctly at a given point in time, the assumption is that they either never received it or have forgotten it (Zaller 1992).

This is not a particularly dynamic view of political knowledge. In response, as I discussed above, scholars have argued that the political environment also shapes the probability of receiving important messages and cues. These alternate takes tend to focus on information about current policies, candidates, and events that may be relevant to a particular choice or debate in a particular period (Bartels 1993; Gilens 2001; Jerit, Barabas, and Bolsen 2006; Gaines et al. 2007; Barabas and Jerit 2009). This approach shifts the focus to the supply of information and the institutions and structures that influence that supply—elections, debates, mass media, and so on.

This alternative take is a welcome one, but it also tends to be overly optimistic and underspecified with regard to the reliability of simple cues. In particular, less attention has been paid to whether citizens are able to recognize the conditions under which these heuristics are likely to be more and less dependable. If the information environment raises the probability of receiving cues that are reliable for some tasks but not others, their availability could actually raise the likelihood that citizens make mistakes (what Tversky and Kahneman (1974) call the availability heuristic). And if this ability to selectively employ cues varies with prior levels of education and sophistication, environmental encouragement to use them may actually reinforce gaps in competence. Such a pattern may even help to explain where those gaps come from in the first place.

Note too that the probability of receiving information has implications for experimental studies. On some issues, public opinion research necessarily lags behind the flow of information. Therefore, in issue areas that are subject to heavy media coverage and high-profile political debates, it is entirely possible that citizens have already incorporated new pieces of information into their judgments. In these cases, the position of the prior belief will already reflect that information, minimizing the likelihood of attitude change.

Chapter two examines the influence of the information environment on the probability that citizens have accurate beliefs about the party identification and issue positions of their senators. I use the context of the 2006 Senate election to document both the benefits and the drawbacks of environments that increase the salience of party cues. Voters are more likely to recall the partisanship of their senator when that senator is up for reelection. The increased availability of party cues does seem to help voters make inferences about the positions of their senator, but only under certain conditions. When those conditions are not met, the availability of party cues can backfire.

The Precision of the Message

Political scientists have yet to identify agreed-upon measures of the precision of messages or prior beliefs. According to Bullock (2009), on these questions the “previous research offers almost no guidance” (1114). This is especially true of the recent research on policy-specific information, where scholars have focused on a diverse array of policy-relevant facts in their efforts to study the learning process. Messages have included: descriptive statistics about the race and socioeconomic status of welfare recipients (Kuklinksi et al. 2000); data about

the crime rate (Gilens 2001); the percentage of the budget spent on various programs (Gilens 2001; Kuklinski et al. 2000); data on teacher salaries and per-pupil funding (Howell and West 2009); predictive appeals about likely policy outcomes (Jerit 2009); the search for WMD in Iraq (Nyhan and Reifler 2010); a hypothetical proposal to change a state-run healthcare program (Bullock 2011); and others.

Testing different types of information is a worthwhile step. But not all types of information are likely to be seen as equally precise. In the absence of accepted measures of precision, a reasonable substitute (though admittedly subjective) is the concept of “diagnostic value” introduced by Kuklinski et al. (2001). The authors posit that information is high in “diagnostic value” when it “clearly and fully conveys the central considerations relevant to a decision or judgment task. The diagnostic value of information is especially high when it denotes both what factors to think about and how to think about them” (Kuklinski et al. 2001, 412). If the question is whether to support a political candidate, their party identification is highly diagnostic, their age perhaps less so. When it comes to attitudes about particular public policies, information about how much they cost or whether they work should be important to most taxpayers. And in the context of market-based social policies, comparative information about the quality and cost of various service providers should have higher diagnostic value than other attributes.

Diagnostic value may not be constant across citizens who often have different values and priorities, and it is also likely to vary across settings. Trends and events may make particular dimensions more salient at some points than others (economic growth versus national security versus budget deficits). But this does not preclude us from being precise about expectations.

In an effort to do so, chapters three and four propose that information about the performance of public programs and institutions—whether they are accomplishing public goals—will have sufficient diagnostic value to change beliefs. In particular, chapter three provides an opportunity to explore variation on two dimensions: the valence of performance information (positive or negative) and the measure of performance (effectiveness vs. efficiency).

The Location and Strength of Prior Beliefs

The distance between respondents’ prior beliefs and the message they receive will help to determine the extent of the attitude change. Existing research has shown that citizens hold serious misperceptions when it comes to basic facts about public policies and the quality of different public service options. Holding all else constant, the wider the gap between the prior and the message, the larger the expected revision of beliefs. In a sense, this is a measure of how surprising or new the information is to the decision-maker.

But responsiveness to new information mainly depends on the strength of the prior. In general, researchers have tended to use proxy variables like party identification and ideology as indicators of prior orientations. This approach has informed discussions of political learning, providing plenty of evidence that prior beliefs are often so strong as to short-circuit the updating process (Bartels 2002). In some cases, partisans appear to actively reject and counter-argue new information that does not conform to their predispositions (Redlawsk 2002).

Here again, however, we would expect information effects to vary across different issues and decisions, not only according to individual-level predispositions. At the risk of oversimplifying, it seems reasonable to suggest that prior beliefs will be strongest on those issues that are most polarized, salient, and regularly in the media. On these issues, repeated messages would be continuously incorporated into beliefs, eventually reducing the level of uncertainty

surrounding the prior.¹ In contrast, on issues that are less regularly drawn into partisan bickering and are only occasionally part of national debates, we would expect prior beliefs to be somewhat weaker. Similarly, consumers in markets characterized by low levels of information and goods that are difficult to evaluate from the outside (e.g. schooling, healthcare, and residential choice) should have less precise preferences that are more easily affected by new data.

The literature on policy information has looked at policies and programs with dramatically different levels of prominence, from the Iraq War and immigration to prison spending and foreign aid. Yet the implicit assumption often seems to be that policy-relevant information will be equally effective in each setting. Chapters three and four explore information effects in areas where prior beliefs are likely to vary from moderately strong (Food Stamps) to weaker (Head Start) to largely uninformed (the quality of local colleges). While some might suggest that these issues and choices are less relevant to politics than more polarized ones, in the conclusion I argue that a focus on only the most polarized issues would miss critical examples of attitude change. After all, how do issues become polarized in the mass public in the first place? Individuals change their minds about policy as they receive information from party elites (Carsey and Layman 2006; Levendusky 2009). Excluding issues that are still on the periphery of partisan debates could lead us to ignore important stages in that evolution.

The Way Forward

The chapters ahead examine the different facets of this rational learning model. Collectively, they provide insight into the circumstances under which citizens can learn relevant policy information, the effect that information can have on beliefs and preferences, and the limits of political learning. Specifically, I tackle the following questions:

- Returning to Bullock’s original question: can providing the relevant “facts” help citizens hold more accurate beliefs and “make better decisions?” Though a growing number of studies have shown that policy information can lead citizens to update their beliefs, much of the existing political behavior research is less sanguine. I use three different data sources to explore information effects across three different venues.
- What political contexts enable citizens to learn relevant information? Research suggests that political learning is, at least in part, a function of whether relevant information is available in the environment. At the same time, studies of campaign effects have produced mixed results. Do political events expose voters to more information about particular candidates and issues? And does that information translate to more accurate beliefs?
- Do information effects vary across different types of policy information? Scholars have defined “policy-specific information” quite broadly, covering everything from policy design to clientele served to how much a program costs. I focus on the potency of performance information—whether government programs and institutions are accomplishing public goals—and explore how effects vary across negative and positive messages.
- How do information effects vary across different groups of citizens? And do they tend to narrow gaps in beliefs across groups? Scholars have vigorously debated whether new information should lead the beliefs of opposing partisans to converge. Likewise, critics of market-based policies have worried that the most advantaged will be the best equipped to

¹ Most analyses of Bayesian models hold the precision of the prior belief constant to simplify the analysis.

learn about their options. In each chapter and in the conclusion, I examine information effects across different groups.

The next chapter asks whether elections inform voters. Specifically, I ask whether citizens in the midst of an election campaign know more about their elected representatives' party identification and issue positions than those who are not. Using the 2006 Cooperative Congressional Election Study (CCES) and both between- and within-subjects analyses, I find that the electoral context does make citizens more likely to recall the party identification of their senator. But this improvement only translates to more accurate perceptions of roll call voting behavior on issues where the senator voted the party line. On issues where the senator defected, electoral context does not help, and may very well detract from, citizen accuracy.

Chapter three shifts the lens from electoral politics to citizen attitudes toward government programs. The literature on policy information has rarely explored the relationship between perceived policy outcomes and public support for that policy. Drawing on data from two survey experiments on the 2010 CCES, chapter three focuses on how this type of policy-specific information shapes support for two social programs: Food Stamps and Head Start. Because the Food Stamp program has been evaluated favorably and Head Start was found to be ineffective, the data allow for an assessment of how individuals with different predispositions respond to different types of policy information. I find that citizens are quite responsive to policy performance information. More importantly, I find no evidence that citizens of different ideological leanings responded differently to information on policy performance.

Chapter four takes us to the final stop on the search for information effects: the market for local public goods. Proponents of market-based policies suggest that informed choice and competition can improve both individual outcomes and the overall quality of public services. Critics often allege that market-based policies will further advantage the most sophisticated citizens because they are best equipped to capitalize on those options. Political scientists have studied the micro-foundations of these market models—namely, whether citizens are capable of making such informed choices (Lowery and Lyons 1989; Teske et al. 1993). In the process, researchers have suggested that decisions to enter and exit jurisdictions is itself form of political participation.

In chapter four, I extend this line of scholarship to the market for local public colleges. Using data from a survey of parents, I explore whether additional information about college characteristics has an effect on citizens' preferences over two colleges. The survey asked parents to choose between a pair of local public colleges based on a set of facts. A subset of respondents received one additional piece of information: the graduation rate at each college. The results show that consumers are quite responsive to information about performance, and that the effects are most prominent among less educated parents.

Lessons Learned

Overall, the findings in chapters two, three, and four suggest that the information environment can affect citizens' level of knowledge and policy preferences in predictable ways. In particular, I find considerable evidence of updating and little evidence of the partisan biases that loom so large in the literature. But it is also true that differences in opinion across partisans remained stable, and that citizens appear to be more responsive to negative information than positive. And while elections raise the salience of party cues, this is not always beneficial.

Chapter five concludes by attempting to reconcile these findings with divides in the existing literature. To do so I return to the rational learning model and argue that divergent

findings in the literature are due, in part, to differences across studies on key dimensions of the model. These differences are only occasionally acknowledged, and variation is rarely embraced as an opportunity to study heterogeneous effects. By sketching out the conditions under which certain patterns hold and where they do not, this study should help to move this divergent literature forward.

Chapter 2

Do Elections Inform Voters? Evidence from the 2006 Senate Elections

In the rational learning model laid out in the introduction, the first stage is the receipt of political messages. In the real world, the probability that citizens receive new information is a function of individual-level characteristics and conditions in the political environment. Traditionally, studies of political knowledge have focused on the importance of individual-level characteristics in explaining whether citizens receive political messages. An alternative body of research has explored the important role that the information environment plays in creating opportunities for political learning. In particular, scholars have argued that the political environment provides citizens with cues that can help them arrive at decisions that reflect their preferences. Here, I test an important environmental influence on the probability of receiving political information: the electoral context. Campaigns and elections provide a near constant flow of messages about candidates and issues. Do those messages reach voters? In other words, is the probability of receiving important political information higher in an electoral context than in a non-electoral one?

These questions are of critical importance to representation and political accountability. Ideal models of representation posit that constituents hold incumbents accountable for the positions they take on important issues. Voters monitor the behavior of their member of congress, compare that behavior to their own policy preferences, and punish or reward the incumbent on election day. In theory, the roll call voting records of incumbents provide an opportunity to engage in this type of electoral accountability. Because incumbents know that citizens are watching, the threat of electoral punishment provides incentive to vote in a way that reflects the preferences of their constituents.

This ideal process encompasses three distinct steps. First, constituents must have strong and stable policy preferences of their own. Second, they must have some knowledge as to how their legislator voted on a particular policy issue that allows them to assess the level of policy agreement. Finally, they must incorporate that assessment into their voting decisions.

Unfortunately, political science research on voter knowledge has provided reason to doubt at each juncture. Citizens are characterized as having “nonattitudes” on important policy issues (Converse 1964; Zaller 1992). They also have low levels of information about politics in general and are particularly poorly informed when it comes to the issue positions of their members of congress (Miller and Stokes 1963; Delli Carpini and Keeter 1996). Instead of policy agreement, voting decisions are more often based on party identification, valence issues, and affective considerations. Even when citizens come into contact with useful policy information, the literature on party identification suggests that it serves as a “perceptual screen” that leads citizens to engage in selective perception and biased updating (Campbell et al. 1960; Rahn 1993; Bartels 2002).

Despite these gloomy findings, recent research has produced evidence that voters do hold legislators accountable for their voting records (Ansolabehere, Snyder, and Stewart 2001; Canes-Wrone, Brady, and Cogan 2002). Aggregate-level studies have found that members of congress who are ideologically “out of step” with their constituents are punished electorally. Improved measurement of citizen perceptions of legislators’ issue positions has also produced evidence of a more informed public. In particular, surveys that ask citizens about actual roll-call votes have found that voters often do have more accurate perceptions than previously thought, and that

elections provide an opportunity to punish incumbents whose views are more extreme than their constituents (Ansolabehere and Jones 2010; Bafumi and Herron 2010).

As I argued in the introduction, these findings present something of a puzzle. If the typical voter is poorly informed about the roll call voting records of her incumbent, how do macro patterns of accountability exist? Some have argued that an army of omniscient citizens is unnecessary: a subset of informed voters can successfully hold members accountable and the errors of uninformed voters will cancel each other out (Erikson, MacKuen, and Stimson 2002). Others have pointed to the importance of cue taking and heuristics as a shortcut to obtaining information about issue positions. By strategically using cues from elites, interest groups, and the broader political environment, under-informed voters can infer where candidates stand on particular issues, approximating the decision-making of their more informed peers (Conover and Feldman 1989; Lupia 1994; Lau and Redlawsk 2007).

A more recent strand of literature has suggested that citizens actually do learn about important policy debates and the issue positions of the parties from media coverage and election campaigns (Lenz 2009; Jerit, Barabas, and Bolsen 2006; Barabas and Jerit 2009). Though election campaigns may not change vote choice in the short-term, they can help clarify party positions and provide information about the stances taken by individual incumbents. Analyses of statewide elections for senate and governor have found that campaigns do tend to inform voters and clarify issue positions (Franklin 1991; Kahn and Kenney 1997; Partin 2001).

This chapter builds on this literature by exploring the contextual forces that shape how much voters know about the roll call voting records of incumbents. The study exploits the staggered terms of U.S. senators as a natural experiment to examine whether constituents whose senator is up for reelection have more accurate beliefs about that senator's party identification and issue positions. Like recent work on the subject, I use a large, unique dataset that asked constituents about real roll call votes taken in recent sessions of congress (the 2006 Cooperative Congressional Election Study (CCES)). The 2006 CCES asked respondents about their own positions on these votes, as well as how they thought each of their senators voted on that issue. The dataset contains more than 36,000 respondents and included questions on seven roll call votes, providing an excellent opportunity to assess the basic research question laid out above.

I find that the election cycle has a clear, significant effect on whether constituents are able to recall the party identification of their senator. The results suggest that citizens are doing more than just guessing, and that the context of an election seems to improve citizens' ability to recall this crucial piece of information. As to whether knowledge of this party cue leads citizens to have a more accurate picture of their senators' issue positions, the results are more mixed. On the whole, electoral context has a modest positive effect on accuracy, but this main effect obscures interesting heterogeneity across different votes. Consistent with a pure cue-taking story, I find that citizens evaluating senators up for reelection are generally more accurate than those evaluating senators who are not up when it comes to votes where the senators toed the party line. However, any informational advantage disappears on votes where the senator defected from the rest of his or her party. Indeed, there is evidence that on these defecting votes, citizens are *less* likely to be correct in the context of an election campaign.

In short, while elections do seem to improve the accuracy of citizen perceptions by clarifying party labels, this information only translates to improved accuracy when party cues predicted the senator's position. These findings have implications for the study of political knowledge and heuristics. In particular, they suggest that voters may not recognize when

heuristics are likely to be useful and when they are not, a point that has been underexplored in the political science literature.

Voters, Issue Information, and Elections

In their classic study of constituency influence in congress, Miller and Stokes (1963) set the tone for decades of research on representation and political knowledge. Their analysis of the 1958 National Election Study (NES) found that the vast majority of voters were unable to correctly identify the positions of their members of congress on pressing issues like social spending, desegregation, and foreign affairs. Miller and Stokes (1963) famously concluded “Far from looking over the shoulder of their Congressmen at the legislative game, most Americans are almost totally uninformed about legislative issues in Washington” (47). More recent work has reiterated this claim, suggesting that when it comes to the level of political information “the mean is low but the variance is very high” (Converse 2000). Studies that examine whether citizens can place their representatives on a particular roll call vote suggest that beliefs are plagued by selective perception and projection (Alvarez and Gronke 1996; Wilson and Gronke 2000). Voters assume that legislators who share their partisanship vote in ways they agree with and ignore information to the contrary. These studies paint a discouraging picture of the micro-foundations of representation.

Others have tempered these conclusions somewhat. For one thing, authors have suggested that Miller and Stokes’ conclusions are based on faulty measurement of constituent opinions and perceptions of member behavior (Achen 1978; Erikson 1978). More recent approaches have presented citizens with actual roll call issues considered in congress, asking them how they themselves would have voted and how they thought their representative voted on it. These types of questions allow for a much more direct examination of policy representation and a clearer assessment citizen accuracy. Using this kind of design, Ansolabehere and Jones (2010) conclude that when presented with actual roll call items, the vast majority of citizens have clear preferences on those issues, hold beliefs about their representatives’ position that are accurate more often than not, and base their approval ratings on these perceived levels of policy agreement.

Research on heuristics, cue taking, and other “cognitive shortcuts” has also lead to a reassessment of whether voters can learn what they need to know. Studies of heuristics argue that voters are “cognitive misers:” rather than exhaustively searching the information environment for every last piece of political information, they “satisfice,” relying on cues that help them to make inferences (Simon 1997). The logic here is that voters need not be encyclopedias of political knowledge, but can use these clues as a shortcut to arrive at a rational decision (Popkin 1991; Lupia 1994; Lupia and McCubbins 1998). Party identification is the most commonly cited cue, and political scientists have shown how voters use it to infer issue positions and make distinctions between candidates. As Conover and Feldman (1989) argue, knowing that a particular candidate is a Democrat or Republican enables voters to place that candidate into a category. They can then attribute the issue positions of that category to the candidate himself.

This mode of reasoning can create problems when candidates’ issue positions do not adhere to the party line. Indeed, Ansolabehere and Jones (2010) find that voters are much more accurate in placing their representatives on votes that followed the party line (about 75 percent correct) than on those where the incumbent voted with the other party (40 to 50 percent correct). However, it is also true that in an increasingly polarized and sorted world, party-based inferences of issue positions are going to be correct more often than not (Levendusky 2009).

But how do voters learn these party cues? Research on media and campaign effects has shown that citizens are capable of learning about candidates and policy issues. In a series of studies, Jerit and Barabas document how the volume and character of media coverage affect voters' knowledge of policy issues (Jerit, Barabas, and Bolsen 2006; Jerit and Barabas 2006; Barabas and Jerit 2009). When it comes to campaigns, political scientists have generally agreed that they rarely change voters' minds, but that they do have a positive effect on how much voters know about candidates (Conover and Feldman 1989; Franklin 1991; Bartels 1993; Ansolabehere and Iyengar 1995). For instance, Conover and Feldman (1989) show that citizens are better able to place the political parties and candidates on the issues as the presidential campaign progresses, and that party cues are a consistent predictor of these placements. In his analysis of the 1980 presidential election, Bartels (1993) found that, contrary to the "minimal effects" thesis, media coverage of the two candidates had large effects on citizen perceptions and levels of information. Similarly, Lenz (2009) argues that voters often learn where the parties stand on particular issues when those issues arise in an electoral campaign.

A subset of this work has looked specifically at the impact of statewide election campaigns on voter information, finding that the electoral setting and campaign intensity affect voter knowledge. Partin's (2001) study of gubernatorial elections finds that campaign intensity increases voters' knowledge of non-incumbents, and that campaign effects are concentrated among the least knowledgeable voters. Similarly, Kahn and Kenney (1997) find that the more intense the campaign, the more likely voters were to use issue positions to evaluate the candidates. Finally, like this study, Franklin (1991) exploits the staggered terms of the U.S. Senate to explore the effect of particular campaign strategies on voters' views of candidate ideology. Using a within-subjects design, Franklin found that voters perceived the ideological position of candidates more accurately when those candidates stressed issue positions during the campaign.

The notion that campaigns can inform citizens about their incumbents' issue positions is in-line with interview-based studies of Congress (Fenno 1978; Kingdon 1989). In these treatments, electoral challengers are the key players: they will draw attention to roll call votes that may be hard for the incumbent to justify, thereby injecting issue information into campaign coverage. For instance, in Kingdon's (1989) classic study of congressmen's voting decisions, when asked whether his constituents would recognize a vote that he had just taken, the congressman replied "No. I know that nobody will notice it right now. People never do. But it may be used against you in the next campaign. I learned that lesson in my first campaign for reelection" (60). Fenno (1978) makes a similar argument, suggesting that it is "the challenger's task to probe for the votes that are most difficult for the incumbent to explain. And that the member cannot be sure before the actual campaign which ones they might be" (143).

In sum, the existing literature suggests that while voter information about roll call voting records is typically low, the political environment often provides opportunities to learn about them. In particular, election campaigns provide party cues, while electoral challengers often highlight particular roll call votes in their efforts to criticize the incumbent. These campaign-driven information flows should raise the probability that voters are aware of party cues and issue positions in the context of an election. Further, presenting survey respondents with actual roll call votes suggests that constituents may be better able assess the positions of their incumbents than was previously thought.

But important questions remain. Though campaigns may help voters to have a more accurate sense of where their incumbent falls on the ideological spectrum, do they actually

provide voters with more information about their positions on individual issues? And, if so, what is the mechanism by which voters become more accurate about these issue positions? Are differences in information levels a function of the accessibility of party cues in an electoral environment? Or do elections provide citizens with deeper knowledge of their incumbents' issue positions that goes beyond a simple party-based inference? This chapter takes a closer look at how citizen accuracy varies across different electoral settings and different types of votes.

Hypotheses

Existing research suggests that election campaigns communicate important political messages, both directly in the form of advertisements and indirectly through media coverage. In particular, elections highlight party cues, helping citizens to categorize candidates and make inferences about their positions (Conover and Feldman 1989). At the very least, then, we should expect citizens to be more likely to recall the party identification of senators who are up for reelection. For the marginal voter who does not follow politics, this information is likely to be more accessible in an electoral context than otherwise.

Hypothesis 1: Citizens will be more likely to recall the party identification of senators who are up for reelection than those who are not.

Existing research also suggests that campaigns are capable of passing on more detailed information about the issue positions of candidates. Franklin (1991) finds that issue-laden campaigns lead to more accurate voter perceptions of where candidates fall on the ideological spectrum. Kahn and Kenney (1997) find some evidence that more intense campaigns communicate more issue information, leading voters to weigh policy concerns more heavily in their evaluations. Lenz's (2009) look at citizen learning also finds that election cycles can successfully place issues on voters' radar screens and effectively communicate the positions of the parties on those issues.

Given what we know about the usefulness of party cues, we would expect that if hypothesis 1 is borne out, knowledge of party identification should translate to improved accuracy when it comes to inferring the roll call positions. In an increasingly polarized legislature, party cues can help voters to correctly place their legislators more often than not. In this case, knowing the partisanship of the senator is the *mechanism* by which elections improve citizen accuracy on roll call votes.

Hypothesis 2: Citizens will be more likely to have accurate beliefs about their senators' issue positions when those senators are up for reelection.

Hypotheses one and two suggest a third hypothesis: if the context of an election campaign leads citizens to have a more accurate view of their senator's issue positions, is this improvement solely a function of the accessibility of party cues? It could also be the case that campaigns teach citizens factual information about the issue positions of individual candidates. In other words, it is possible that the electoral environment has an independent effect on knowledge of candidate positions above and beyond party cues. As Kingdon's (1989) and Fenno's (1978) interviewees pointed out, roll call votes can become prominent during the campaign, which could inform voters about incumbents' issue positions.

Luckily, the roll call votes analyzed here provide an opportunity to test these two explanations. Because many of the roll call issues on the 2006 CCES were not straight party-line votes, the dataset includes a sizable number of instances where senators defected to vote with the opposite party. If voters in the midst of a senate campaign have information about issue positions that goes beyond a cue-based inference, we would expect them to also be more accurate on defecting votes when their senator is up for reelection.

Hypothesis 3: Electoral context has an independent effect on voter information about roll call votes that goes beyond partisan cue taking.

Of course, it is entirely possible that voters have specific information about party-line votes and are not simply making cue-based inferences. What might look like cue-based improvements in accuracy could actually be the result of more “systematic processing,” where voters learn about specific roll call positions (Chaiken 1980). Unfortunately, it is not possible to tease out these two alternative explanations in the current dataset; improved accuracy on party line votes is a case of observational equivalence. However, assuming that citizen accuracy on party line votes is entirely cue-driven provides a conservative test of the effects of the electoral context. Further research should endeavor to distinguish inferences from real knowledge on party line votes.

Data and Methods

The 2006 CCES Common Content provides an excellent opportunity to examine these questions (Ansolabehere 2007). The survey presented respondents with a detailed description of seven roll call votes that had taken place in a previous session of congress. The survey then asked all respondents how they would have voted on each issue and how they thought each of their senators voted. The roll call issues on the survey were: the partial birth abortion ban; federal funding of stem cell research; an amendment calling for the mandatory withdrawal of troops from Iraq; comprehensive immigration reform; raising the minimum wage; extending the capital gains tax cut; and the Central American Free Trade Agreement (CAFTA). The roll call questions provided respondents with a description of the issue that borrowed from actual legislative debate and coverage of the proposals, providing a high degree of external validity. Respondents were also asked whether they recalled the party identification of each of their senators.

This set of votes is particularly useful because a handful of them divided one of the parties in the Senate, meaning a nontrivial number of members defected to vote with the other party. Federal stem cell research split Republicans, with 19 Republicans joining Democrats in voting for the measure. The ban on partial birth abortion split Democrats, with 12 Democrats breaking ranks to vote for the ban. The free trade bill split both parties, with 11 Democrats joining Republicans in voting for the bill and 11 Republicans joining the Democrats in voting against it. I leverage these defections from the party line to test Hypothesis 3. I define a “defection” as any vote that does not match the majority of the senator’s party.

The 2006 election cycle featured 33 Senate elections, 4 of which featured retiring incumbents (open seats) and 29 of which featured incumbents running for reelection. Because I am interested in the effect of the electoral context on knowledge about the incumbent, I analyze the 29 races featuring an incumbent as the “electoral contexts.” In the analysis below, I exclude incumbents from the open-seat races and limit comparisons to those who were up for reelection and those who were not. The “election” variable is the predictor of interest, coded one for a senator who is up for reelection and zero otherwise.

I analyze the data in three different ways. First, I pool across senators, including each respondent’s evaluation of each senator as an observation (this dataset “stacks” senators, such that each respondent appears twice in the dataset). Second, I pool the data across roll call issues, such that the observation is the respondent-Senator-issue trio. These are “between-subjects” comparisons because they compare those evaluating a senator up for reelection to those evaluating one who is not. Finally, I confirm the patterns found in the first section by conducting “within-subjects” comparisons. In every case, standard errors are adjusted to reflect the pooling

of respondents across senators and issues, and the appendix contains some robustness checks for the hypothesis tests reported here.

In deciding which outcome variable to use, decisions have to be made about how to characterize citizen accuracy. The recent paper by Ansolabehere and Jones (2010) limits the analysis of voter accuracy to only those who expressed a belief about how their representative voted, thereby eliminating any respondents who said that they were unsure about the vote. As you will see below, this decision tends to inflate the picture of citizen accuracy because it selects on the dependent variable. Only those respondents who were confident enough to provide an answer are included in their analysis, and this subset is likely to look quite different from the universe of respondents.

Recent research has suggested that when respondents say that they “don’t know,” more often than not they legitimately lack the information (Luskin and Bullock 2011). This is particularly important for the current study. If the goal is to assess whether citizens are more likely to have an accurate perception in an electoral context, then neither an incorrect answer nor reporting that you “don’t know” constitute accurate perceptions. I therefore include the “unsures” in my analyses and examine variation in the proportion of respondents who provided an accurate answer.

The analyses below use differences in proportion and linear probability models to examine the effect of electoral context. The linear probability model includes controls for respondent-level covariates and senator characteristics. The covariates include the respondent’s level of education, a dummy variable indicating whether they were a member of the Senator’s party, and an indicator for the respondent’s general level of political knowledge. The political knowledge variable is the sum of two questions about the respondent’s member of congress: were they familiar with that representative’s name and could they correctly recall his or her partisanship. The regression models do not control for attitudes toward the senator (e.g., approval rating or ideological placement) because these variables are likely “post-treatment” in that they are affected by the electoral context as well. Including post-treatment variables in a regression would bias the estimate of electoral context effects.²

The regression models also control for senator characteristics. I include measures of the senator’s ideology (measured by their DW-Nominate score for the 2005-2006 session of Congress, running from -1 to 1), their ideological extremity (the absolute value of their Nominate score, running from 0 to 1), their party unity score from the previous session of congress (the proportion of times the Senator voted with the majority of her party, scaled 0 to 1), and the number of years they have served in the Senate (the number of years in office).

I also endeavor to sort out whether increases in accuracy are the effect of lucky guessing (Luskin and Bullock 2011). To do so, I examine the difference in the percentage of respondents who say they are “unsure” between the electoral and non-electoral conditions and then compare that to the differences in the percentages that were accurate and inaccurate. If respondents engage in random guessing, we would expect the difference in the percent unsure to be split evenly across the incorrect and correct categories.³ If the effect is truly positive, then there

² In addition, the rates of non-response on these senator-specific attitudes is related to the knowledge that individuals have about that Senator. Those who cannot place them on the ideological spectrum are likely less informed than their peers who can. Therefore, using these items as statistical controls tends to omit less informed citizens. For instance, fully 4,000 respondents answered “not sure” to the Senator approval question. List-wise deletion of these cases would bias the sample in the direction of more informed respondents.

³ Following Luskin and Bullock’s (2011) reasoning: if the electoral context encourages individuals who would not otherwise answer to provide a guess, we would expect the percentage of people getting it right by guessing correctly

should be an overwhelming move from unsure to correct. If the effect is truly negative, then there should be the opposite.

In the analyses where respondents are pooled across issues, it became clear that some of the effects were being disproportionately shaped by the prevalence of one Senator in the dataset. Florida Democrat Bill Nelson, who was up for reelection in 2006, defected from the Democratic line on three of the seven issues. Given the large number of respondents from Florida in the dataset, Nelson's constituents had a disproportionate influence on the estimated effects. More specifically, because Nelson's voting record does not toe the party line, we might expect constituents to be less accurate. This makes for a conservative test of the hypotheses because it will drag the overall estimate of citizen accuracy downward. In order to make sure that evaluations of Nelson are not driving the results, I replicated the analyses without Nelson's constituents in the dataset. Those results are provided in the appendix and cited occasionally in the notes. I also replicated the analyses where I weighted each Senator equally rather than by the state's sample size in the 2006 CCES. The results do not differ appreciably from those reported here.

I begin the results section with an overview of citizen accuracy on the items in question before moving on to discuss differences across electoral context. I then explore the three hypotheses in turn before concluding with a discussion of the findings.

Results

Table 2.1 displays the percentage of respondents that were correct, incorrect, or unsure about each senator's partisanship and their positions on the seven roll calls in the dataset. The results show that people are generally accurate when it comes to the partisanship of their senator. They are correct about 77 percent of the time, and are rarely wrong: just 3 percent of respondents put their senator in the wrong party. However, it is worth pointing out that even on this most basic piece of information, almost 20 percent of respondents report that they are unsure. The low number of incorrect responses suggests that respondents who were less than certain on this question chose not to answer rather than hazard a guess. Having a "don't know" option on these factual questions certainly dampened the number of incorrect answers. In terms of the dependent variable assessed below, the proportion of accurate responses was 77 percent (percentage correct vs. percentage incorrect or unsure).

to be $G = (DK(\text{non-election}) - DK(\text{election})/\# \text{ Response categories})$. Because the number of response categories is two ("for" or "against"), then if the difference in the percentage getting it right between electoral and non-electoral context is greater than G , then the gains in accuracy are not a function of pure guessing.

Table 2.1: Citizen Accuracy on Senator’s Party Identification and Roll Call Positions

	Correct	Incorrect	Unsure
Party ID Recall	77.4	3.4	19.2
95% C.I.	(77.0 to 77.8)	(3.3 to 3.6)	(18.8 to 19.5)
Partial Birth Abortion	50.9	12.8	36.3
	(50.5 to 51.4)	(12.5 to 13.1)	(35.8 to 36.7)
Stem Cells	51.5	11.6	36.9
	(51.1 to 51.9)	(11.4 to 11.9)	(36.5 to 37.3)
Iraq War Resolution	54.2	10.9	34.9
	(53.8 to 54.6)	(10.6 to 11.1)	(34.5 to 35.3)
Immigration	41.5	15.1	43.4
	(41.1 to 41.91)	(14.79 to 15.39)	(43.0 to 43.8)
Minimum Wage	50.9	11.7	37.4
	(50.5 to 51.3)	(11.4 to 11.95)	(37.0 to 37.82)
Capital Gains	50.0	9.3	40.7
	(49.6 to 50.5)	(9.1 to 9.5)	(40.3 to 41.1)
CAFTA	29.6	15.0	55.5
	(29.2 to 29.9)	(14.7 to 15.3)	(55.0 to 55.9)

Table Note: Cells contain the percentage of respondents answering correctly, incorrectly, and unsure on the item in question and the 95% confidence interval for that estimate. The dataset pools respondents across each of their Senators. Calculations reflect sampling weights included in the 2006 CCES.

Despite being able to name their senator’s partisanship at a relatively good clip, citizens are far less accurate when it comes to the specific roll call votes in question. Accuracy ranges from a high of about 54 percent on the Levin Amendment to withdraw troops from Iraq to a low of just under 30 percent on CAFTA. The percentage answering incorrectly ranges from a high of 15 percent (on comprehensive immigration reform and CAFTA) to low of 9 percent on the vote to extend the capital gains tax. In general, citizens are the least likely to be incorrect or unsure on votes about high profile, partisan issues (taxes, the war in Iraq). On issues like immigration and international trade, where party lines are less clear, citizens are more likely to be unsure or incorrect. In a world where party cues help citizens learn about politics, voters have more trouble on issues that are cross-cutting.

It is also important to note that the decision to include the “unsures” in the analysis is extremely consequential. On each and every item, more than 1/3 of respondents are unsure of their senator’s position. On CAFTA and immigration, the percentage unsure is larger than the percentage that answered correctly. Clearly, removing the unsures would lead us to believe that

citizens have much higher levels of knowledge about roll call votes than they actually do. In reality, rather than getting it right about 75 percent of the time (Ansolabehere and Jones 2010), the rate is rarely any better than 50 percent, and often far worse.

Differences Across Contexts: Party Identification

Table 2.1 serves as the backdrop for examining the effect of electoral context. The first question is whether citizens are better able to recall their senator’s party identification when he or she is up for reelection. Given the salience of party labels in an election cycle, this information should be more accessible to those evaluating a senator locked in a reelection campaign.

Table 2.2 provides descriptive evidence that this is indeed the case. The top row displays the percentage of respondents who correctly recalled the partisanship of senators who were not up for reelection, along with the percentage who got it wrong or said that they were unsure (and the associated 95% confidence interval for the estimate). The bottom row reports the analogous percentages for senators who were up for reelection.

Table 2.2: Citizen Accuracy on Party Identification by Electoral Context

	Correct	Incorrect	Unsure	Difference
Non-Election	74.1	3.8	22.1	9.7
	(73.6 to 74.5)	(3.6 to 4.0)	(21.7 to 22.6)	(9.0 to 10.3)
Election	83.7	2.7	13.6	
	(83.1 to 84.3)	(2.4 to 2.9)	(13.2 to 14.3)	

Table Note: Cells contain the percentage of respondents answering correctly, incorrectly, and unsure on the party identification recall item and the 95% confidence interval for that estimate. “Difference” refers to the difference in percent correct across electoral settings. The dataset pools respondents across each of their Senators. Because the data are pooled, the confidence interval on the difference reflects clustering at the level of the individual. Calculations reflect sampling weights in the 2006 CCES.

As is clear from the table, respondents were significantly more likely to recall the partisanship of senators who were up for reelection—on the order of almost 10 percentage points. Fully 84 percent of respondents evaluating a senator up for reelection were able to correctly recall their partisanship, compared to 74 percent of those evaluating senators in an off year. Respondents evaluating senators up for reelection were also somewhat less likely to be incorrect and much less likely to be unsure about the senator’s party identification.

There is little evidence that this gap is due to lucky guessing. The difference in the percentage reporting that they were unsure is about 8.5 percentage points; if these additional respondents were simply guessing at random, we would expect them to be split evenly across the correct and incorrect categories. Given that the gap in the percentage answering correctly is 9.7 percentage points, the pattern is clearly not the result of lucky guessing. These descriptive data suggest that the electoral context generates a real increase in citizens’ ability to recall the party identification of their senator.

In order to validate these differences in proportion and acknowledge potential confounders, Table 2.3 displays the results from a linear probability model of citizen accuracy. In each regression model, the dependent variable is dichotomous with “1” corresponding to an accurate answer and “0” to respondents who were incorrect or unsure. Each row reports a regression coefficient and its standard error. Because individuals appear twice in the stacked

dataset, standard errors are clustered at the level of the individual respondent. The coefficients correspond to the change in the probability of answering accurately for a one-unit change in the independent variable. The intuition is the same as in Tables 2.1 and 2.2, but coefficients are expressed as a change in proportion rather than a change in percentage.

Table 2.3: Regression Model of Citizen Accuracy on Party Identification

	Coefficient	
Election	0.09	***
(s.e.)	(0.003)	
Education	0.04	***
	(0.001)	
Pol. Knowledge	0.47	***
	(0.006)	
Co-Partisan	0.04	***
	(0.004)	
Ideology	-0.02	***
	(0.004)	
Extremity	-0.03	
	(0.018)	
Party Unity	0.10	*
	(0.000)	
Tenure	-0.00	*
	(0.000)	
Constant	0.20	*
	(0.025)	
<hr/>		
R2	0.22	
N	70,821	

Table Note: Cells contain unstandardized coefficients and standard errors from a linear probability model where the dependent variable was scored 1 if the respondent recalled correctly and 0 if the respondent was incorrect or unsure. Standard errors were clustered at the level of the individual respondent. *** = $p < .001$, ** $p < 0.01$, * $p < 0.05$.

The results reveal that even after controlling for respondents' education and level of political information as well as a set of senator characteristics, electoral context translates to about a 9 percentage point increase in the probability of being able to recall a senator's party identification. Not surprisingly, basic political knowledge is also a strong predictor of accuracy, as is the respondent's level of education. Moving from zero on the political knowledge scale (incorrectly identifying the party of your Representative and not recalling their name) to one (getting both questions "right") corresponds to an increase in accuracy of almost 50 percent. Interestingly, a senator's ideology is negatively related to the probability of a correct answer, suggesting that voters had slightly more trouble recalling the party of conservative lawmakers (the Nominat scale runs from -1 to 1, with 1 being the most conservative).

In short, respondents are significantly more likely to know the partisanship of senators who were up for reelection even after controlling for their level of knowledge, their education and Senator characteristics.

Differences Across Contexts: Placing Senators on Roll Call Issues

Does the positive effect of electoral context on party identification recall translate to increased accuracy in placing senators on important roll call issues? At the very least, the literature on party cues would suggest that knowing the partisanship of the senator provides important information as to what their issue positions are likely to be. If a respondent knows the partisanship of their senator and knows where the parties stand on policy issues, then they can make an educated guess as to the position of their senator on those issues. It could also be the case that the campaign serves to inform voters about particular roll call votes.

For now, we leave those competing explanations to the side and focus on whether respondents are more accurate in placing senators who are up for reelection on the seven roll call issues. Table 2.4 displays the descriptive results in the same format as Table 2.2, while Table 2.5 reports the results for roll call-specific regression models.

It is important to note that the regression models do not control for whether the respondent correctly identified the party identification of their senator. While the accessibility of party information may be the *mechanism* by which citizens improve the accuracy of their beliefs, the variable itself is also affected by the election campaign. To include it would introduce post-treatment bias into the estimation. Instead, the third section uses an alternative strategy to examine the relationship between learning about partisanship and learning about issue positions.

Table 2.4 reveals that the effect of electoral context on citizen accuracy varies considerably by issue. In every case, individuals evaluating senators up for reelection are less likely to report being unsure. In some cases, this pattern translates to an increase in the probability of placing the senator correctly. In others, respondents evaluating senators who are up are more likely to be both correct and incorrect, suggesting that some guessing is going on.

The most prominent differences between electoral and non-electoral contexts are on partial birth abortion, stem cells, and comprehensive immigration reform. On these issues, respondents evaluating a senator up for reelection are about four to four-and-a-half percentage points more likely to place them on the issue accurately. For abortion and immigration reform, respondents are about five percentage points less likely to be unsure in the electoral context, and the vast majority of this difference accrues to the correct column. The same is true for stem cell research, though to a lesser extent; more than two-thirds of the difference in the percent unsure winds up in the correct column.

For other issues, these descriptive data provide little evidence that electoral context improves citizen accuracy. On the extension of the capital gains, the effect of electoral context is negative and insignificant. On both the Iraq war resolution and the minimum wage, respondents were not systematically more accurate. On the minimum wage, the gain in accuracy is more than offset by the increase in the percentage who are incorrect on the issue (4 percentage points more likely to be incorrect). And on the Iraq war resolution, the pattern is consistent with pure guessing; the difference in the percent unsure is divided equally between the incorrect and correct column. CAFTA is an interesting case in that the levels of accuracy are quite low to begin with, but electoral context does lead respondents to be more accurate and less likely to be incorrect. The effect is small but statistically significant.

Table 2.4: Citizen Accuracy on Roll Call Issues by Electoral Context

		Correct	Incorrect	Unsure	Difference
Partial Birth Abortion	Non-Election	49.4 (48.9 to 50.0)	12.7 (12.3 to 13.1)	37.9 (37.3 to 38.4)	3.9 (3.0 to 4.8)
	Election	53.4 (52.7 to 54.1)	13 (12.5 to 13.5)	33 (32.9 to 34.3)	
Stem Cell Research	Non-Election	49.9 (49.4 to 50.4)	10.9 (10.6 to 11.2)	39.2 (38.7 to 39.8)	4.6 (3.8 to 5.5)
	Election	54.5 (53.8 to 55.2)	13.1 (12.6 to 13.5)	32.4 (31.7 to 33.1)	
Iraq War Resolution	Non-Election	53.3 (52.8 to 53.8)	10 (9.7 to 10.3)	36.7 (36.2 to 37.2)	2.6 (1.7 to 3.4)
	Election	55.9 (55.2 to 56.6)	12.5 (12.1 to 13)	31.6 (31.0 to 32.3)	
Immigration	Non-Election	40.0 (39.5 to 40.5)	15.1 (14.6 to 15.6)	44.9 (44.4 to 45.4)	4.2 (3.3 to 5.0)
	Election	44.2 (43.5 to 44.9)	15.1 (14.7 to 15.4)	40.6 (39.9 to 41.3)	
Minimum Wage	Non-Election	50.2 (49.7 to 50.7)	10.4 (10.1 to 10.8)	39.3 (38.8 to 39.8)	2.1 (1.2 to 2.9)
	Election	52.3 (51.6 to 53.0)	14.1 (13.6 to 14.6)	33.6 (33.0 to 34.3)	
Capital Gains Tax	Non-Election	50.2 (49.7 to 50.7)	8.3 (8.0 to 8.6)	41.5 (41.0 to 42.0)	-0.5 (-1.3 to 0.4)
	Election	49.7 (49.0 to 50.4)	11.2 (10.7 to 11.6)	39.1 (38.4 to 39.8)	
CAFTA	Non-Election	28.4 (28.0 to 29.0)	15.1 (14.8 to 15.5)	56.4 (55.9 to 56.9)	3.4 (2.6 to 4.2)
	Election	31.8 (30.1 to 32.5)	14.7 (14.2 to 15.2)	53.5 (52.8 to 54.2)	

Table Note: Cells contain the percentage of respondents answering correctly, incorrectly, and unsure on the item in question and the naïve 95% confidence interval for that estimate (not adjusted for heteroskedasticity). Results based on a dataset that pools respondents across Senators. “Difference” refers to the difference in the percentage answering correctly across electoral settings. Confidence interval for estimated difference reflects standard errors clustered at the individual level to reflect pooling. Calculations reflect sampling weights provided by Polimetrix.

The regression results reveal a similar pattern of small but statistically significant differences in accuracy by electoral context. Because of the large sample size, most coefficients are statistically significant, but many are substantively small (and few pass a more conservative

statistical test).⁴ As was the case above, the largest effects appear on the partial birth abortion ban, stem cells, and comprehensive immigration reform. Interestingly, after controlling for respondent and senator characteristics, electoral context appears to have a more pronounced impact on accuracy regarding the Iraq war resolution. The effect of electoral context on the capital gains tax and minimum wage is small, about half the size of the effects for abortion and immigration reform.

Table 2.5: Regression Models of Citizen Accuracy on Roll Call Issues

	Abortion	Stem Cells	Iraq War	Immigration
	Coeff.	Coeff.	Coeff.	Coeff.
Election	0.07***	0.06***	0.06***	0.07***
(s.e.)	(0.004)	(0.004)	(0.004)	(0.004)
Education	0.04***	0.04***	0.03***	0.03***
	(0.002)	(0.002)	(0.002)	(0.002)
Pol.Knowledge	0.29***	0.35***	0.35***	0.26***
	(0.007)	(0.007)	(0.007)	(0.007)
Co-Partisan	-0.07***	-0.03***	-0.05***	-0.07***
	(0.005)	(0.005)	(0.005)	(0.005)
Ideology	0.20***	-0.16***	0.03***	-0.05
	(0.006)	(0.005)	(0.005)	(0.005)
Extremity	0.24***	0.23***	0.15***	0.28***
	(0.023)	(0.020)	(0.021)	(0.021)
Party Unity	0.27***	0.13**	0.64***	0.31***
	(0.042)	(0.041)	(0.041)	(0.041)
Tenure	-0.00	-0.00***	0.00***	-0.00***
	(0.000)	(0.000)	(0.000)	(0.000)
Constant	-0.18	-0.04	-0.46***	-0.24***
	(0.033)	(0.032)	(0.031)	(0.031)
<hr/>				
r-squared	0.11	0.14	0.10	0.08
N	61,536	70,648	70,417	69,760

Table Note: Cells contain unstandardized coefficients and standard errors from a linear probability model where the dependent variable was scored 1 if the respondent placed their Senator correctly and 0 if the respondent was incorrect or unsure. Results based on a dataset that pools respondents across Senators. Standard errors were clustered at the level of the individual respondent to reflect pooling across Senators. *** = p<.001, ** p<0.01, * p<0.05.

⁴ Running these same regression models with standard errors clustered by Senator leads to significant coefficients on electoral context on: partial birth abortion, immigration, and CAFTA. The other effects are statistically insignificant.

Table 2.5: Regression Models of Citizen Accuracy on Roll Call Issues (Continued)

	<u>Min. Wage</u>	<u>Capital Gains</u>	<u>CAFTA</u>
	Coeff.	Coeff.	Coeff.
Election	0.03***	0.04***	0.05***
(s.e.)	(0.004)	(0.004)	(0.004)
Education	0.04	0.05	0.03
	(0.002)	(0.002)	(0.002)
Pol.Knowledge	0.33***	0.35***	0.21***
	(0.007)	(0.007)	(0.006)
Co-Partisan	-0.06***	-0.07***	-0.07***
	(0.005)	(0.005)	(0.004)
Ideology	-0.23***	0.12***	0.11***
	(0.005)	(0.005)	(0.005)
Extremity	0.13	0.09***	-0.06***
	(0.021)	(0.021)	(0.019)
Party Unity	0.68***	0.65***	0.00***
	(0.041)	(0.042)	(0.000)
Tenure	-0.00***	0.00***	0.000***
	(0.000)	(0.000)	(0.000)
Constant	-0.44***	-0.51***	-0.09***
	(0.031)	(0.032)	(0.029)
<hr/>			
r-squared	0.16	0.11	0.05
N	69,762	68,841	69,317

Table Note: Cells contain unstandardized coefficients and standard errors from a linear probability model where the dependent variable was scored 1 if the respondent placed their Senator correctly and 0 if the respondent was incorrect or unsure. Results based on a dataset that pools respondents across Senators. Standard errors were clustered at the level of the individual respondent to reflect pooling across Senators. * p<0.05, ** p<0.01, *** = p<.001.

It is interesting to note the consistent, negative coefficient on the variable that measures whether the respondent was a member of the same party as the senator (the row labeled “co-partisan”). This pattern may suggest that co-partisans are more prone to projecting their own views onto members of their own party than they are to projecting disagreement onto members of the opposite party. The coefficient on senator ideology is also interesting, as it seems to serve as a proxy for whether the issue divided the Republicans. On issues where Republicans were divided (stem cells, immigration reform, and minimum wage), the more conservative the lawmaker the harder it was to place them.

I also pooled the data across issues to examine the overall effect of electoral context on the accuracy of roll call placements. The top rows of Table 2.6 replicate the analysis using this dataset, where the observation is the respondent-senator-issue. Overall, the results reveal that electoral context has a small but statistically significant effect on citizen accuracy across issues

(about three percentage points).⁵ The results are marginally better than we would expect if pure guessing were at work, but not by much.

The leftmost column of Table 2.7 contains the analogous regression results. After controlling for other respondent and senator characteristics, electoral context is associated with a five percentage-point increase in the probability of correctly placing the senator. The results reveal some support for hypothesis 2, but it is limited: electoral context does increase in accuracy on some roll call votes, but these effects are generally small.

Table 2.6: Citizen Accuracy across Electoral Conditions and Type of Vote
(Pooling Across Issues)

		All Votes			
		Correct	Incorrect	Unsure	Difference
Non-Election		45.8	11.8	42.4	3.1
		(45.4 to 46.2)	(11.6 to 12.0)	(41.9 to 42.9)	(2.8 to 3.4)
Election		48.9	13.4	37.7	
		(48.5 to 49.4)	(13.1 to 13.6)	(37.2 to 38.2)	
		Party Line Votes			
		Correct	Incorrect	Unsure	Difference
Non-Election		48.6	9.6	41.9	6.0
		(48.1 to 49.0)	(9.4 to 9.8)	(41.4 to 42.4)	(5.4 to 6.4)
Election		54.5	9.4	36.1	
		(54 to 55.0)	(9.2 to 9.6)	(35.6 to 36.6)	
		Defecting Votes			
		Correct	Incorrect	Unsure	Difference
Non-Election		23.1	30.1	46.7	-3.3
		(22.5 to 23.7)	(29.5 to 30.8)	(45.9 to 47.6)	(-4.1 to -2.5)
Election		19.8	34.1	46.1	
		(19.2 to 20.4)	(33.3 to 34.8)	(45.3 to 47.0)	

Table Note: Cells contain percentage of respondents answering correctly, incorrectly, and unsure on the item in question and the naïve 95% confidence interval for that estimate (not adjusted for heteroskedasticity). “Difference” refers to the difference in the percentage correct across electoral settings. Results based on a dataset that pools across Senators and issues. Confidence intervals for the estimated differences do reflect standard errors that are clustered at the level of the individual respondent. Calculations reflect sampling weights provided by Polimetrix.

⁵ The effect is larger—about 5 percentage points—if we exclude Bill Nelson’s constituents from the analysis.

**Table 2.7: Regression Models of Citizen Accuracy by Vote Type
(Pooling across Issues)**

	All Votes Coeff.	Party Line Coeff.	Defecting Coeff.
Election	0.05***	0.07***	-0.04***
(s.e.)	(0.002)	(0.002)	(0.004)
Education	0.04***	0.04***	0.01***
	(0.001)	(0.001)	(0.002)
Pol. Knowledge	0.31***	0.34***	0.11***
	(0.005)	(0.005)	(0.006)
Co-Partisan	-0.06***	-0.07***	0.04***
	(0.003)	(0.004)	(0.005)
Ideology	-0.00	0.00	-0.03***
	(0.004)	(0.004)	(0.006)
Extremity	0.12***	0.08***	0.02
	(0.014)	(0.015)	(0.036)
Party Unity	0.43***	0.29***	-0.30***
	(0.026)	(0.033)	(0.048)
Tenure	0.00	-0.00***	0.00***
	(0.000)	(0.000)	(0.000)
Constant	-0.29***	-0.14***	0.30
	(0.093)	(0.114)	(0.117)

R2	0.08	0.09	0.03
N	480,281	419,570	60,234

Table Note: Cells contain unstandardized coefficients and standard errors from a linear probability model where the dependent variable was scored 1 if the respondent placed their Senator correctly and 0 if the respondent was incorrect or unsure. Respondents are pooled across Senators and issues. Standard errors were clustered at the level of the individual respondent to reflect pooling. * $p < 0.05$, ** $p < 0.01$, *** = $p < .001$.

Party Line vs. Defecting Votes

Party cues can help voters place senators on a given issue if they know where the parties stand on the issue and the senator votes in line with the party. If the Senator bucks the party line, then voters who only know the other two facts (party position and senator's party identification) will make an incorrect inference. If, however, campaigns provide voters with more detailed information about their senator's issue positions—either actual knowledge of the roll call or a sense that their senator does not share the party's views in a given domain—then we would expect voters in an electoral context to be more accurate on defecting votes as well.

The bottom two-thirds of Table 2.6 report the effect of electoral context disaggregated across each type of vote. The results reveal that the small, positive effect of electoral context documented above actually masks two distinct patterns. On party-line votes, electoral context has a strong and significant effect on accuracy. Respondents placing a senator who is up for reelection on a vote where the senator took the party line are 6 percentage points more accurate

than they are in a non-electoral context. The story for non-party line votes is almost the opposite. Electoral context has a negative effect on accuracy, decreasing the probability of being correct by just over three percentage points.

These patterns suggest that party cues are driving most of the increase in accuracy, not deeper knowledge of roll call votes. When senators stick with their party, citizens who know their legislators' party identification can more easily infer where the Senator stood on that issue. And because electoral context seems to make the party identification of incumbents more accessible, this translates to more accurate placements—but only on votes that adhered to the party line. In contrast, accessible information about the senator's partisanship is of little help when the vote defies expectations. In fact, the results suggest that it may even hurt. Note, too, that individuals facing non-party line votes are about equally likely to say that they are unsure. So the increase in incorrect answers is not just a function of unlucky guessing. If respondents were just blindly guessing on these issues, we would expect the percentage correct and incorrect to be about even on these issues. Instead, in both electoral and non-electoral settings, far more respondents are wrong than right. Those who answer are actively making mistakes.

The second two columns of Table 2.7 tell the same story, this time using regression results. The overall effect of electoral context is somewhat larger after controlling for respondent and senator characteristics (about five-and-a-half percentage points). But this overall effect masks the fact that electoral context has a pronounced, positive effect on accuracy when it comes to party line votes and a significant, negative effect on non-party line votes. Constituents evaluating a senator who is up are about 7 percentage points more likely to know the roll call position on party line votes. On non-party line votes, they are 4 percentage points less likely to be correct than those evaluating a Senator who is not up.⁶

These results reveal two things. First, while the test is admittedly conservative, there is little evidence that constituents have roll-call knowledge that is independent of party cues. Second, party-based rules of thumb can lead voters to make mistakes when votes are not in the expected direction. The electoral context appears to make voters more prone to such mistakes, perhaps because the campaign has made party cues more salient. This negative effect is not particularly large and is somewhat sensitive to the collection of senators included in the analysis (excluding Bill Nelson decreases the negative effect to about -3.5 percentage points) but it remains significant when standard errors are clustered at the respondent level.

However, what is clear is that any apparent gain in citizen accuracy that may result from an information-rich campaign environment disappears when a senator's vote defies party labels. Based on this conservative test, the results provide little support for hypothesis 3.

Confirming the Patterns: Within Subjects Comparison

The Senate's staggered terms present an opportunity to look at some of these differences using a "within subjects" design. In any given election cycle, some constituents have one senator up while the other is off. By comparing their perceptions of the senator who is up to perceptions of the senator who is not, we can analyze differences across contexts while holding any individual-level characteristics constant (Franklin (1991) uses a similar design).

In an effort to validate the patterns seen above, I looked at differences in accuracy "within subjects" on the party identification question and across the different types of votes. In

⁶ The regression results are robust if we cluster the standard errors by senator. The coefficient on electoral context remains positive and significant for both the overall and party-line models, and is marginally significant ($p=0.06$) in the model of defecting votes.

the dataset are only those respondents who had one senator up for reelection. I compare how accurate they were in placing the campaigning senator to how accurate they were in placing the other. I pool respondents across issues. Table 2.8 displays the results.

Table 2.8: Within Subjects Comparisons (Pooling Across Issues)

Party Identification Recall				
	Correct	Incorrect	Unsure	Difference
Non-Election	76.5	4.0	19.6	7.3
	(75.8 to 77.1)	(3.7 to 4.3)	(18.9 to 20.1)	
Election	83.7	2.7	13.6	
	(83.1 to 84.3)	(2.4 to 2.9)	(13.1 to 14.1)	

Issue Positions (All Votes)				
	Correct	Incorrect	Unsure	Difference
Non-Election	47.6	12.0	40.4	1.0
	(47.4 to 47.9)	(11.8 to 12.2)	(40.1 to 40.6)	
Election	48.6	13.5	37.9	
	(48.3 to 48.9)	(13.3 to 13.6)	(37.6 to 38.2)	

Party Line Votes				
	Correct	Incorrect	Unsure	Difference
Non-Election	50.9	9.5	39.6	3.8
	(50.3 to 51.4)	(9.2 to 9.7)	(39.1 to 40.2)	
Election	54.7	9.1	36.2	
	(54.2 to 55.2)	(8.8 to 9.3)	(35.7 to 36.8)	

Defecting Votes				
	Correct	Incorrect	Unsure	Difference
Non-Election	26.9	27.6	45.5	-7.4
	(26.1 to 27.7)	(26.8 to 28.4)	(44.4 to 46.5)	
Election	19.56	34.4	46.0	
	(19.0 to 20.2)	(33.7 to 35.2)	(45.1 to 46.9)	

Table Note: Cells contain the percentage of respondents answering correctly, incorrectly, and unsure on the item in question and the naïve 95% confidence interval for that estimate. “Difference” refers to the difference in the percentage correct across electoral settings. Results based on a dataset that pools across Senators and issues. Confidence interval on the difference across groups reflects standard errors clustered at the level of the individual respondent to reflect pooling. Calculations reflect sampling weights provided by Polimetrix.

The results confirm that individuals have an easier time recalling the partisanship of their senator when he or she is up for reelection. Citizens are seven percentage points more likely to be correct on this question and less likely to be wrong or unsure. There is little evidence that this is just the result of lucky guessing. Interestingly, the within-subjects analysis produces no evidence that electoral context increases citizen accuracy across all roll call items. Respondents are both more likely to be right and more likely to be wrong when evaluating their senator who is up for reelection, and the increase in incorrect answers outweighs the increase in correct ones.⁷ As the results of the previous section suggest, however, there are two distinct patterns across the different types of votes. When we disaggregate, we see that citizens are somewhat more accurate in on party line votes (about four percentage points) but are much less likely to be accurate on questions where the senator voted with the other party. Respondents are about seven percentage points less likely to accurately place the senator up for reelection on these types of votes. Even after excluding Bill Nelson's constituents, respondents are still four-and-a-half percentage points less likely to be accurate on defecting votes.

The within-subjects results seem to confirm two findings. First, the context of a campaign does lead citizens to be more informed about the partisanship of their incumbent Senator. But the strong tendency to be incorrect on defecting votes suggests that these salient party cues can lead citizens to make mistakes, and that these mistakes are more likely in an electoral context.

Discussion

Thus we have good news and bad news for citizen competence. The good news is that constituents are more likely to know their Senator's party identification when they are up for reelection. The second bit of good news is that this seems to lead constituents to be somewhat more accurate in placing their Senator on roll call votes, but only when those votes follow the party line. The bad news for citizen competence is that blindly relying on cues can sometimes hurt accuracy as much as it helps. Psychologists have argued that heuristics often lead individuals astray, but political scientists have typically touted their positive effects (though see Kuklinski and Quirk (2000) for an exception). The truth is, the political world is simply too complex for a simple rule-of-thumb to be right all the time.

One need only look to the data for anecdotal evidence of how this plays out. While Senator Kay Bailey Hutchison, a Republican from Texas, voted with the Democrats in favor of stem cell research, 53 percent of her constituents presumed that she voted against the bill. Just 10 percent placed her accurately (the rest were unsure). While Senator Mike DeWine, a Republican from Ohio, voted in favor of a minimum wage increase, 50 percent of his constituents assumed that he had voted in the opposite direction, and just 11 percent placed him correctly.

What process is driving these systematic inaccuracies, and why might the electoral context exacerbate them? The data do not allow us to examine the second part of the question in any detail, but they can shed some light on the first. The literature on partisan projection suggests that individuals will infer their member's positions on important issues from their own position and from the position of the party (Wilson and Gronke 2000). Using this logic, co-partisans will presume that a member of their own party voted in a way they agree with and vice versa for a member of the opposing party. In a partisan, sorted world, this rule of thumb is likely to be correct (Levendusky 2009). But projection clearly poses a problem on votes where the legislator defects from his or her party.

⁷ This is partly a function of Florida voters. If I remove them from this analysis, the improvement in accuracy rises to about 2.3 percentage points, and the difference in the percentage incorrect drops to almost even.

It seems plausible that voters in an electoral setting might be more prone to projection than those who are not. The facts necessary to engage in projection are all likely to be more accessible in the midst of a campaign. Knowledge of the legislator’s partisanship, the position of the party on the issue, and citizens’ own position on the issue are all likely to be more salient in an electoral context. It is possible that the availability of these facts leads voters to be even more prone to projection in an electoral context.

To test this hypothesis, I used the data to examine whether partisans were more vulnerable to projection on these defecting votes when they were in the midst of a campaign. The data are drawn from the between subjects dataset where respondents are pooled across senators and issues. Table 2.9 displays the results.

Table 2.9: Projection on Defecting Votes by Electoral Condition

Opposing Partisans						
Non-Election				Election		
Actual:				Actual:		
Perceived:	Disagree	Agree		Perceived:	Disagree	Agree
Disagree	28.5 (26.8 to 30.2)	44.0 (42.9 to 45.2)		Disagree	31.0 (29.1 to 33.0)	53.3 (52.0 to 54.6)
Agree	24.8 (23.2 to 26.5)	20.2 (19.3 to 21.1)		Agree	22.8 (21.1 to 24.7)	16.0 (15.1 to 17.0)
Unsure	46.8 (44.9 to 48.7)	35.8 (34.7 to 36.9)		Unsure	46.1 (44.0 to 48.3)	30.7 (29.5 to 31.9)

Co-Partisans						
Non-Election				Election		
Actual:				Actual:		
Perceived:	Disagree	Agree		Perceived:	Disagree	Agree
Disagree	28.1 (27.1 to 29.21)	20.2 (18.8 to 21.6)		Disagree	20.1 (19.0 to 21.1)	19.9 (18.4 to 21.4)
Agree	33.0 (31.9 to 34.1)	29.9 (28.3 to 31.5)		Agree	40.4 (39.1 to 41.7)	32.0 (30.3 to 33.8)
Unsure	38.9 (37.7 to 40.1)	49.9 (48.2 to 51.7)		Unsure	39.6 (38.3 to 40.9)	48.1 (46.2 to 50.0)

Table Note: Table divides respondents according to whether they actually agreed or disagreed with their senator on a roll call issue (columns) and whether they perceived agreement or disagreement with the senator on that issue (rows). Each cell indicates the proportion of respondents in each condition. Only defecting votes are included here. The cells highlighted in gray correspond to partisan projection in the text: co-partisans perceiving agreement when they actually disagree, and opposing partisans perceiving disagreement when they actually agree. Respondents are pooled across Senators and issues.

The tables are organized as follows: the columns correspond to whether the respondent actually agreed with the Senator on the issue in question. The rows correspond to whether the respondent perceived disagreement or agreement with the senator's vote. In other words, the diagonals from left to right are correct answers (perceive disagreement when you actually disagree and vice versa), while the off-diagonals are incorrect answers (perceiving disagreement when you actually agree and vice versa). The tables disaggregate by whether the respondent was a member of the Senator's party or not (opposing partisans are the top table, co-partisans are the bottom), and electoral context (elections on the right, non-elections on the left). The percentage engaging in projection is shaded in gray. For opposing partisans, a projection error occurs when they perceive disagreement but actually agree, and for co-partisans it is when they perceive agreement but actually disagree.

Table 2.9 indicates that on defecting votes, individuals evaluating an incumbent Senator up for reelection are more prone to projection than others—by about seven-and-a-half to nine percentage points. Put another way, on defecting votes, respondents in an electoral context are seven to nine percentage points more likely to be incorrect. This is distinct from a decline in the proportion answering correctly (though the tables show some evidence of that as well). When it comes to counter-stereotypical votes, voters evaluating an incumbent who is running for reelection appear to be more prone to projection errors than those evaluating an incumbent who is not up.

It is interesting to think about these findings in the context of research on candidate ambiguity (Tomz and Van Houweling 2009). Franklin's (1991) study of senate campaigns found that when senators stressed issues, constituents had a clearer sense of where they stood ideologically. When incumbents avoided the issues, voters had far less clarity. It seems plausible that part of the process underlying the relationship between electoral context and partisan projection could be due to candidate ambiguity on these cross cutting issues. In an effort to appeal to both their base and more moderate voters, incumbent candidates may seek to stake an ambiguous position on issues that divide their constituency, providing fertile ground for projection by partisan voters.

Conclusion

Some of these results seem to be yet another testament to the frailty of citizen competence. Citizens are only marginally more accurate in placing their senators on roll call votes when those senators are up for reelection, and they are prone to systematic mistakes on votes that run counter to party labels.

On the other hand, I see two reasons why the limited good news unearthed here may outweigh the bad. First, in a polarized, sorted world, citizens using a simple rule of thumb based on party cues are likely to be accurate more often than not. Indeed, given the trends of the last fifteen years, their rate of accuracy will only improve going forward. If congressional campaigns provide voters with information about their senator's party identification and the positions of the parties, then the electoral context will help more often than it hurts. If most of the votes go in a predictable direction, then voters will have fewer and fewer misperceptions.

Second, in a general election, party line votes are likely to garner the most attention. After all, these positions are more likely to distinguish the incumbent from a challenger. Instead, defecting votes are more likely to come up in primary challenges from more extreme candidates, as was the case with former Utah Republican Robert Bennett, Alaska Republican-turned-Independent Lisa Murkowski, and Connecticut Democrat Joe Lieberman. In each case, these incumbents were attacked for defecting from their party on particular roll call votes. But

primaries are notoriously low turnout affairs, meaning rank and file voters probably have less opportunity to learn from the messages they transmit.

This discussion also raises the possibility that the absence of positive information effects on defecting votes results from a lack of campaign messages about those votes. If media coverage and advertisements do not cover issues on which the incumbent broke with his party, then voters should not be any more likely to know about these issues positions in an electoral context. That being said, we would not expect them to know less, either.

Finally, this work points out the need for further research on two topics. First, political scientists tend to tout the positive benefits of shortcuts and heuristics, assuming that they help voters. But the results here suggest that voters are not particularly good at knowing when to use party cues and when they may be less helpful. If voters in the electoral context recognized that heuristics were less useful on defecting votes, we would not have seen significantly lower rates of accuracy. Political scientists should take a harder look at whether voters can distinguish the conditions under which commonly used cues and heuristics are likely to be more or less useful.

Second, the results suggest a need to look at these questions with a finer-toothed comb, one that examines individual campaigns and the roll call issues that became salient. The issues that become salient clearly vary by campaign. Such a study was beyond the scope of this article, but occasionally examples came to the fore. For instance, we know that stem cell research came up repeatedly in the 2006 Missouri Senate race between incumbent Jim Talent (R) and challenger Claire McCaskill (D). It was the subject of a controversial pro-McCaskill advertisement starring Michael J. Fox, who voiced his support for stem cell research. Rush Limbaugh made national news when he criticized the ad. Perhaps not surprisingly, the data show that fully 75 percent of Jim Talent's constituents knew that he was opposed to stem cell research. Just 50 percent could accurately place Talent's fellow incumbent, Republican Kit Bond, on the same issue. Further research should use the 2006 CCES to take a more detailed look at variation in issue content across campaigns (see Barabas and Jerit (2009)).

Chapter 3

Policy Performance and Public Support

In the budget battle of early 2011, even the most prized government programs were on the chopping block. In a politically risky move, House Republicans set their sights on making cuts to the popular Head Start program, which has provided early childhood education and social services since the mid-1960s. As the *New York Times* reported, of the programs the Republicans targeted for cuts, Head Start was “the most visible and popular”: “During townhall meetings . . . many voters across the country, some carrying children, pleaded for the program. Editorial boards have zeroed in on cuts to Head Start as draconian” (Steinhauer 2011, A14).

Republicans argued that targeting Head Start was not an example of indiscriminant cutting, but a strategic decision to reduce funding for a program that had not accomplished its goals. They cited a disappointing 2010 evaluation by the Department of Health and Human Services and a recent investigation by the Government Accountability Office (GAO) that uncovered evidence of enrollment fraud. As a spokesman for Rep. Duncan Hunter (R-CA) told reporters, “The program has benefited from sizable increases in federal funding in recent years, all the while the program continues to underperform in certain aspects” (Steinhauer 2011, A14).

In light of the popularity of Head Start, it is not clear that this argument would resonate with the public, particularly the program’s traditional allies in the Democratic Party. This vignette raises a broader question for students of political behavior: do voters’ policy preferences reflect their view of how well that policy performs? Theories of retrospective voting suggest that voters keep their eye on how conditions change and base their votes on whether things have gotten better or worse. Whether the same logic applies to individual policy attitudes is less well understood. Do messages about a social program’s success or failure cause voters to update their attitudes? Or are these beliefs subject to the same ideological and partisan biases that lead to motivated reasoning and biased updating?

Political behavior research has typically analyzed policy preferences in the mass public at a more general level than specific programs or policies. Estimates of policy preferences typically come from abstract issue scales, and the work on political sophistication and constraint has often asked how these general policy preferences hang together (Converse 1964). While these general attitudes are certainly important to the study of representation and party politics, the research on “issue publics” and “operational liberalism” suggests that individuals often have preferences about individual policies that may be distinct from their overarching ideology (Krosnick 1990). Given that members of Congress often face up or down votes on specific policies, “policy-specific” attitudes, and the information on which they rest, are also critical to public accountability and representation (Gilens 2001).

Fortunately, a recent strand of research has focused on the importance of “policy-specific information” in shaping citizens’ attitudes. Specifically, studies have asked how knowledge of policy facts—like the amount spent on a program or the types of citizens the program serves—relate to the way people feel about those policies (Gilens 2001; Kuklinski et al. 2000). Others have examined the process by which citizens learn these facts, highlighting the effect of media coverage on voters’ policy-relevant knowledge (Barabas and Jerit 2009). And a third strand of work has asked whether members of the mass public will correct misperceptions about policies when confronted with the facts (Kuklinski et al. 2000; Gaines et al. 2007; Nyhan and Reifler 2010).

This growing literature has highlighted the important influence that policy-specific information has on policy preferences. But we know less about which types of policy-specific information are more likely to affect voters' attitudes and preferences. Drawing on data from two large surveys, this chapter focuses on how messages about whether an existing government program is working as intended affect attitudes toward that program. The surveys asked respondents about their perceptions of and support for two prominent social programs: Food Stamps and the aforementioned Head Start. Prior to being asked about these programs, a random subset of respondents received data from actual government evaluations, providing an opportunity to identify the causal effect of performance information.

The data allow for a test of two different messages, one positive and one negative. After decades of criticism for fraud and abuse, the Food Stamp program has implemented a strict quality control system that has led to a marked increase in efficiency. In contrast, Head Start has remained a cornerstone of Great Society-era social programs despite evaluations that have raised questions about the program's effectiveness. The question is not only whether these different messages have an impact on citizen attitudes, but also whether respondents of different ideological leanings respond in different ways. Do liberals discount negative findings about a cherished program? Will conservatives downplay positive data about one they typically dislike?

The results suggest that performance information has a significant impact on citizens' perceptions of program effectiveness and their policy preferences. Respondents who received word that the Food Stamp program had become more efficient over the past decade became less likely to label the program "ineffective" or to call for "major changes." Those who learned about Head Start's discouraging evaluation results revised their beliefs downward. Counter to some of the research on ideological bias, I find no evidence that ideological groups responded differently to these messages. Liberals and conservatives and Democrats and Republicans tended to revise their beliefs in parallel, maintaining existing levels of disagreement but shifting the distribution in the direction of the new information.

The findings have implications for theories of congressional policymaking. Arnold (1992) argues that members of Congress must be sensitive to the "potential preferences" of their constituents when confronting a roll call vote on a given policy. In their retrospective voting calculations, constituents take note of policy outcomes and then attempt to link those outcomes to their legislator's votes. These perceptions of policy outcomes are therefore likely to shape the way that members vote on both new policies and those that come up for reauthorization. If citizens update their beliefs in response to information about policy outcomes, then members of congress may have incentive to maintain "good" policies and fix "bad" ones. If citizens do not, then policymakers may feel less pressure to ensure that government programs are effective.

The analyses also raise questions about legislators' ability to shape perceptions and preferences through the strategic use of evaluations and investigations. By embedding evaluation requirements in legislation, requesting studies by the Government Accountability Office, or holding hearings that publicize the results of a given program, members can send messages about policy performance to the public. This information can then be used to justify congressional action in the future. This chapter sheds some early light on the micro foundations that might explain larger patterns in congressional policymaking.

The chapter proceeds as follows. First I review some of the existing literature on "policy-specific information" and link it to research on congressional policymaking. In section two I describe the data, methods, and the programs in question. The third section describes the results of each study, and the fourth discusses the findings.

Policy Information and Preferences

Since the study of public opinion began, researchers have been preoccupied with how much citizens know about policy issues. Early studies of representation found that rank and file voters rarely knew where their members of congress stood on particular issues (Miller and Stokes 1963). Converse's (1964) classic study of citizen beliefs found that few exhibited ideological "constraint"—a sense of "what goes with what"—and that many often failed to hold consistent beliefs from one point in time to another (212). More recent work on political knowledge has found that sophistication has changed little since these initial studies, despite gains in educational attainment and the availability of information (Delli Carpini 2005). Political scientists have argued that the skewed distribution of political knowledge has real consequences for policymaking and representation. Scholars have used simulations to show that attitudes and voting behavior would look quite different if all voters were "fully informed" (Bartels 1996; Althaus 1998).

In the early 2000's, a set of political scientists broadened our conception of political awareness beyond general knowledge about politics and government. After all, knowing who holds a majority in the House of Representatives or sits on the Supreme Court is unlikely to directly shape policy attitudes or voting behavior. Instead, these authors argue that "policy-specific information" (Gilens 2001) or "policy-relevant facts" (Kuklinski et al. 2000) are critical to understanding policy preferences. In his study of policy-specific information, Gilens (2001, 380) concludes

policy-specific ignorance may well have a greater influence on political preferences than a lack of general knowledge as measured by political information scales. . . . the results suggest that much of what separates actual political preferences from hypothetical 'enlightened preferences' is due to ignorance of specific policy-relevant facts, not a lack of general knowledge or the cognitive skills or orientations that measures of general political information reflect.

When citizens were provided with policy-specific facts, Gilens (2001) found that the messages led them to update their preferences, and that policy information was a far more important predictor than general political information.

Subsequent work has explored how the political environment impacts the level and character of policy-specific knowledge. Research on the "information environment" has found that citizens can and do learn about policy when the media covers policy issues. In a series of studies, Jerit and Barabas (2006), Jerit, Barabas, and Bolsen (2006), and Barabas and Jerit (2009) document how media coverage of public policy issues affects the public's policy-specific knowledge. In their study of 50 political issues Jerit et al. (2006) found that while education is an important predictor of policy knowledge, that relationship varies significantly depending on the extent of media coverage. When media attention is significant, even less educated citizens pick up on elite messages. Using a within-subjects design and variation in media coverage, Barabas and Jerit (2009) show that policy specific messages had a strong causal effect on individuals' knowledge of gun control, Social Security, and healthcare issues.

Whether this knowledge leads to changes in attitudes is another story. Kuklinski et al. (2000) found that citizens with the most inaccurate perceptions of welfare were also the most confident in their beliefs, and that additional factual information had mixed effects on policy preferences. Similarly, Gaines et al. (2007) use panel data to document how citizens updated their attitudes about the Iraq War in light of new information. While Republicans and Democrats

generally held similar factual beliefs about the state of the war effort, they interpreted those facts differently in order to rationalize their prior beliefs.

Recent studies have used experimental survey designs to assess the effect of additional policy information, but they have often come to mixed conclusions. Like Gilens (2001), some have found evidence that policy-specific facts can shift policy preferences. For instance, Howell and West (2009) found that providing facts about educational spending and teacher salaries affected individuals' policy preferences, but that these effects varied significantly across different groups of respondents.

Other experiments have found little evidence of updating. In their study of immigration attitudes, Sides and Citrin (2007) provided respondents with information about the number of immigrants and the annual immigration rate; neither piece of information affected attitudes. Nyhan and Reifler (2010) explored whether voters updated their factual beliefs when provided with new information that explicitly corrected misperceptions about the Iraq War, the Bush tax cuts, and stem cell research. The authors found that not only were ideologues on both sides less likely to correct inaccurate beliefs when messages ran counter to their priors, but that providing such information actually led conservatives to believe the claims "even more fervently than those who did not receive a correction" (Nyhan and Reifler 2010, 320).

However, Bullock's (2011) recent paper provides a direct test of whether policy-specific information can affect policy attitudes in the presence of overt partisan cues. Respondents were provided with a newspaper article about a proposed change to healthcare policy. Some received information about a planned expansion, others about a plan to reduce benefits. A subset of respondents also received a partisan cue. Some of these cues were predictable (Democrats support an expansion or oppose a reduction), while others were "counter-stereotypical" (Democrats oppose an expansion or support a reduction). In contrast to earlier work, Bullock (2011) found that while partisan cues affected respondent attitudes, these effects were generally "swamped" by the effect of policy information (500). On balance, Bullock (2011) concludes that these experiments "[bring] a wealth of evidence to bear on the idea that party cues inhibit processing of policy information. The idea does not hold up well" (509).

Is All Policy Information Created Equal?

This recent literature has shown that facts about policy can have an important influence on preferences under some conditions but not others. The divergence is due, in part, to a failure to specify which type of policy information is likely to affect attitudes. Researchers have tested all manner of facts, from data on the crime rate to the proportion of the budget that goes to welfare to the effects of the Bush tax cuts. While these messages clearly vary considerably in terms of the values and priorities they highlight, the implicit assumption is that all messages should lead all voters to update their beliefs.

But as the introductory chapter discussed, not all messages are created equal. Recall that the key quantity in Bayesian models is the precision of the new message relative to the precision of the prior belief. The more precise the message, the larger the effect it will have on the posterior belief. If messages about policy are perceived as imprecise—perhaps they are vague, misleading, or focus on dimensions that are not important to voters—then they will have less of an impact. In contrast to earlier studies, this chapter hypothesizes that a particular type of information is quite precise: data on the effectiveness and efficiency of government programs.

Why might this kind of information be particularly potent? For decades, political scientists have argued that information about policy outcomes and government performance are

central to citizens' retrospective voting calculations (Key 1966). Arnold (1992, 44) lays out the logic clearly:

An appraisal could look something like this. A citizen notices that some condition that matters to him is either deteriorating or improving. He searches for governmental responsibility, asking how the federal government contributed to the changed conditions. He then searches for the contribution of his own representative to that governmental action. Finally he rewards or punishes his representative.

Intuitively, messages about policy outcomes are likely to be both easy to understand and highly diagnostic. In his theory of retrospective voting, for instance, Fiorina (1981) argues that in contrast to information about policy design, voters are both more concerned about and better able to understand information about policy outcomes. Policy performance messages also provide voters with both a criteria by which to judge policies (whether they are accomplishing public goals) and the data with which to evaluate them.

As the introduction pointed out, this type of outcome information has also become more abundant as policymakers have lurched toward an “evidence-based” approach to policymaking. In an era of tight public budgets, a bipartisan mix of policymakers has argued that funding should be reserved for programs that have been shown to work (Bornstein 2012; Liebman 2013; Brass, Nunez-Neto, and Williams 2006). Those that do not pass a cost-benefit test should be reformed or abandoned altogether. Presidents Bush and Obama both adopted these arguments, to the chagrin of many policy advocates. Policymakers have increasingly used outcomes-based arguments to make the case for policy reform.

But we know less about how voters are likely to respond to such messages. The research on policy-specific information has typically focused on facts about policy design—who receives benefits or how much it costs—rather than policy outcomes. Gilens' (2001) look at whether data about the crime rate affects preferences on prison spending comes close, but the causal chain is actually reversed: a drop in the crime rate is not evidence that prison policy is working. Jerit's (2009) paper on “predictive appeals” also touches on the topic, but she tests the impact of predictions about outcomes, not information about actual outcomes. The recent studies of public opinion on the Iraq War by Gaines et al. (2007) and Nyhan and Reifler (2010) did look explicitly at perceptions of policy outcomes, but these beliefs were tightly coupled to feelings about President Bush, setting the stage for partisan bias. As such, the findings may not be generalizable to other policy domains.

The current study seeks to fill this gap in the political information literature in three ways. First, the surveys analyzed here focused citizens' attention explicitly on policy performance—does a given policy work as intended, or is it ineffective, wasteful, or some combination of the two? The items focused on two well-known social policies and provided respondents with real information from long-term, government-sponsored evaluations of those programs.

Second, the data allow for an assessment of how positive and negative messages about policy performance might affect attitudes. Over the past decade, the Food Stamp program has become more efficient, resulting in savings for the federal government. In contrast, recent evaluations of Head Start have raised concerns about fraud and ineffective results.

Third, the messages analyzed here were likely to run counter to the beliefs of conservatives in one case (Food Stamps) and liberals in the other (Head Start). As such, the design allows me to explore whether these distinct groups were equally willing to update their beliefs.

Hypotheses

As the literature review suggests, efforts to assess the effects of policy information have produced mixed results. However, information on policy outcomes should be particularly high in what Kuklinski et al. (2001) call “diagnostic value.” The messages studied here provided respondents with a clear measure of how a policy was performing on the federal government’s own metrics. Therefore, I hypothesize that subjects will be responsive to information about policy performance:

Hypothesis 1: (Responsiveness) Respondents in the treatment group will significantly shift their policy attitudes when provided with policy performance information.

Of course, citizens with different predispositions may interpret policy information in different ways. On the one hand, we might expect information that runs counter to existing beliefs to be new to the respondent, and therefore more likely to generate attitude change. Confirmatory information may bolster existing support or solidify opposition, but discordant information could be more potent in compelling citizens to revisit their beliefs. However, the large literature on “motivated reasoning” and bias would suggest an alternative hypothesis: people will reject information that does not jibe with their preexisting beliefs (Zaller 1992). Partisans and ideologues may even “double-down” on inaccurate beliefs in the face of contrary information (Nyhan and Reifler 2010; Bullock 2007).

The policies at the center of this analysis are likely to be less polarizing and divisive than those examined in other work. For one thing, they are not as closely linked to political figures as issues like the Iraq War, tax cuts, or the Affordable Care Act. As such, it is more difficult to predict how these conditional effects will play out. Given the large literature on ideological bias, however, I hypothesize that the effect of performance information will be conditional on respondents’ ideological predispositions.

Hypothesis 2: (Conditional Information Effects) Information effects will be conditional on ideology. Ideologues will be less responsive to information that runs counter to predispositions.

Research in social psychology has uncovered a “negativity bias” in the way individuals respond to information (Ito et al. 1998). As Soroka (2006) explains, “the effect of a 1-unit increase in negative news is not simply the opposite of a 1-unit decrease; rather, increases in bad news may matter a good deal more” (372). Political scientists have documented this phenomenon the way citizens evaluate candidates, the effects of negative campaigning, and asymmetries in the effect of positive and negative economic news on issue salience (Lau 1985; Soroka 2006).

Though the current dataset does not allow for a direct exploration of these asymmetries, we can compare “effect sizes” across the studies. Based on existing literature, I hypothesize that negative performance information will be more potent than positive in shifting attitudes and preferences.

Hypothesis 3: (Negativity Bias) Negative policy performance information will have a larger effect on attitudes than positive information.

Hypothesis 3 is particularly important as it relates to political strategy. If negative evaluations are more potent in undermining public support than positive messages are in cultivating it, it may create an incentive to use policy evaluations as a justification for spending cuts. I return to this theme in the conclusion.

The Objects: Food Stamps and Head Start

In the pantheon of federal social programs, few have been around as long as the Food Stamp and Head Start programs. The federal experience with the Food Stamp Program, or the

Supplemental Nutritional Assistance Program as it was recently renamed, dates back to 1939, when the Department of Agriculture created a small food assistance program. The Kennedy and Johnson administrations rekindled the Food Stamp program, resulting in a pilot from 1961-1964 and the passage of the Food Stamp Act in 1964. Head Start was also formally enacted in 1964 as part of the War on Poverty's Economic Opportunity Act. It was expanded over the course of the 1970s and was reauthorized individually via the Head Start Act of 1981.

In terms of public perceptions the two programs have very distinct profiles. Americans generally perceive the Food Stamp program negatively. Writing in the pre-welfare reform period, Smith (1987) argued that "while one might assume that the [Food Stamp] program's face association with preventing hunger might encourage public support, it does not apparently have any more appeal than welfare. Like welfare, it may be tarred by images as a wasteful, mismanaged program open to abuse" (82). The American National Election Study (ANES) included an item on Food Stamps each year between 1984 and 2000. Across those surveys, between 11 and 18 percent of respondents believed that spending on Food Stamps should be increased, while between 30 and 46 percent of respondents argued that spending should be cut.⁸

In contrast, Head Start has a long track record of public popularity. A 2003 survey found that 92 percent of Americans either strongly supported or somewhat supported Head Start; 89 percent of Republicans joined 91 percent of Democrats in expressing support for the program (Pax World Funds and National Head Start Association 2003). The data below suggest that the majority of Americans believe Head Start effectively prepares low-income students for kindergarten.

Data and Methods

The data analyzed here come from the 2010 Cooperative Congressional Election Study (CCES), administered by Polimetrix in the fall of 2010. The survey experiments were included on a module commissioned by the Program on American Citizenship at the American Enterprise Institute. It asked a range of questions about education policy, civic education and social policy attitudes. The Program on American Citizenship shared the de-identified dataset and documentation for this analysis upon request.

The module was administered to 1,000 respondents. Two-thirds were randomly assigned to either receive information about Food Stamp efficiency or to be in a pure control group (N= 655; 306 treated, 349 control). The Head Start experiments were included on the post-election survey and covered all respondents who completed the second wave (N = 831). Post-election respondents either received information about the Head Start evaluation (N = 277), information about the investigation of Head Start fraud (N=274), or a generic description of the program (the control group, N=280).

The pre-election survey included questions designed to measure respondents' prior views of program performance. For instance, respondents were asked to give their best guess as to whether they thought Food Stamp error rates had increased, decreased, or stayed the same over the past decade. After providing a brief description of Head Start, the survey asked whether respondents thought low-income children enrolled in Head Start did better in elementary school than similar low-income children who did not enroll.

The different informational prompts are described below. The dataset includes three outcome variables of interest: perceptions of the program's effectiveness (a four-point scale

⁸ ANES item asked in each year from 1984-2000. For period 1992 to 2000, the percentage favoring increase ranges from 11 to 18 percent; percentage favoring cuts from 28.7 percent to 46 percent.

anging from extremely ineffective through extremely effective), whether federal funding for the program should be increased, decreased, or kept the same (three-point scale), and whether Congress should eliminate the program, make major or minor changes to it, keep it “as is”, or expand the existing program (a five-point scale). The policy change measure is coded to reflect support for the existing program; those who wanted to expand it are at the top of the scale (five), while those who wanted to eliminate it are at the bottom (one).

The analyses below estimate information effects using simple differences in means and Ordinary Least Squares (OLS) regression models. Because of the categorical nature of the outcome variables, I also replicated the results using ordered Probit models. For ease of interpretation, I report the differences in means and OLS results below, which mirror the results for the Probit models almost exactly. The analysis of main effects reported here does not include controls for demographic and political variables, but the results were robust to those alternative specifications. Adding additional control variables made the patterns even clearer.

In analyzing subgroups, I estimate group-specific treatment effects using a fully saturated model that contains dummy variables for each subgroup and the interaction between those dummies and an indicator of the treatment.⁹ These models are estimated without a constant. To simplify the analysis, I used un-weighted data in estimating all effects.

Results

Prior Perceptions

Pre-treatment perceptions mirrored the distinct patterns of support for these programs found in prior research. In general, individuals were either skeptical or not willing to hazard a guess as to whether the Food Stamp program had improved its error rates over the past decade. Table 3.1 displays answers to this question across the entire sample and broken down by partisanship and ideology. The figures in bold represent the percentage of respondents in each response category (including the “don’t know”), while those in italics calculate those percentages among those who provided an answer other than “don’t know.”

Among those who answered, about 60 percent of respondents guessed that Food Stamp error rates had gotten worse, and 18 percent thought that they had stayed the same. Not surprisingly, perceptions of the program’s recent track record divided along partisan and ideological lines: Republicans and conservatives were much more likely to believe that error rates had gotten worse than Democrats and liberals. This divergence is even more prominent among those respondents who provided answer. Interestingly, conservatives and Republicans also seem somewhat more likely provide an answer, suggesting that they have stronger perceptions of the program.

Pre-treatment assessments of Head Start were quite different. Table 3.2 reveals that the vast majority of respondents believe that Head Start students are likely to do better in first grade than similar students who do not enroll. Even though Head Start is generally popular, however, there are important differences across partisan identifiers. While few Republicans believe that Head Start students do any worse, they are much more likely to believe that enrollees do about as well. Democrats, liberals, and moderates overwhelmingly believe that Head Start students do better than those who do not enroll.

⁹ The model of subgroup effects takes this form: $Y = B(\text{Liberal}) + B(\text{Liberal} * T) + B(\text{Moderate}) + B(\text{Moderate} * \text{Treatment}) + e$, where X is a dummy variable that corresponds to groups (conservatives, moderates, liberals), X*T is an interaction of the dummy variable and the treatment indicator, and e is an error term.

Table 3.1: Respondent’s Perceptions of Food Stamp Performance

	Food Stamp Error Rates Have:			
	Gotten Worse	Stayed Same	Gotten Better	Don't know
Overall	41.8	12.5	15.6	30.2
Answers Only	<i>59.9</i>	<i>17.8</i>	<i>22.3</i>	
Democrats	28.6	17.2	18.6	35.6
	<i>44.4</i>	<i>26.8</i>	<i>28.9</i>	
Republicans	56.4	12.9	6.6	24.2
	<i>74.4</i>	<i>16.9</i>	<i>8.6</i>	
Liberals	26.6	18.6	21.3	33.5
	<i>40.0</i>	<i>28.0</i>	<i>32.0</i>	
Moderates	33.6	15.4	16.1	35
	<i>51.6</i>	<i>23.7</i>	<i>24.7</i>	
Conservatives	59.9	12.5	5.3	22.3
	<i>77.1</i>	<i>16.1</i>	<i>6.8</i>	

Table Note: Cell entries are percentages of that group answering “gotten worse,” “stayed the same,” or “gotten better.” Numbers in italics calculate those percentages after excluding the “don’t know” responses. Rows may not sum to 100 percent due to rounding.

Table 3.2: Respondent’s Perceptions of Head Start Performance

	Head Start Students Do:			
	Worse	The Same	Better	Don't Know
Overall	3.2	23.0	63.3	10.5
Answers Only	<i>3.6</i>	<i>25.6</i>	<i>70.8</i>	
Democrats	0.9	12.9	78.2	7.9
	<i>1.0</i>	<i>14.1</i>	<i>84.9</i>	
Republicans	6.3	31.6	50.1	12.0
	<i>7.1</i>	<i>35.9</i>	<i>57.0</i>	
Liberals	2.3	12.5	77.7	7.6
	<i>2.5</i>	<i>13.5</i>	<i>84.0</i>	
Moderates	0.7	17.1	72.8	9.4
	<i>1.0</i>	<i>18.9</i>	<i>80.4</i>	
Conservatives	5.8	34.5	48.8	11.0
	<i>6.5</i>	<i>38.8</i>	<i>54.8</i>	

Table Note: Cell entries are percentages of that group answering “worse,” “the same,” or “better.” Numbers in italics calculate those percentages after excluding the “don’t know” responses. Rows may not sum to 100 percent due to rounding.

As Tables 3.1 and 3.2 reveal, the positive information about the Food Stamp program and the negative information about Head Start generally run counter to prevailing public attitudes, and this is particularly true among partisans and ideologues who are predisposed to support or be skeptical of the program in question.

Study 1: Food Stamp Error Rates

In the first study, one-third of the sample was randomly chosen to receive information about the trend in Food Stamp error rates, while one-third was assigned to a pure control group. Prior to voicing their opinions about the Food Stamp program, respondents in the treatment group received the following message:

The Department of Agriculture monitors how many errors states make in giving out their Food Stamps. Overpayments are more common than underpayments and waste public money.

Food Stamp error rates were cut in half between 1999 and 2008, from almost 10 percent in 1999 to about 5 percent in 2008. This reduction resulted in financial savings for the government.

All respondents were then asked whether they believed that the Food Stamp program was effective in helping low-income families, whether spending should be increased or decreased, and what they thought Congress should do the next time it considers the program.

Table 3.3 displays the difference in means between treatment and control groups on each of the outcome variables. The first column reveals that information about Food Stamp error rates had a significant, positive effect on perceptions of the program's effectiveness. The positive information caused treated respondents to perceive the program as being more effective than the control group. The probability that a respondent in the treatment group labeled the program "effective" (somewhat or extremely) was more than 9 percentage points higher between the treatment and control groups.

Though information on error rates had a significant effect on perceptions of program effectiveness, it had only mixed effects on policy preferences. On the spending outcome, the information effect is in the right direction, but it falls short of statistical significance. A lower proportion of respondents said that program funding should be decreased in the treatment group than in the control (34 percent vs. 44 percent), but a slightly higher proportion of control group respondents thought that funding should be increased (24 percent vs. 20 percent). On the policy change outcome, there is evidence of a significant information effect, though it is on the smaller side. Respondents in the treatment group were significantly less likely to recommend that Congress eliminate or make major changes to the program than those in the control group. Nearly 50 percent of the control group chose one of those two options (47.6), while 38 percent of those in the informed group recommended that course of action.

Table 3.3: Effect of Food Stamp Error Rate Information on Outcome Variables

	Effectiveness		Spending		Policy Change	
	Mean	95% c.i.	Mean	95% c.i.	Mean	95% c.i.
Control	2.67	(2.56 to 2.77)	-0.20	(-0.29 to -0.11)	2.87	(2.74 to 3.01)
Treatment	2.84	(2.73 to 2.94)	-0.15	(-0.23 to -0.06)	3.09	(2.94 to 3.23)
Effect	0.17	(-0.32 to -0.02) *	0.05	(-0.08 to 0.18)	0.21	(0.02 to 0.41) *
N	600		555		588	

Table Note: Table contains mean ratings and the difference in means between on each outcome scale for respondents who received information about the improvement in Food Stamp error rates over the past decade and those who did not. Symbols correspond to statistical significance levels (+ = $p < 0.1$, * = $p < 0.05$, ** = $p < 0.01$, *** = $p < 0.001$).

However, the results provide little evidence that effects were conditional on ideology. The first pattern to note is the expected relationship between ideology and beliefs in the control group. In the absence of information, conservatives' views of the Food Stamp program are considerably less positive than liberals or moderates. However, it is not the case that conservatives and liberals responded differently to positive information about the program. When it came to perceptions of program effectiveness, conservatives significantly increased their rating, as did liberals. Just over 40 percent of conservatives in the treatment group labeled the program "ineffective," compared to 55 percent in the control group. The comparison of coefficients reveals that the information effects for these two groups were statistically indistinguishable. Interestingly, the positive information effect among conservatives was significantly different from effect among moderates ($p < 0.1$). As a result, the information halved the gap in perceptions of effectiveness that existed between moderates and conservatives in the control group.

Table 3.4 evaluates hypothesis 2, which suggested that the information effects would be conditional on ideological leanings. Each table contains the treatment effects for each ideological subgroup, as well as tests of the hypothesis that the treatment effects are significantly different across groups. If conservatives responded to the information differently than liberals, we would expect to see significantly different treatment effects across the two groups. When it comes to positive information about the Food Stamp program, we might expect conservatives to be less likely to revise their beliefs because they are less supportive of welfare programs and social spending. In contrast, Democrats and liberals might feel even more enthused about the program when told of its improved efficiency.

Table 3.4: Effect of Food Stamp Error Rate Information on Outcomes, By Ideology

	<u>Effectiveness</u>		<u>Spending</u>		<u>Policy Change</u>	
	coeff.	s.e.	coeff	s.e.	coeff	s.e.
Liberal	3.07	0.09	0.39	0.08	3.53	0.12
Lib*Info	0.26	0.13 +	-0.13	0.11	0.19	0.18
Moderate	2.80	0.09	-0.11	0.07	2.96	0.11
Mod*Info	-0.03	0.13	0.11	0.10	0.32	0.17 *
Conservative	2.26	0.08	-0.63	0.06	2.29	0.10
Con*Info	0.28	0.11 *	0.05	0.09	0.09	0.15
N	571		529		559	
r-squared	0.91		0.29		0.88	

<u>Comparisons Of Treatment Effects</u>						
	<u>Effectiveness</u>		<u>Spending</u>		<u>Policy Change</u>	
	coef	s.e.	coeff	s.e.	coeff	s.e.
Con-Lib	0.02	0.18	0.18	0.15	-0.10	0.23
Con-Mod	0.31	0.17 +	-0.06	0.14	-0.23	0.22
Lib-Mod	-0.29	0.19	-0.24	0.15	-0.13	0.24

Table Note: The cells contain coefficients and standard errors from three separate regression models, one for each dependent variable: the 4-point effectiveness scale (extremely ineffective, somewhat ineffective, somewhat effective, extremely effective), the three-point spending scale (spending should be increased, kept the same, or decreased) and the 5-point policy scale (Congress should eliminate the program, make major changes, make minor changes, leave the program as-is, or expand the existing program). The results are from a saturated OLS model without a constant that features interactions for each subgroup and treatment status. The “Comparisons of Treatment Effects” tests the hypotheses that group-specific treatment effects were significantly different from one another. Asterisks highlight statistically significant differences: + = $p < .10$, * = $p < .05$, ** = $p < .01$, *** = $p < .001$.

The results for the other two outcome variables also show little evidence that conservatives responded differently than the other ideological groups. None of the groups significantly changed their spending preferences. If conservatives were biased in updating their spending preferences, so were liberals. One area where there may be evidence of conditional effects is on the policy change variable. Conservatives did not revise their attitudes in light of error rate information, while moderates did; the effect for liberals was in the correct direction but not statistically significant. Again, however, we cannot reject the null hypothesis that these treatment effects were of the same magnitude, as shown by the comparisons in the far right corner of Table 3.4.

In sum, the results for the Food Stamp study provide mixed support for hypothesis 1 and no support for hypothesis 2. On two out of three outcome variables, the typical respondent revised his or her perception of the Food Stamp program and was less likely to recommend that Congress eliminate the program or make major changes. However, the effects on these variables were quite small, especially when compared to those discussed below. And there was little evidence that conservatives responded differently than other ideological groups. In the case of

program effectiveness, conservatives actually revised their views significantly more than moderates.

It seems plausible that respondents interpreted gains in efficiency and reductions in waste as a clear indicator of a more effective use of resources, but not a justification for increased spending. This may have been especially true for conservatives, who might welcome efficiency gains but do not see them as a justification for more spending.

Study 2: Evidence of Head Start Effectiveness

In 2010, the Administration for Children and Families at the Department of Health and Human Services released the results from a rigorous, 10-year study of the Head Start program. The results were less than inspiring: researchers found that whatever benefit Head Start students received while enrolled in the program faded by the end of first grade (Administration for Children and Families 2010). In the pithy words of the former director of the Institute for Education Sciences at the Department of Education, “that is not good” (Whitehurst 2010). As the introduction alluded, Republican policymakers have cited this research to bolster arguments for reform, while the Obama Administration has developed a plan to fund Head Start centers on the basis of their performance.

The survey analyzed here randomly assigned 1/3 of the post-election sample to receive information about the evaluation (N = 277). The other two thirds of the sample received information about an investigation of enrollment fraud in Head Start (described in the next section) or a generic prompt describing the program (the pure control).

All respondents received a generic prompt describing Head Start, the number of students it served in 2009, and the amount of money spent on it. The prompt also highlighted the fact that “In recent years, policymakers have debated how well Head Start works.” Treated respondents then received the following message:

In 2010, the Department of Health and Human Services released results from a 10-year evaluation of Head Start.

Researchers found that by the end of first grade, children who had attended Head Start did not perform any better on tests of academic skills than children who had not attended the program.

The results suggest that the benefits of participating in Head Start disappear by the end of first grade.

All respondents then answered the same series of outcome questions: ratings of effectiveness, spending preferences, and preference for policy change.

Table 3.5 reveals that the Head Start information did have a large and significant impact on each of the three outcome variables. On the effectiveness question, the mean rating was significantly lower in the treatment group, and the effect (0.67 on the four-point scale) represents more than half a standard deviation. The proportion of respondents labeling the program “extremely ineffective” increased from 11 percent in the control group to 26 percent in the treatment group. At the other end of the scale, this ratio was reversed: the probability that a treated respondent would label Head Start as “extremely effective” fell from 28 percent in the control group to 11 percent in the treated group.

Table 3.5: Effect of Head Start Evaluation Information on Outcome Variables

	Effectiveness		Spending		Policy Change	
	Mean	Std. Error	Mean	Std. Error	Mean	Std. Error
Control	2.91	(2.79 to 3.03)	0.00	(-0.10 to 0.11)	3.16	(2.99 to 3.33)
Treatment	2.28	(2.15 to 2.40)	-0.40	(-0.50 to -0.29)	2.49	(2.31 to 2.66)
Effect	-0.63	(-0.81 to -0.46) **	-0.40	(-0.54 to -0.26) **	-0.67	(-0.92 to -0.43) **
N	469		451		473	

Table Note: Table contains mean ratings and the difference in means between on each outcome scale for respondents who received information about the HHS evaluation of Head Start and those who did not. Symbols correspond to statistical significance levels (+ = $p < 0.1$, * = $p < 0.05$, ** = $p < 0.01$, *** = $p < 0.001$).

The results for spending preferences exhibit the same pattern. Respondents significantly reduced their support for increasing Head Start spending. The information shifted spending preferences by 0.4 on the three-point spending scale, exactly half of a standard deviation. The proportion of respondents who said that spending should be reduced increased from 30.5 percent to 57.7 in the treatment group. The treatment also had a significant effect on respondents' desire for policy change, with half of the treated respondents recommending either eliminating the program or making major changes to it. Just over one-third of respondents in the control group felt the same way. In short, there is clear, powerful evidence for hypothesis one.

What of hypothesis two? Table 3.6 displays the effects across different ideological groups and is identical in format to Table 3.4. Note again the predictable gradient between the opinions of liberals and those of conservatives. But the main pattern that stands out is just how reactive liberals were when they were told about the results of the Head Start evaluation. They significantly shifted their perceptions of effectiveness, their spending preferences, and their desire for policy change when told of the evaluation. In each case, the information effect for liberals was as large as the effect on conservatives. None of the treatment effects are significantly different from one another, providing no evidence for hypothesis two.

For instance, while the percentage of liberals choosing one of the two "ineffective" (extremely or somewhat) responses was just 8 percent in the control group, fully 35 percent of treated respondents chose that option. On the two policy preference variables the story is similar, with large, significant declines in support for Head Start. On spending, the proportion of liberals calling for more funding dropped from 63 percent in the control group to 39 percent in the treatment group. Not a single liberal called for eliminating the program in the control group, and just 6.5 percent thought it needed major changes; in the treatment group, 9 percent argued for elimination, and nearly 40 percent total thought it needed either a major overhaul or to be eliminated. Moderates and conservatives were similarly reactive across the board, with significant information effects on each outcome variable. Moderates were particularly reactive, especially in comparison to their weaker responses in the Food Stamp experiment.

Table 3.6: Effect of Head Start Evaluation Information on Outcomes, By Ideology

	<u>Effectiveness</u>		<u>Spending</u>		<u>Policy Change</u>	
	coeff.	s.e.	coeff	s.e.	coeff	s.e.
Liberal	3.41	0.11	0.57	0.08	4.07	0.15
Lib*Info	-0.70	0.15 ***	-0.45	0.11 ***	-0.83	0.20 ***
Moderate	3.04	0.11	0.14	0.08	3.38	0.14
Mod*Info	-0.77	0.16 ***	-0.52	0.12 ***	-0.69	0.20 **
Conservative	2.48	0.09	-0.49	0.07	2.40	0.12
Con*Info	-0.57	0.13 ***	-0.31	0.09 **	-0.69	0.17 ***
N	457		440		460	
r-squared	0.9		0.38		0.86	

<u>Comparisons Of Treatment Effects</u>						
	<u>Effectiveness</u>		<u>Spending</u>		<u>Policy Change</u>	
	coef	s.e.	coeff	s.e.	coeff	s.e.
Lib-Con	-0.13	0.20	-0.14	0.15	-0.14	0.26
Con-Mod	0.19	0.21	0.21	0.15	0.00	0.26
Lib-Mod	-0.06	0.22	0.07	0.17	-0.13	0.29

Table Note: The cells contain coefficients and standard errors from three separate regression models, one for each dependent variable: the 4-point effectiveness scale (extremely ineffective, somewhat ineffective, somewhat effective, extremely effective), the three-point spending scale (spending should be increased, kept the same, or decreased) and the 5-point policy scale (Congress should eliminate the program, make major changes, make minor changes, leave the program as-is, or expand the existing program). The results are from a saturated OLS model without a constant that features interactions for each subgroup and treatment status. The “Comparisons of Treatment Effects” tests the hypotheses that group-specific treatment effects were significantly different from one another. Asterisks highlight statistically significant differences: + = p<.10, * = p<.05, ** = p<.01, *** = p<.001.

In sum, information about the ten-year Head Start evaluation provided clear support for hypothesis one, and no support for hypothesis two. Liberals were just as likely to update their beliefs as conservatives and moderates. The experiment also furnished suggestive evidence of hypothesis three.

Study 3: Evidence of Head Start Fraud

The remaining third of the post-election respondents received information about the GAO’s 2010 “secret shopper” investigation. Investigators from the GAO tried to enroll fictitious students who should not have been for the program because their incomes were too high. They found that Head Start employees at more than half of the centers investigated “fraudulently misrepresented” the fictitious parent’s income and employment information in order to fulfill eligibility requirements (Government Accountability Office 2010b). These results led the GAO to conclude that the program is vulnerable to enrollment fraud.

After receiving the same opening prompt as all other respondents, the treated group (N=274) received the following message before moving on to the outcome questions:

The Government Accountability Office recently investigated enrollment fraud in Head Start.

Investigators tried to enroll fictitious children who should not have been eligible for the program because they were not from low-income families.

At more than half of the Head Start centers that were investigated, the staff “fraudulently misrepresented” parents’ income and employment information so that children would qualify for the program.

The GAO concluded that Head Start is vulnerable to enrollment fraud.

Did learning of the GAO investigation lead respondents to re-consider their views of Head Start?

Table 3.7 displays the differences in means in the same format as above. In general, providing information about the GAO’s investigation did lead respondents to revisit their perceptions and attitudes, providing support for Hypothesis 1. On each outcome variable, treated respondents were less supportive of Head Start and more likely to call for changes to the program and its level of funding.

Table 3.7: Effect of Head Start Fraud Information on Outcome Variables

	Effectiveness Scale		Spending Scale		Policy Change Scale	
	Mean	Std. Error	Mean	Std. Error	Mean	Std. Error
Control	2.91	(2.79 to 3.03)	0.00	(-.098 to 0.11)	3.16	(2.99 to 3.33)
Treatment	2.52	(2.39 to 2.65)	-0.22	(-0.33 to -0.11)	2.68	(2.50 to 2.86)
Effect	-0.39	(-0.57 to -0.22) ***	-0.23	(-0.37 to -0.08) ***	-0.48	(-0.73 to -0.23) **

Table Note: Cells contain mean ratings on each outcome scale and the difference in means between respondents who received information about the GAO investigation of enrollment fraud and those who did not. Symbols correspond to statistical significance levels (+ = p<0.1, * = p<0.05, ** = p<0.01, *** = p<0.001).

The results are not as impressive as those for the Head Start evaluation, but they are statistically and substantively significant. Just 29 percent of respondents in the treated group rated Head Start “extremely effective” compared to 52 percent in the control group. The information about fraud led far fewer respondents to call for increased funding (43 percent in the control to 63 percent in the treatment). As was the case in experiment two, negative information about program performance changed opinions so fully that the program went from having solid majority support in the control group to having only a minority who supported increased funding. In the real world, these results would spell trouble for Head Start.

There is also no evidence of biased updating on the part of liberal respondents in these results. Table 3.8 displays the treatment effects among ideological subgroups. The effects across liberals, conservatives, and moderates are comparable, and in no case can we reject the null hypothesis that the results for each subgroup are identical.

Like the first two studies described above, informing citizens about the investigation of Head Start fraud provided no credible evidence of the motivated reasoning and ideological bias that has occupied such a prominent place in the political behavior literature.

Table 3.8: Effect of Head Start Fraud Information on Outcomes, By Ideology

	<u>Effectiveness</u>		<u>Spending</u>		<u>Policy Change</u>	
	coeff.	s.e.	coeff	s.e.	coeff	s.e.
Liberal	3.41	0.11	0.57	0.08	4.07	0.15
Lib*Info	-0.39	0.17 *	-0.23	0.13 +	-0.45	0.23 *
Moderate	3.04	0.11	0.14	0.08	3.38	0.14
Mod*Info	-0.39	0.15 *	-0.15	0.12	-0.58	0.20 **
Conservative	2.48	0.09	-0.49	0.07	2.40	0.12
Con*Info	-0.41	0.13 **	-0.17	0.10 +	-0.35	0.17 *
N	446		431		456	
r-squared	0.91		0.33		0.86	

Comparisons Of Treatment Effects

	<u>Effectiveness</u>		<u>Spending</u>		<u>Policy Change</u>	
	coef	s.e.	coeff	s.e.	coeff	s.e.
Lib-Con	-0.02	0.21	-0.06	0.16	-0.10	0.28
Con-Mod	-0.02	0.20	-0.02	0.15	0.22	0.26
Lib-Mod	0.00	0.23	-0.08	0.17	0.13	0.30

Table Note: The cells contain coefficients and standard errors from three separate regression models for each dependent variable: the 4-point effectiveness scale (extremely ineffective, somewhat ineffective, somewhat effective, extremely effective), the three-point spending scale (spending should be increased, kept the same, or decreased) and the 5-point policy scale (Congress should eliminate the program, make major changes, make minor changes, leave the program as-is, or expand the existing program). The results are from a saturated OLS model without a constant that features interactions for each subgroup and treatment status. The cells in the lower half of the table test the hypotheses that group-specific treatment effects were significantly different from one another. Asterisks highlight statistically significant differences: + = $p < .10$, * = $p < .05$, ** = $p < .01$, *** = $p < .001$.

Which Type of Negative Information is More Potent?

The data suggest that information about effectiveness was more potent than information about fraud. To look at this question more closely, I regressed each outcome variable on two dummy variables, one for each message, and then compared the magnitude of the coefficients.

On this question the data are definitive: across each outcome, informing respondents about the Head Start evaluation had a larger effect than informing them about the investigation of fraud. This asymmetric effect is generally true across almost every subgroup as well. These results suggest that evidence of program effectiveness has higher diagnostic value than information about fraud. In the language of the rational learning model, the information about the evaluation is more precise, and therefore weighed more heavily in the updating process.

What of Hypothesis Three?

Unfortunately, because the “valence” of the information (ie: positive or negative) did not vary within these studies, we cannot directly test whether the information effects in the Food Stamp experiment are larger or smaller than those in the Head Start studies. However, the results

clearly suggest that this is the case. To get a sense of how different these effects are, I look at a simple measure of effect size across the studies—the Cohen’s *d* (Cohen 1988). Cohen’s *d* is widely used in meta-analyses to aggregate and compare effect sizes. It expresses inter-group differences in terms of the standard deviation on the outcome. It is a dimensionless quantity which facilitates comparisons. Though these measures are only occasionally employed in political science (often in meta-analyses of topics; see Lau, Sigelman, and Rovner. 2007), they are useful for examining the magnitude of treatment effects in a way that is comparable across studies.

To assess hypothesis 3, I calculate Cohen’s *d* for each outcome variable in each study. Cohen (1988) suggests that a *d* between 0.2 and 0.5 is considered “small”, one between 0.5 and 0.8 is “medium”, and above 0.8 is “large”. Note that these rules of thumb are typically applied to large-scale behavioral interventions (e.g. medical trials). The effects here are likely to be somewhat less intense. Nonetheless, we can use these general rules of thumb to get a sense of how the treatment effects compare. Table 3.9 displays the Cohen’s *d* for each information effect discussed above.

Table 3.9: Effect Sizes for Each Type of Information

Treatment	Outcome Variable		
	Effectiveness	Spending	Policy Change
Food Stamps	0.19	0.07	0.18
Head Start Evaluation	0.63	0.50	0.49
Head Start Fraud	0.40	0.28	0.35

Table Note: Cells contain estimates of Cohen’s *d*, which expresses the difference in means between treated and control groups in terms of the standard deviation of the outcome variable (measured across the two groups). Cohen (1988) suggests that *d*’s between 0.2 and 0.5 are “small,” between 0.5 and 0.8 are “medium,” and above 0.8 are “large.”

The results reveal that information about the Head Start evaluation had medium effects on the three outcome variables, while effects from the Food Stamp study fall just below the “small” threshold. The information on Head Start fraud fell somewhere in between. Note that “small” does not mean “trivial” in any sense; these effects are statistically significant. But they clearly do not erase differences between ideological groups.

In the settings explored here, then, the evidence suggests that negative information about policy performance will cause a larger shift in attitudes than information with a positive valence. Clearly, information effects are likely to be a function of both the policy in question and the type of information provided. Food Stamps are a notoriously controversial program, meaning respondents may have stronger prior beliefs that are harder to change.

But the same logic applies to Head Start, only in the opposite direction. As a program with immense popularity and good will, it is somewhat surprising that citizens responded as actively as they did. These patterns dovetail with the large body of research on “negativity bias” in psychology, though more research is necessary to thoroughly explore these patterns.

Discussion

Contrary to some of the earlier findings in the “policy information” literature, I find that citizens do update their perceptions and attitudes in response to messages about how well those policies are working. Rather than ignoring new evidence and adhering to prior beliefs, the evidence suggests that partisans and ideologues reacted in predictable, even rational ways to new information. The data about Food Stamp error rates produced weaker and more uneven results, but even here conservatives revised their views of program effectiveness. On most outcome variables, liberals and conservatives tended to revise their beliefs in tandem. But these attitude changes did not lead opinions to converge. Instead, groups adjusted their assessments in the same direction and to roughly the same degree, leaving gaps between ideological groups mostly intact.

To illustrate these patterns visually, I have plotted mean ratings on the effectiveness and policy change scales for each of the treated and control groups and disaggregated the results across liberals and conservatives. Figure 3.1 displays the results.

The panels in Figure 3.1 illustrate how the attitudes of liberals and conservatives tended to move in parallel, and that these shifts were often quite similar in magnitude. In other words, there is little evidence of the biased updating or polarization that has received so much attention in the literature. But opinions did not converge, either. Though there were cases where the gap in the mean ratings of liberals and conservatives narrowed slightly in the presence of policy information, the differences of opinion were still large and highly significant in the treatment group. Providing extra information often led opposing ideologues to revise their beliefs, but these revisions did not lead them any closer to agreement.

On this point, the current study contributes to the ongoing debate in political science about what “unbiased Bayesian updating” actually looks like. As Gerber and Green (1999) argue, lasting disagreement across partisans is compatible with Bayesian updating: “The Bayesian hypothesis holds that new information moves people with different partisan affinities (but similar levels of prior information) in the same direction and to approximately the same extent” (192). In contrast, Bartels (2002) insists that unbiased Bayesian updating should lead citizens whose prior beliefs are different to “converge” when they are provided with the same information. In Bartels’ view, the fact that the opinions of Democrats and Republicans rarely converge—even on questions of fact—is evidence of biased updating. In the latest and most comprehensive salvo in this debate, Bullock (2009) uses a standard Bayesian model to show that “enduring disagreement among Bayesians is no ‘paradox’: it is the normal state of affairs” (1122). While providing new information can foster lower levels of disagreement between individuals with different priors, there is nothing in Bayes’ theorem to guarantee that opinions converge. The opinions of the most rational Bayesians may diverge in response to clear, precise new messages.

The patterns displayed in Figure 3.1 are clearly a case of the “sustained disagreement” that Bullock (2009) and Gerber and Green (1999) highlight. And as those authors contend, these gaps do not seem to be a function of biased information processing on the part of one group or the other. Even in cases where conservatives did not appear to significantly change their beliefs, the pattern among liberals was the same. While the different types of information were not equally potent, that potency did not vary across respondents in the way hypothesis two suggested.

Figure 3.1: Mean Responses and Treatment Effects, By Ideology




Figure Note: Black bars correspond to the mean rating among respondents in the control group on the effectiveness and policy change scales. Gray bars correspond to the mean rating on those scales among respondents in the treated group. White bars measure the treatment effects. Results are disaggregated across respondents who self-identified as liberals and conservatives.

Implications for Policy

From a systemic perspective, we should not lose sight of the fact that even though opposing partisans and ideologues may continue to disagree after receiving new information, that information did shift the overall distribution in one direction. Such aggregate shifts in opinion have implications for policymaking, policy insulation, and congressional behavior. For instance, while liberals were still much warmer toward Head Start than conservatives after the treatment, the information did momentarily dampen their enthusiasm and lead them to shy away from unabashed support for the program. At the same time, conservatives became even more skeptical of the program's merits. Discouraging news about a program's performance can provide policymakers with a justification for reform that may resonate with even the most ardent supporters.

This was the logic that many ascribed, rightly or wrongly, to the Bush administration's attempt to systematically evaluate federal programs using the "Program Assessment Rating Tool," or PART. As one journalist remarked in 2005, "underneath the rhetoric, PART is inherently designed to undercut progressive programs" (Grim 2005). Gilmour and Lewis (2006) hint that this strategy worked, at least in the short term: PART scores had a more important influence on budgeting for programs in agencies traditionally associated with the Democratic Party.

The results outlined here provide some micro-level evidence as to why this strategy may be less politically risky than a head-on attempt to cut spending. Attempts to reform or cut existing programs are almost always controversial, but these patterns suggest that calling for rigorous evaluation can help legislators lay the groundwork for change. Evidence of program performance that resonates with the public will make it more difficult to argue for business as usual.

The question raised by hypothesis 3 is whether these dynamics can work in reverse. The Food Stamp analysis is probably a "hard case" with which to test this phenomenon. Like welfare, Food Stamps have been perennially unpopular, so much so that the government has now changed the name to the "Supplemental Nutrition Assistance Program," or SNAP in an effort to improve its branding. Existing public opinion data seem to suggest that respondents were more likely to have stronger and more precise priors about Food Stamps than they were about Head Start (Smith 1987). As such, it is probably not that surprising that while citizens were willing to adjust their perceptions of the program's effectiveness, the information did not lead them to change their spending preferences. While Food Stamps are often part of the national dialogue on welfare, Head Start was, until recently, not really a part of that same dimension.

Implications for Research

But there is also likely a distinction here between the effects of two different measures of performance—effectiveness versus efficiency. The information about the Head Start evaluation was a definitive statement about whether the program was accomplishing its goals. In other words, to the extent voters are consequentialist and their support for policies depends on the likely or realized consequences of those policies, we would expect this information to be particularly high in diagnostic value. If a program is ineffective, why spend more money on it or leave it intact at all? Meanwhile, information about efficiency is also likely to have effects on perceptions of program effectiveness, but perhaps not on spending preferences. Efficiency measures actually say nothing about whether the program is helping the people it is designed to reach. A measure of the cost-benefit ratios of various programs gets closer to conveying both efficiency and effectiveness, and future research should examine these dynamics in more detail.

Future research should also exploit variation in the performance of a given policy on different dimensions—effectiveness, efficiency, equity, and even normative concerns. Liberals may be more concerned about equity goals than efficiency or even effectiveness, while conservatives may reverse that order. And both camps may disagree about the appropriate response to evidence of program success or failure. In the case of failure, some may believe the program needs more funding to be successful, while others will argue that policy design is the issue, not inadequate resources. A carefully designed study could assess the impact of different measures of policy performance and how those effects vary across respondents.

Concerns About External Validity

Survey experiments are always vulnerable to concerns about external validity. These analyses are no different. In particular, it is possible that the effects reveal evidence of priming rather than learning. In some sense, however, this is the dynamic of interest, and priming is likely a prerequisite to the kind of “learning” discussed here (Lenz 2009). If there were no effects, it would suggest that such information is unlikely to affect opinions in the real world.

Also, the messages were almost certainly clearer and simpler than is often the case in rough and tumble world of politics. In the real world, this information would likely be accompanied by interpretations from party elites and advocates. In politics, almost all facts are subject to debate. Indeed, shortly after the Head Start study was released, the main advocacy group began a campaign to clarify exactly what the study found. Because these sources may be more recognizable or trusted than government bureaucracies, it could be that the interpretations and not the facts are retained and learned (see Gaines, et al. 2007). While there is no way to know how these results would differ if they included party cues or counter-arguments, some of the stronger effects would be difficult to erase completely. And as Bullock (2011) contends, clear policy information can get past countervailing party cues.

Conclusion

The patterns uncovered here raise new questions about the role and importance of “policy-specific information” in shaping citizen attitudes. Contrary to earlier efforts, I find that the partisan and ideological biases that loom large in the literature are not present in this context. These results suggest that existing policies with long track records may be less vulnerable to the kind of motivated reasoning that typically surrounds the hot-button issues of the day. It is also true that voters’ priors may be weaker on issues that are only occasionally brought into polarized debates, leaving them open to learning. These processes have implications for politics and policymaking. The way citizens perceive the effectiveness and quality of government programs is a key constraint on the environment in which policymakers operate.

Chapter 4

Increasing College Knowledge: Information Effects in Social Policy Markets¹⁰

Since the 1980s, market-based ideas have played an increasingly large role in American social policy. Frustrated by the performance of traditional welfare programs, where services are provided directly by government-owned monopolies, policymakers in the United States and the United Kingdom began to experiment with consumer-centered, market-oriented approaches to social service provision. In general, these reforms use publicly funded vouchers to create “quasi-markets” where citizens are empowered to choose from among an array of public and private service providers (Le Grand 1991; Chubb and Moe 1990; Osborne and Gaebler 1992). In the U.S., these markets now stretch from cradle to grave; Americans can use vouchers to purchase child care, food assistance, housing, job training, postsecondary education, and medical services.

In theory, market-based policies work on two levels. First, endowing citizen consumers with choice will promote a better match between the service that consumers need and the service they can access. Second, in a quasi-market where consumers are free to “vote with their feet,” providers will have incentive to improve their service in order to compete for customers. Under the right circumstances these competitive pressures will help to raise the quality of service across the board.

In reality, however, the market-based model demands a lot of citizen-consumers. It assumes that sufficient numbers of citizens have the motivation and ability to collect information on the costs and benefits of various options, weigh that data carefully, and then choose the provider that best fits their preferences. If enough citizen-consumers make informed choices, then the demand-side of the market-based logic should work as intended.

But these are big ifs. A sizable literature in political science has explored whether these micro-foundations are present in the mass public. The informational assumptions have received the most scrutiny: are citizens informed about the costs and benefits of different service options, and do they use that information to “vote with their feet” (Lowery and Lyons 1989; Teske, Schneider, Mintrom, and Best 1993; Bickers and Stein 1998; Bickers, Salucci, and Stein 2006)? Survey-based tests of these assumptions have produced mixed results. Investigations of Tiebout’s (1956) model of residential sorting have often found that citizens lack information about the tax and service packages offered by local governments (Lowery and Lyons 1989; Lowery, Lyons, and DeHoog 1995).

But others have found that subsets of citizens—active choosers and those with strong preferences for particular services—are quite well informed (Teske et al. 1993; Bickers, Salucci, and Stein 2006). Even those with low levels of knowledge can use heuristics to arrive at informed choices (Bickers and Stein 1998). More recently, political scientists have turned this lens on school choice markets and come to similar conclusions. Some citizens are knowledgeable about the quality of local public schools (Teske, Fitzpatrick, and Kaplan 2006; Chingos, Henderson, and West 2012) and use that information to evaluate incumbent officials (Berry and Howell 2007). But many are not, and information levels are particularly low among low-income citizens (Howell 2006; Schneider et al. 1998).

Skepticism about citizens’ ability to satisfy these informational demands dovetails with broader discussions of citizen competence and the growing literature on informational “nudges”

¹⁰ This chapter was adapted from a paper co-authored with Mark Schneider of the American Institutes for Research. It is included here by permission of Professor Paul Pierson, chair of my dissertation, and the Graduate Division of the University of California, Berkeley.

(Thaler and Sunstein 2009). Citizens are often poorly informed on basic questions of politics and public policy (Converse 1964; Zaller 1992; Delli Carpini and Keeter 1996; Kaiser Family Foundation 2013). And even if citizens come into contact with factual information, well-established cognitive limitations and biases can conspire to lead them away from informed choices. Individuals tend to be “boundedly rational” and rely on heuristics when faced with complex choices (Simon 1997; Kahneman, Slovic, and Tversky 1982). Judgments are also subject to systematic biases, including loss aversion, overconfidence, and social identities that affect information processing (Bartels 2002; Jolls, Sunstein, and Thaler 1998). Research has uncovered similar information problems in markets for education (Buckley and Schneider 2007; Horn, Chen, and Chapman 2003), healthcare (Kaiser Family Foundation 2013; McFadden 2006), and personal finance (see Knoll 2010 for a review).

More troubling for critics of market-based policies is the fact that these micro-foundations are often weakest among less educated and lower-income citizens. Income and education are powerful predictors of citizens’ levels of political knowledge (Price and Zaller 1993; Delli Carpini 2005), their awareness of the options that are available to them under existing policies (Howell 2006; McFadden 2006), and their level of accuracy in assessing the quality of local services (Schneider, Teske, Marschall and Roch 1998). Market-based policies may stack the deck in favor of sophisticated citizens. As Steuerle (2000) warns, “competition by its very nature does imply some amount of inequality. If not all are getting the same product, then those with the least information or competence in choosing are more likely, all other things being equal, to get a lower-quality good or service” (20).

Clearly, important questions remain as to whether market-based policy ideas are built on solid micro-foundations of citizen competence. The level and malleability of citizen knowledge lie at the center of these debates. Are citizens aware of their options, and are their preferences informed? Can additional information about quality and cost inform them? And do information effects vary across citizens from different socioeconomic groups?

This study examines these questions in the context of the market for higher education. In contrast to policy debates in K-12, where market-based remedies like private school vouchers and charter schools have proven controversial, contemporary higher education policy has embraced a voucher-driven, market-based approach. However, concerns about ballooning student debt, stagnant attainment rates, and stubborn racial and socioeconomic gaps have raised questions as to whether consumers are capable of making informed choices in this market.

Using data from a 2010 survey of parents with a child in high school, this study examines how preferences toward local public colleges respond to information about those colleges. Specifically, the survey asked parents to choose between two public colleges based on basic information about them. One subset of respondents, chosen at random, received the same basic set of facts plus information about the percentage of students that graduated within six years at each campus. Because the different versions of the question were randomized, we are able to isolate the effect of the additional information and examine whether those effects vary across different socioeconomic groups.

The analysis reveals that graduation rate information had a statistically and substantively significant effect on parental preferences. It shifted the probability of choosing the college with the higher graduation rate by more than 16 percentage points. The effect is consistent even after controlling for other potential confounders, such as the respondents’ level of education and income. An exploratory analysis of the effects among subgroups finds suggestive evidence that information effects were largest among the least-educated consumers. Though further research is

needed, the findings suggest that targeted information may be especially helpful to citizen-consumers who are less familiar with the market in question.

The paper concludes by discussing the implications of these findings for the informational assumptions that underlie market-based policies. If citizen-consumers do not make effective use of information when it is provided to them, it suggests a need for alternative approaches to correct apparent market failures in social policy. If, on the other hand, targeted approaches to informing consumers can shape their preferences, then current efforts to “nudge” citizens toward better decisions hold promise.

Citizens, Information, and the Market for Education

According to Le Grand (1997), quasi-markets in social policy are built on a set of assumptions about recipients and providers that differ dramatically from those underlying the centralized, bureaucratic welfare state. The traditional model sees beneficiaries as “passive pawns” that “were supposed to be content with a universal, often fairly basic, standard of service” (Le Grand 1997, 156). Those who ran the system were seen as benevolent “knights,” motivated by a sense of duty and the public good. In contrast, proponents of quasi-markets see both beneficiaries and providers as “knaves” or “active agents in pursuit of their self-interest” (Le Grand 1997, 159). The goal of the market approach, therefore, is to “harness the knavery—or to put it less pejoratively, the self-interest—of those working in the system to the public good” (Le Grand 1997, 159).

To do so, policies typically provide recipients with a publicly funded voucher that allows them to choose a service provider from an array of approved organizations. In theory, citizens’ ability to vote with their feet should enhance the match between their preferences and the services they purchase. Collectively, these choices can also create competitive pressure that compels service providers to improve service.

But these optimistic predictions assume that citizens recognize the opportunity to choose and know enough about their options to make an informed decision. Political scientists have debated whether the typical citizen can satisfy these information assumptions and behave as market models predict. A good deal of this scholarship has focused on Tiebout’s (1956) model of residential sorting, which argued that individuals choose between localities based on the different “bundles” of taxes and services they offer. If citizens are perfectly mobile and perfectly informed, Tiebout (1956) predicted that they would “vote with their feet” and find the jurisdiction that matches their preferences.

Most studies have found that the majority of citizens fail to recognize differences in tax and service packages and rarely make exit or entry decisions based on such information (Lowery and Lyons 1989; Lowery, Lyons, and DeHoog 1995). Based on their comparative study of two metropolitan areas, Lowery and Lyons (1989) conclude that the variety of tax and service bundles “just do not seem to influence attitudes and behaviors in the gross manner suggested by the Tiebout model” (95-95). However, other studies have argued that although the average citizen may not be well informed about government service options, some groups have high levels of knowledge (Teske et al. 1993; Percy, Hawkins and Maier 1995). In their study of residential choice and public schooling in Long Island, New York, Teske et al. (1993) admit that most citizens do not hold accurate beliefs about public school spending. But they identify a set of active choosers who do “shop around” on the basis of information about school expenditures and tax rates (702). Similarly, Percy, Hawkins and Maier (1995) found that citizens who moved from one town to the other tended to choose their new residence based on their tax and service preferences. Even when citizens lack objective factual knowledge, others have found that they

can use heuristics to find their way to choices that match their preferences (Bickers and Stein 1998; see also Schneider et al. 1999).

A more recent strand of political science research has examined these information assumptions in the context of school choice markets. In their seminal argument for school choice, Chubb and Moe (1990) proposed that providing parents with the power to choose among schools would enhance public accountability compared to the traditional, democratically controlled system. When citizens have the power to exit, leaders of public schools will have incentive to respond to parents' needs in order to maintain enrollments. In contrast, parents' power to influence schools via school board elections is severely diluted because they are one constituency among many.

Again, the school choice logic is compelling in the abstract, but it quickly rubs up against reality. Parents often lack accurate information about their own child's school, let alone other options in the district. These information deficits are especially evident among low-income citizens. In their study of parents in two school districts, Schneider et al. (1998) found that respondents were rarely accurate when asked to estimate the percentage of students reading on grade level, the racial composition of the schools, and the number of disciplinary incidents. While they found higher levels of accuracy among parents who were "active choosers," they conclude "on average, low-income parents have very little accurate information about objective conditions in the schools" (Schneider et al. 1998, 769). Buckley and Schneider's (2007) study of the charter school market in Washington, DC uncovered similarly low levels of parent accuracy. And in his examination of parent knowledge in Massachusetts, Howell (2006) found that most parents were unaware when their child's school was placed on the state's list of underperforming schools. These patterns have led others to worry that school choice markets only serve to advantage the most sophisticated actors, leading to a more segregated educational system (Weiher and Tedin 2002).

Importantly, however, other studies have found that parents and citizens can make informed judgments about school quality when government policy makes such information readily available. For instance, Howell and Berry (2007) found clear evidence of retrospective voting in South Carolina school board elections after the state had implemented a school accountability system that reported on the performance of local schools. In a more direct test, Chingos et al. (2012) examined perceptions of school quality in Florida, where schools are rated on an A-F performance scale. They found that citizen perceptions of performance were highly correlated with state ratings, particularly among parents. Using a regression discontinuity design, they found significant gaps in perceptions of schools that were just above and just below the cut-off between letter grades (also see Friesen et al. (2012) on the effect of school-level "report cards" in Canada).

In sum, existing research on residential sorting and educational choice suggests that the average citizen fall short of the informational demands of market theory, that some groups are better informed, and that transparency policies can increase information levels. Are the same patterns evident in the market for postsecondary education? Political scientists have paid less attention to these informational assumptions, in part because quasi-markets have not been politically controversial in higher education. Since the mid-1960s, federal student aid programs provide vouchers that allow students to shop around, making it the longest-running voucher experiment in American public policy (Hauptman 2000). Indeed, Friedman's (1955) influential ideas about educational choice actually drew inspiration from the design of the Servicemen's

Readjustment Act of 1944 (the “GI Bill”), which provided vouchers to veterans to cover tuition payments at accredited colleges.

Despite its maturity, concerns about citizen competence in the higher education market have come to the fore in recent years. Increasing costs, stagnant completion rates, and growing student debt loads have led policymakers to raise questions as to whether individuals are equipped to make informed choices. Prospective students and their families typically lack accurate information about the cost and quality of available options, and these deficits are most pronounced among low-income and first-generation students. As Bok (2003) has argued, “[applicants] rarely possess either the time or the information to explore all the promising options available to them and usually have only a limited basis for comparing the options they do consider” (161). A study by federal researchers found that less than one-third of parents with children in eleventh and twelfth grade were able to provide an accurate estimate of college tuition (Horn, Chen, and Chapman 2003).

These information problems make a complex decision more difficult, especially for families of low socioeconomic status. Parents who have already gone to college are generally more knowledgeable about college options and the availability of financial aid than those who have not (Vargas 2004; Hossler, Schmit, and Vesper 1999). Low-income and minority parents also lack social networks that expose them to accurate information about college going (Perna and Titus 2005). Not surprisingly, research has found that low-income and minority parents are often unable to offer an estimate of the cost of one year of college, and those who do tend to overestimate the cost by a factor of two or three (Horn, Chen, and Chapman 2003; Grodsky and Jones 2004). One study of Latino parents found that on an eight-item test of “college knowledge,” low-income parents got an average of 2.6 questions correct, while high-income parents got about 4.8 questions right (Tornatzky, Cutler, and Lee 2002).

Enrollment patterns reflect these information deficits. Students tend to prefer campuses that are closer to home (Long 2004) and those that boast the best amenities (Jacob, McCall, and Stange 2013). Among low-income students, researchers have documented a pattern known as “under-matching,” or the tendency of high-achieving low-income students to enroll in a college that is less selective than they are qualified to attend (Hoxby and Avery 2012). Under-matched students are less likely to graduate than similarly qualified peers who attend more selective campuses (Bowen, Chingos, and McPherson 2009). Likewise, millions of low-income students have flocked to expensive for-profit colleges with low completion rates (Government Accountability Office 2010a).

Informing Citizen-Consumers

In light of these outcomes, policymakers and advocates have tried to provide consumers with better information about their postsecondary options. Most of this effort has focused on the “supply side”—collecting more data and packaging it in ways consumers can use. For instance, the Obama administration created a “College Scorecard” and a “Financial Aid Shopping Sheet,” both designed to provide prospective students with comparable information about different colleges.

A handful of interventions have tried to affect the “demand side” by providing additional information directly to consumers via field and survey experiments—what Thaler and Sunstein (2009) would call informational “nudges.” In North Carolina, for instance, a recent field experiment found that providing parents with a brochure on college options significantly increased their knowledge of college costs and financial aid when compared to a control group (College Board & College Foundation of North Carolina 2013). In another experiment,

Oreopolous and Dunn (2013) randomly assigned a group of Canadian high school students to watch a short video on college costs and access a financial aid calculator. They found that the modest intervention increased students' self-reported aspirations, and that the effects remained three weeks later. And in the largest experiment to date, Hoxby and Turner's (2012) "Expanding College Opportunities" (ECO) project found that sending personalized information about college options to high-achieving low-income students had positive effects on application and enrollment behavior.

However, questions remain about the efficacy of information-only strategies. For instance, another field experiment randomly assigned low-income families to receive information about financial aid opportunities as part of filing their taxes at H&R Block (Bettinger et al. 2012). Some families in the treatment group received detailed information about their likely award; for others, the tax professional actually filled out and submitted the financial aid application for the family. Information alone (the former condition) did not have a significant impact on enrollment behavior. But when the tax adviser actually submitted the financial aid application directly for the families, it had a significant effect on enrollment behavior. An evaluation of workforce training vouchers (Individual Training Accounts) found a similar dynamic; simply making information and counseling available to voucher users was less effective than linking the voucher to mandatory counseling (McConnell et al. 2006).

Hypotheses

What effect would we expect graduation rate information to have on parental preferences? Theory and existing research suggest that the additional piece of information should have some effect on preferences. In addition, respondents were presented with a pair-wise choice, which simplifies the decision considerably (Thaler and Sunstein 2009). While the paired colleges were quite similar, the graduation rates differed, providing respondents with a piece of information that distinguished one school from the other. In this setting, the graduation rates should serve as useful heuristics.

Hence, we derive the following basic hypothesis: parents who have access to graduation rate information will be more likely to choose the college with the higher graduation rate than parents who do not.

Research also suggests that different subgroups may respond to new information in different ways. In particular, a rational learning model would imply that, *ceteris paribus*, consumers who know less about their options should exhibit the largest information effects. Here, the graduation rate information is likely to be "new," leading to a more significant revision of prior beliefs than if the information was already known or inferred. In contrast, one piece of information is less likely to have a significant effect on the judgment of those with a large store of existing information. In this study, the most basic measure of familiarity with the college-going process is the respondents' level of education. Those with a bachelor's degree or above have already experienced the college-going process and are likely to be embedded in social networks where information about college options is prevalent. As Hossler et al. (1999) write, "[p]arents who have gone to college are familiar with the experience and are better equipped to explain to their children how the college system is structured, how it works, and how students can prepare for it" (26).

Hence, we derive the following hypothesis: parents with no college degree will be the most reactive to graduation rate information.

Of course, these hypotheses implicitly assume that citizen-consumers are rational processors of new information. However, research on information processing from political

science and behavioral economics provides reason to doubt this assumption (Jolls et al. 1998). In the current setting, four potential issues stand out: limited cognitive capacity, reliance on heuristics, motivated reasoning, and over-optimism.

First, individuals have finite capacity to take in and process all the information they might need to make a decision. Instead, decision-makers exhibit “bounded rationality” and do the best they can with the information they can find (“satisficing”) (Simon 1997). Having too many choices or too much information about available choices can quickly overwhelm decision-makers (Iyengar and Lepper 2000). Though the choice analyzed here was made in a relatively simple decision-making context, the presence of multiple attributes could overwhelm the impact of any one piece of information.

Second, if respondents engaged in “heuristic” rather than “systematic” processing, their cognitive shortcuts may have led them to ignore graduation rate information (Chaiken and Maheswaran 1994). There is a large literature on heuristics, some of which suggests they are useful (Lupia 1994; Lupia and McCubbins 1998), some of which argues that they can lead to predictable, systematic mistakes (Tversky and Kahneman 1974; Kuklinski and Quirk 2000). In particular, individuals tend to attach more weight to attributes that are more easily recalled from memory—perhaps because they happened more recently (availability heuristic) or because they are more familiar (familiarity heuristic) (Kahneman, Slovic, and Tversky 1982). In the current study, reliance on heuristics may lead respondents to prefer colleges with which they were familiar regardless of the additional information. However, it could also be the case that the graduation rates themselves serve as useful heuristics.

Third, respondents may engage in “motivated reasoning” when they encounter new information about colleges and discard data that does not comport with their prior beliefs (Lodge and Taber 2000). In politics, for instance, Lodge and Taber (2000) argue that individuals have an affective summary judgment about political objects—a measure of how they feel about them—and are motivated to ensure that new information jibes with their prior judgment. When messages do not fit with these existing judgments, they are marginalized, counter-argued, or ignored entirely (Redlawsk 2002). In the postsecondary market, motivated reasoning may blunt the effect of new information as attachments to alma maters or local sports teams can lead respondents to downplay information that runs counter to those feelings.¹¹

Finally, even if individuals have statistical information about the likelihood of particular outcomes, they may be overly optimistic in estimating the likelihood that such an outcome applies to them. As Jolls, Sunstein, and Thaler (1998) argue, “people may have reasonably adequate information about the risks of smoking, but this does not at all imply that they have adequate perceptions of the risks of smoking that *they themselves face*” (1542). For instance, nearly all parents (94 percent) expect their children to attend college even though fewer than 70 percent of high school graduates actually do so (Taylor et al. 2011). Such overconfidence may lead parents to assume that their child will be among the proportion of students who finish a degree. This calculation could mute the effect of graduation rate information.

Data and Methods

The data are drawn from a 2010 survey of parents’ higher education attitudes and preferences commissioned by the American Enterprise Institute for Public Policy Research and administered by Polimetrix, an online polling firm that maintains a nationally representative

¹¹ Though note also that a rational updating model could also fit with this pattern. If sophisticated consumers have most informed preferences, then they are least likely to respond to new information.

panel of respondents. The sample consists of 1,000 parents with children between the ages of 12 and 19, two hundred from each of the five largest states (California, Florida, Illinois, New York, and Texas). The survey asked parents about their aspirations for their child, their knowledge of the college application process, and the characteristics that they felt were important in a college. The American Enterprise Institute has made the de-identified dataset and documentation available to researchers upon request.

The outcome variable analyzed here asked parents to choose between two public universities in the state based on a list of attributes. Each county in a region received the same pair of colleges. The survey presented all respondents with a table that contained the following data from the National Center for Education Statistics' Integrated Postsecondary Education Data System (IPEDS) and the *Barron's Guide to Colleges*: college name, location, admissions selectivity (according to *Barron's*), whether the school was in a city, town, or rural area, enrollment, demographics (percent white, Hispanic, African American, and Asian), and cost of attendance. A randomly selected subset of respondents received one additional piece of information: the most recent six-year graduation rate reported to NCES. The graduation rates of the colleges differed by at least eight percentage points. Respondents who received the additional information are the treatment group.

From here the analysis refers to the college with the higher graduation rate as "College A" and the other as "College B" for the sake of simplicity. All respondents were then asked: "If you had to choose one of the following options, which of these schools would you recommend [your child] apply to?" The outcome variable of interest is the college preference, which can be expressed as the probability of choosing one school over the other. For ease of interpretation, the outcome is coded as a dichotomous variable where "1" equals choice of College A.

Because the graduation rate information was randomly assigned, the analyses below estimate treatment effects using simple differences in means. However, because the objects of the choice and the specific information provided varied by county, we also estimated effects using a regression model that controls for regional differences. Standard errors and statistical tests reflect the fact that respondents are clustered in regions. The results are robust to these different methods. The analyses reported here do not use the sampling weights included in the dataset.

Results

Before examining treatment effects, Table 4.1 provides some descriptive statistics for the full sample and disaggregated across the treated and control groups. Table 4.1 also provides evidence of how these characteristics vary across the treatment and control groups. To measure the balance across the two groups, the fifth column displays the results of a statistical test. The table reveals that the sample has slightly more respondents with "no college" than with a bachelor's degree or higher, but the proportion of each in the treatment and control groups are not significantly different. Hispanic parents make up about 9 percent of the sample, and the control group has a slightly higher proportion of Hispanic parents than the treatment. These differences are not statistically significant. The main difference relates to the percentage of respondents who report being "very" or "fairly well informed" about the college application process. Respondents in the treated group were more likely to report that they were informed, and the difference in proportions is significant at the 0.05 level. The results below are robust even after controlling for this imbalance.

Table 4.1: Descriptive Statistics and Balance

Variable	Overall	Control	Treated	Difference	P-Value
Male	50.2%	48.5%	51.9%	3.4%	0.28
Percent Bachelor's Degree or Higher	31.8%	32.0%	31.6%	-0.3%	0.92
Percent no college	34.9%	35.3%	34.5%	-0.8%	0.79
Median Income (categorical)	\$60-69,000	\$60-69,000	\$60-69,000	0	0.76
White	75.7%	74.2%	77.3%	3.1%	0.25
African American	7.2%	6.7%	7.7%	1.0%	0.54
Hispanic	9.1%	10.5%	7.7%	-2.7%	0.13
Percent "Very or fairly well-informed"	46.0%	41.2%	49.3%	8.1%	0.01
Difference in Distances (miles)	45.1	47.8	42.3	-5.5	0.50
Number of Respondents	1,000	507	493		

Table Note: Cells contain proportion of respondents in each category across the whole sample and in the treatment and control groups. The fifth column tests the difference in proportions across the control and treatment groups. Conventional levels of statistical significance apply ($p < 0.05$).

Table 4.1 also contains the means for a derived variable in the dataset that measures the relative proximity of each college to the respondents' zip code. Because students typically enroll in colleges that are close to home, proximity is likely to be correlated with parental preferences. Because respondents faced a pair-wise choice, the critical variable is the difference in distances, or how much closer or farther away College A was than College B. Positive numbers indicate that College A was farther away. The table indicates that College A was on average about 45 miles farther away than College B, but relative proximity is not significantly different across treatment and control groups.

Prior Levels of Information and Preferences

Existing research suggests that the level of "college knowledge" is generally low and is highly correlated with socioeconomic status. Two items asked parents how informed they felt about various aspects of the college choice process. The first item asked how well-informed they were about the college application process; the second how well-informed they were about college costs. Respondents answered both on a four-point scale: very well informed, fairly well informed, a little informed, not at all informed.

Table 4.2 displays the proportion of respondents answering "very" or "fairly well informed" to each question and disaggregates those results by respondents' level of education. Less than half of the sample reported that they were very or fairly well informed on either question, but levels of information vary predictably with educational attainment. Respondents with a graduate degree appear to be quite well-informed, but less than a quarter of those with a high school diploma or below report a similar level of knowledge. Even among those with a bachelor's degree, between one-third and one-half reported that they were only a little or not at all informed. Respondents with "some college" are difficult to place on an ordered continuum, as some may have multiple years of experience at a four-year college, others a semester at a community college.

Table 4.2: Proportion of Parents Who Felt Very or Fairly Well Informed

	Application Process	College Costs
Overall	45.2%	43.3%
N=998	(42.1 to 48.2)	(40.2 to 46.4)
H.S. Diploma or Less	22.6%	23.1%
N=349	(18.8 to 27.7)	(18.6 to 27.5)
Some College	47.8%	44.7%
N=255	(41.6 to 54.0)	(38.6 to 50.8)
Two-year Degree	39.7%	38.5%
N=78	(28.6 to 50.8)	(27.4 to 49.5.)
BA	62.6%	58.7%
N=206	(56.0 to 69.2)	(52.0 to 65.5)
Graduate Degree	79.5%	78.4%
N=112	(71.9 to 87.1)	(70.6 to 86.1)

Table Note: Cells contain proportion of respondents who reported being “very” or “fairly well” informed about the college application process and college costs as well as the 95 percent confidence interval for those proportions.

These results mirror existing research and have implications for the analysis of information effects below. If high-SES parents have a more sophisticated understanding of the college choice process than their lower-SES peers, we might expect additional information to have a smaller effect on their preferences. An updating model would suggest that one additional piece of information may not be enough to affect preferences among those with a large store of existing knowledge. If prior beliefs are strong and precise, an additional fact will be unlikely to affect a large change in preferences. The more sophisticated may also be the most likely to ignore or marginalize new information if it does not fit with their prior beliefs (Lodge and Taber 2000).

Information effects may also depend on what citizens are looking for in a service provider. Providing information on a dimension that consumers do not see as important would be unlikely to shape preferences. The survey included items that asked parents whether particular dimensions were important, somewhat important, or not important in choosing a college. Table 4.3 shows a summary of those responses, both for the overall sample and disaggregated by a dichotomous measure of educational attainment (parents with a bachelor’s degree or above and those without a BA). The results show both the percentage reporting that a given characteristic was “important” (the column labeled “important”) and the sum of those labeling it important and somewhat important (in the “total” column).

Table 4.3: Proportion of Parents Rating Each Dimension as Important

	Overall		Less than a BA		BA or Above	
	Important	Total	Important	Total	Important	Total
Quality of Academic Programs	82.9	97.9	82.1	97.8	84.6	98.1
Cost of Attendance	69.6	95.8	74.2	96.3	59.8	94.7
Availability of Financial Aid	70.7	91.4	77.0	94.0	57.2	85.9
Students Graduate on Time	46.5	87.3	50.9	89.0	37.1	83.7
Close to Home	24.6	75.1	29.0	79.3	15.1	66.0
Selectivity	20.9	68.7	18.6	64.4	25.8	78.0
Socioeconomic Diversity	25.0	60.2	27.3	61.4	20.1	57.6
Racial Diversity	21.0	52.6	21.4	51.8	20.1	54.4

Table Note: Cells contain the percentage of respondents labeling each attribute “important” and the total percentage labeling it either “important” or “somewhat important” (the column labeled “total”). Percentages omit respondents who did not answer the item.

The results reveal that parents are quite concerned about the quality of academic programs, the cost of college, and the availability of financial aid. Less than half of the parents rated whether students graduate “on time” as important. However, very few (13 percent) said that on-time graduation was “not important.” Interestingly, just one-quarter or less of parents thought that a college’s proximity to home, admissions selectivity, or racial or socioeconomic diversity were “important.” Even when you combine the response categories, just over three-quarters of parents label “close to home” as at least somewhat important, compared to 87 percent who say the same about on-time completion.

The data also suggest some differences across respondents with different levels of education. In general, parents with a BA or an advanced degree are less concerned with college costs and the availability of financial aid. When it comes to whether students graduate on-time, a slight majority of parents with less than a BA (51 percent) rated “students graduate on-time” as an important characteristic, while 37 percent of those with a BA or above said the same. At the same time, parents with less than a four-year degree were twice as likely to rate “close to home” as important (29 percent) than their more educated peers (15 percent).

These priors have implications for the analysis below. First, “students graduate on-time” clearly ranks near the middle when it comes to the characteristics that parents believe are important in choosing a college. This means that while graduation rate information may have an effect, it is unlikely to completely trump other concerns like costs and the quality of academic programs. Second, it seems likely that parents who believe on-time graduation is an important characteristic should be particularly responsive to information related to that measure. Those who do not see it as an important characteristic might downplay completion rates. If Table 4.3 is a guide, one might expect parents with less than a bachelor’s degree to be more responsive.

Information Effects

Table 4.4 displays the proportion of respondents choosing College A across the treatment and control condition, as well as the 95 percent confidence interval for the difference in those proportions. The difference is the between-subjects treatment effect. Table 4.4 also contains the treatment effects for respondents in each state. Just under half of respondents in the control group

preferred College A in the absence of graduation rate information. This suggests that respondents did not display an overwhelming preference for either college in the control group.

The results also reveal that the presence of graduation rate information had a significant effect on preferences. Specifically, 61 percent of parents in the treatment group chose College A, translating to a treatment effect of 16.7 percentage points. The change is statistically and substantively significant. But it is not so large that it suggests a complete reversal of preferences. A substantial minority of parents still chose College B even in the presence of graduation rate information.

Table 4.4: Information Effects (Differences in Means)

	Control	Treatment	Difference	
Overall	0.45	0.61	0.17	**
CI	(0.40 to 0.49)	(0.57 to 0.66)	(0.09 to 0.24)	
N=993	504	489		
California	0.49	0.60	0.11	
CI	(0.39 to 0.59)	(0.50 to 0.70)	(-0.20 to 0.42)	
N=199	94	105		
Florida	0.36	0.67	0.30	***
CI	(0.27 to .46)	(0.57 to 0.76)	(0.17 to 0.44)	
N=198	102	96		
Illinois	0.42	0.53	0.12	+
CI	(0.32 to 0.51)	(0.43 to 0.63)	(-0.02 to 0.26)	
N=200	106	94		
New York	0.53	0.67	0.14	
CI	(0.43 to 0.63)	(0.57 to 0.77)	(-0.20 to 0.49)	
N=196	102	94		
Texas	0.44	0.60	0.16	
CI	(0.34 to 0.54)	(0.50 to 0.70)	(-0.19 to 0.51)	
N=200	100	100		

Table Note: Cells contain the proportion of respondents in each group that chose College A, the difference in those proportions, and the 95 percent confidence interval for each difference. The symbols correspond to statistical significance levels (+ = $p < 0.1$, * = $p < 0.05$, ** = $p < 0.01$, *** = $p < 0.001$).

The information effects clearly vary across states, and only two of the state-specific estimates register as statistically significant because of the small sample size (~200). However, the effects are generally similar in magnitude to the overall treatment effect, with the exception of Florida, where it is considerably larger. The information effects in Illinois, New York, and Texas are all in the twelve to sixteen percentage point range, though none of them are

statistically significant at the 0.05 level (the effect in Illinois is significant at the $p < 0.1$ level). It is also worth noting that the overall effect is not driven by one particular state. Excluding respondents from Florida, where information had the largest effect, reduces the estimated effect slightly to 13 percentage points, still statistically significant. Excluding the respondents from California produces an effect of 18 percentage points.

Table 4.5: Regression Models of Treatment Effects

	Model 1	Model 2	Model 3	Model 4
Information	0.16	0.16	0.16	0.16
Std. Error	0.04	0.03	0.02	0.03
Confidence Interval	(0.08 to 0.25)*	(0.08 to 0.24)*	(0.11 to 0.2)*	(0.10 to 0.23)*
Difference in Distance		-0.10	-0.10	-0.10
Std. Error		0.03	0.03	0.03
Confidence Interval		(-0.17 to -0.03)*	(-0.17 to -0.03)*	(-0.17 to -0.03)*
BA or Above			-0.04	-0.04
Std. Error			0.05	0.04
Confidence Interval			(-0.14 to 0.07)	(-0.16 to 0.01)
Income			0.00	0.00
Std. Error			0.01	0.01
Confidence Interval			(-0.01 to 0.01)	(-0.02 to 0.02)
African American			-0.11	-0.11
Std. Error			0.06	0.07
Confidence Interval			(-0.25 to 0.04)	(-0.26 to 0.04)
Hispanic			0.02	0.02
Std. Error			0.04	0.04
Confidence Interval			(-0.08 to 0.12)	(-0.08 to 0.11)
Informed				0.01
Std. Error				0.03
Confidence Interval				(-0.06 to 0.09)
Controls for Region	Yes	Yes	Yes	Yes
N	993	993	921	921

Table Note: Cells contain estimated coefficients, standard errors, and confidence intervals from a series of linear probability Ordinary Least Squares regressions. Standard errors Symbols correspond to levels of statistical significance (* = $p < 0.05$, ** = $p < 0.01$, *** = $p < 0.001$).

Because the additional information was randomly assigned, these differences in means should produce an unbiased estimate of the treatment effect. Table 4.5 assesses the robustness of those treatment effects by running a series of linear probability regression models that control for confounders. The first column of Table 4.5 shows the results of a model that included controls for the pair of colleges the respondent received (which corresponded to the region of the state they resided in). The second and third columns add additional individual-level characteristics to this basic model (income, whether the respondent had a BA or advanced degree, whether she was Hispanic or African American, and whether she felt very or fairly well informed about the college application process).

Even after controlling for differences across college pairs, demographics, and self-reported information levels, the estimated treatment effect remains stable at about 16 percentage points. Table 4.5 also reveals the relationship (or lack thereof) between other covariates and the preferences expressed by parents. The coefficient on the relative distance measure, which corresponds to the change in probability of choosing College A per 100-mile increase, suggests that relative proximity is an important predictor of preferences. For every 100 miles further away College A was than College B, respondents were 10 percentage points less likely to choose it. Interestingly, three measures that we would expect to be highly correlated with college knowledge—education, income, and self-reported knowledge—did not surface as significant predictors of preferences. The results in Tables 4.4 and 4.5 provide strong support for the first hypothesis.

Exploratory Subgroup Analysis

The second hypothesis predicted that information effects would vary meaningfully across different groups of respondents, and that respondents with less education and lower levels of awareness would be more responsive. Here we test that hypothesis by disaggregating the treatment effects across respondents with different levels of education and self-reported familiarity with the college application process. Because the randomization was done within subgroups, these analyses are necessarily exploratory. Nonetheless, identifying patterns among subgroups can guide future research.

To assess the second hypothesis, we estimated regression models where the outcome variable was regressed on dummy variables for each of the subgroups and an interaction term between each of those dummy variables and the treatment indicator. The models are saturated and estimated without a constant for ease of interpretation. The coefficient on the dummy variable corresponds to the mean for that subgroup in the control group, while the coefficient on the interaction terms corresponds to the treatment effect on that particular subgroup. For the education analysis, respondents were placed into three groups: those with no college degree, those with some college or a two-year degree, and those with a BA or above. The information groupings are based on the four categories of the self-reported information measure: very well informed, fairly well informed, a little informed, and not at all informed.

Tables 4.6 and 4.7 display these regression results, as well as a series of statistical tests to assess whether the subgroup coefficients were significantly different from one another. The coefficients in Table 4.6 show that graduation rate information had a significant impact on the preferences of parents with no college degree and those with some college or a two-year degree. The treatment effects for these two groups are large—on the order of 21 percentage points. Meanwhile, among parents with a BA or higher, the information effect is just 7 percentage points and is not significant. In both cases, the effects on less educated respondents were significantly different from those for respondents with a bachelor's degree or above. For instance, the

treatment effect was 14 percentage points larger among those with no college experience than it was among those with a BA, and that difference is significant at the $p < 0.1$ level. The difference is similar between those with some college or a two-year degree and the most educated in the group (a 13 percentage point gap, significant at the 0.05 level). These results provide evidence for the second hypothesis.

Table 4.6: Treatment Effects By Level of Education

	Coeff	SE	Sig.
HS or Below	0.42	0.04	
HS or Below*Info	0.21	0.05	**
Some College/Two-Year	0.44	0.04	
Some College/Two-Year*Info	0.21	0.05	**
BA or above	0.49	0.04	
BA or above*Info	0.07	0.06	
N	993		
R-Squared	0.54		
Comparisons of Treatment Effects			
H.S. - BA or above	0.14	0.08	+
Some/Two-Year - BA or above	0.13	0.05	*
H.S. - Some/Two-year	0.01	0.08	

Table Note: Cells contain regression coefficients and standard errors from a saturated linear probability model featuring dummy variables for each educational group and interaction terms between those group dummies and treatment status. The rows labeled “Comparisons of Treatment Effects” tests the hypothesis that treatment effects were different across groups. Symbols correspond to levels of statistical significance (+ = $p < 0.1$, * = $p < 0.05$, ** = $p < 0.01$, *** = $p < 0.001$)

When it comes to the results for the subgroups disaggregated by self-reported information, the pattern is somewhat different, as revealed in Table 4.7. Graduation rate information had large, significant effects on the preferences of parents with the lowest reported levels of information. Receiving graduation rate information shifted the preferences of the least informed parents by about 25 percentage points. However, significant information effects are also evident among respondents who reported higher levels of information. And none of the formal tests of the coefficients can reject the null hypothesis that the coefficients are of the same magnitude.

Table 4.7: Treatment Effects By Self-Reported Level of Information

	Coeff	SE	Sig
Not at all informed	0.38	0.06	
Not at all informed*Info	0.25	0.08	*
A little informed	0.46	0.05	
A little informed*Info	0.16	0.03	***
Fairly well informed	0.45	0.05	
Fairly well informed*Info	0.12	0.05	*
Very informed	0.49	0.08	
Very informed*Info	0.17	0.05	*
N	993		
R-squared	0.55		
<u>Comparisons of Treatment Effects</u>			
Not at all informed - very informed	0.09	0.11	
Not at all informed - fairly well informed	0.13	0.09	
Not at all informed - A little informed	0.09	0.08	
A little informed - very informed	-0.01	0.07	
A little informed - fairly well informed	0.03	0.05	
Fairly well informed - very informed	-0.04	0.06	

Table Note: Cells contain regression coefficients and standard errors from a saturated linear probability model featuring dummy variables for each educational group and interaction terms between those group dummies and treatment status. The rows labeled “Comparisons of Treatment Effects” tests the hypothesis that treatment effects were different across groups. Symbols correspond to levels of statistical significance (+ = $p < 0.1$, * = $p < 0.05$, ** = $p < 0.01$, *** = $p < 0.001$)

Thus, the results provide mixed but suggestive evidence for the second hypothesis. When divided by their level of education—arguably the most reliable measure of whether an individual is familiar with the college-going process—the results provide support for hypothesis two. When respondents were categorized according to their self-reported level of information, less informed respondents do appear to be quite reactive, but the effects are not significantly different from those reporting higher levels of information.

Discussion

Critics of market-based approaches to social policy worry that citizen-consumers are ill equipped to make informed choices about their health and welfare, particularly those from disadvantaged backgrounds. In response to these warnings, researchers, policymakers, and advocacy groups have taken efforts to flood social policy markets with useful information. The movement has been particularly remarkable in education, where governments and third-party organizations now collect and publicize important data on the costs and quality of providers. The question now is whether these “supply side” efforts to improve the quantity and quality of information can affect the preferences and behavior of citizen consumers.

The data analyzed here provide suggestive evidence that under certain conditions, they can. Collectively, the results show that graduation rate information had a statistically significant

and substantive effect on the college parents preferred. The main information effect was robust to different estimation methods and was in the predicted direction in each of the states. Overall, parents responded to additional information in predictable and meaningful ways.

More interesting, perhaps, is the evidence that the graduation rate information was most potent among parents with lower levels of education. These are the very citizens that critics argue will be at a disadvantage in public goods markets. Some of the data analyzed here suggest these fears are well-founded: parents with less than a college education were the least likely to report that they were informed about the college application process or about college costs. However, consistent with a rational updating model, these under-informed parents appear to have seized on the additional information in making their “choice,” such as it was. Parents with some college education, but no four-year degree, were similarly responsive. In other words, while critics are right to raise concerns about low levels of information in equilibrium, the results here suggest that simple informational interventions may actually pay large dividends among the citizen-consumers that need information the most. Equally interesting is the fact that information failed to move the most educated parents—those who are arguably most familiar with local college options. The absence of information effects within this group could simply be a function of these parents having the largest store of existing information about their options (as the self-reported information measure suggests). While a Bayesian updating model might have predicted the findings for these subgroups, political scientists know well that deviations from this ideal are commonplace.

Future research should test these subgroup effects more conclusively by administering informational interventions to a variety of different types of consumers. Most of the existing research on information effects in postsecondary markets has focused on improving the outcomes of low-income students, but plenty of middle class students also make suboptimal choices each year.

Addressing External Validity Concerns

These results build on the growing body of work on how informational interventions can shape preferences and behaviors in various areas of policy. However, there are legitimate concerns about external validity. Rather than actually “choosing,” respondents were asked to express a preference for which campus they would rather their child apply to. Clearly, this oversimplifies the college choice process. It limits the choice set to a pair of public colleges, truncates the dimensions that prospective students might use to make decisions, and asked respondents to express a preference in a costless, low-stakes environment. It is also true that the options in the choice set were similar on most of the dimensions other than the graduation rate, which may have led respondents to weigh that attribute more heavily than they would have otherwise. In short, there are obvious reasons to question whether the effects documented here would translate to lasting changes in preferences, let alone changes in application or enrollment behavior.

While there is no way to fully address external validity concerns, there are reasons to believe that the information effects documented here are realistic. First, the fact that the information did not provoke a radical reversal of preferences actually adds to its validity. Adding graduation rates to the information set did not cause parents to completely disregard any other dimensions (such as the relative proximity to home). An exceptionally large treatment effect that swamped all other attributes would have been harder to believe.

Second, recent field experiments speak to the power of informational interventions in the postsecondary search process. In the most comprehensive study to date, Hoxby and Turner

(2012) found that sending high-achieving, low-income students personalized college guides with graduation rate, financial aid, and other information had large effects on application and enrollment behavior. Intention-to-treat estimates suggest that receiving the guide increased the probability of applying to a very selective, private college by 17 percentage points and enrolling in one of those colleges by 14 percentage points. Importantly, students who received the packet also applied to campuses with significantly higher graduation rates than their peers in the control group. The three-minute informational intervention created by Oreopolous and Dunn (2013) produced similarly large effects on students' aspirations (though not on application or enrollment behavior). Like the current study, the information effects were largest among students who were unsure about their postsecondary aspirations. These other studies suggest that the information effects documented here could translate to real world settings.

Third, the choice set, which consisted of two public colleges, may actually be realistic for many students, particularly those from disadvantaged backgrounds. Nearly two-thirds of all first-time students enrolled in four-year colleges are at public institutions. Forty percent of students at public institutions report applying to three or fewer institutions (Pryor et al. 2011). The image of the student applying to 10 or more colleges with vastly different characteristics is actually quite rare. It is also true that the campuses in each pair could both conceivably be in a student's choice set; both were in-state and had the same level of admissions selectivity.

Finally, it is important to point out that graduation rate information may actually be most useful when students are constructing their choice set rather than when they are trying to find the best match within it. In other words, the type of information provided here would be most useful during what Kahneman and Tversky (1979) call the "editing phase": choosers conduct an early analysis of their options in order to simplify the more detailed evaluation that comes later. In the editing phase, decision-makers assess options relative to a reference point and discard those options that are clearly dominated by others. Survey evidence suggests that families do engage in such editing, often eliminating schools on the basis of their cost of attendance (Taylor et al. 2011). Graduation rate information—a simple summary measure of student outcomes—could also be a particularly useful heuristic with which to eliminate "bad" choices from the batch. In short, graduation rate information need not have a direct impact on the final enrollment decision to play an important role in the decision-making process. To the extent information helps prospective students in the editing phase, it is playing an important role.

The Importance of Choice Architecture

The findings may also speak to the importance of what Thaler and Sunstein (2009) call "choice architecture." The question provided respondents with a set of comparable information, and those in the treated group were equipped with a simple measure of student outcomes. They were also presented with a pair-wise choice and asked to express a preference (the item asked "if you had to choose"). In other words, the informational intervention embodied many of the characteristics that "choice architects" aspire to: it imposed a simple structure on what is otherwise a complex process, it clarified the mapping from choice to welfare (by highlighting the probability of completing a bachelor's degree within six years), and it mimicked a required choice (Thaler and Sunstein 2009).

In other words, these findings do not show that "more information" is an effective strategy for improving higher education markets. Rather, the results suggest that information that is transmitted in a way that catches consumers' attention and packaged in a way that consumers can use can have an effect on expressed preferences. Similarly, the success of the Hoxby and Turner (2012) intervention was likely due to the fact that the information was personalized, sent

directly to students in a brightly colored envelope, and allowed for easy ranking and comparison. Supply side information efforts would be wise to consider how that information is packaged and transmitted. Even if policymakers build a user-friendly database, there is no guarantee consumers will come. And if objective third parties do not create the choice architecture for consumers, you can be sure that self-interested parties—like the colleges themselves—will create it for them, which could leave consumers vulnerable to manipulation.

Conclusion

Proponents of market-based approaches to education ignore consumers' information problems at their peril. Citizens have low levels of information about the cost and quality of schools, and that knowledge tends to be correlated with socioeconomic status. Even when citizens are diligent, the nature of experience goods and the opacity of pricing make it difficult to make informed choices. All of these issues raise serious doubts as to whether sufficient numbers of citizens are competent enough to make choices that promote their well-being and generate the kind of competitive market that proponents envision. Heated debates about school vouchers, charter schools, college admissions, and health insurance markets tend to revolve around the same fear: if the most sophisticated citizens take all of the best seats, then quasi-markets will exacerbate the very inequalities social welfare efforts are designed to erase.

Attention to the relationship between imperfect information and inequality is a consistent theme of political science and public policy research. Political scientists have voiced remarkably similar warnings about the implications of citizen competence for representation and democratic accountability. If voters are uninformed, biased in their acquisition and use of information, or vulnerable to faulty rules of thumb, electoral outcomes can lead to representational disconnects. To the extent the distribution of political knowledge is skewed toward the top of the socioeconomic spectrum—and most evidence suggests that it is (Converse 1964; Zaller 1992)—electoral and policy outcomes will likely reflect the preferences of the most sophisticated (Bartels 2009; Schlozman, Verba, and Brady 2012).

For some, the response to these patterns is simply “sometimes life isn't fair.” For others, the response is to call for a retreat from market-based approaches to any public policy problem. But another group has asked whether these information problems are amenable to policy solutions. While most of the political science research on public goods markets has focused on existing levels of citizen competence, a newer strand of work—consisting mainly of field and survey experiments—has sought to examine the malleability of political knowledge. The question is no longer whether citizens are equipped to act as savvy choosers, but under what conditions they can become equipped to do so. The answers thus far are by no means uniformly positive. Even in the lab, citizens tend to hang onto misinformation and ignore new data with impressive doggedness. But other studies have painted a more sanguine picture of the opportunities to help citizens learn more and make better choices.

This study fits broadly into this latter category. The findings suggest that information does lead citizen-consumers of all stripes to revise their judgments about college options. Less-educated parents were particularly likely to seize on new information. But the study has also unearthed important questions for future research. How do information effects play out when the choice set is broadened to more than two options? How do consumers manage trade-offs and multi-dimensional comparisons? Most importantly, under what conditions can informational “nudges” help to level the playing field among the advantaged and the disadvantaged? Answers to these questions can help policymakers improve upon existing efforts and lay the groundwork for smarter, more equitable social policies.

Chapter 5 Conclusion

This study began by highlighting the tension between optimistic predictions about the benefits of government transparency and the persistently discouraging findings on citizen competence. Proponents of open government promise that transparency will help citizens exert greater control over their government. The logic is simple: increasing the supply of information makes it easier for citizens to monitor the decisions of public officials and the quality of the public services they receive. Voters and consumers can then use this information to ensure that their choice of candidates, issue positions, or service providers rewards those who serve their interests and punishes those who do not.

But skeptics can point to a half-century of research from political science, psychology, and behavioral economics that cautions against such optimism. Citizens are not only “rationally ignorant” when it comes to politics (Downs 1957); they also exhibit deep-seated biases in the way they consume, process and assimilate new information in formulating beliefs. Partisans confronted with the same facts often come to radically different conclusions, while the beneficiaries of public services often fail to comparison shop in the way rational economic models predict. More troubling, perhaps, is the finding that the most sophisticated citizens are the most likely to receive new information, potentially deepening existing political and economic inequalities. If this literature serves as a guide, efforts to augment the supply of information will be hard-pressed to fundamentally change the level of citizen competence.

As the three preceding chapters illustrate, neither the optimistic nor the pessimistic story is entirely correct. Rather than ignorant ideologues or omniscient super-citizens, the findings show a public that can use new information to update their beliefs about politics and policy under favorable conditions. Information affected citizens in predictable ways when it was accessible and simple to understand, when it was diagnostic of the object or choice in question, and when it acted on weaker prior beliefs. In most of the cases examined here, these conditions were present, and significant information effects emerged.

Of course, the key phrase here is “under favorable conditions.” Under alternative circumstances, information effects were weaker, nonexistent or perverse. For instance, when party cues failed to predict a Senator’s position on a roll call issue, the increased accessibility of those cues in the context of an election led constituents to be less accurate in inferring issue positions. Likewise, the most educated parents did not significantly shift their preferences when confronted with graduation rate information, perhaps because the graduation rate information ran into strong, well-informed prior beliefs. Rather than proving the rational learning model wrong, these exceptions furnished important information about the conditions under which we should expect information to have an effect in predicted directions.

How do we reconcile these findings with the conflicting conclusions from existing literature on information effects? This is hardly the first study to highlight the importance of the information environment (Bartels 1993; Jerit, Barabas, and Bolsen 2006), policy-specific information (Kuklinski et al. 2001; Gilens 2001; Bullock 2011), the diagnostic value of new information (Kuklinski et al. 2001) or targeted informational “nudges” (Thaler and Sunstein 2009; Hoxby and Turner 2012). But this existing research provides less insight into why there is such a divergence between studies that find evidence of information effects and those that find ignorance, resistance, and misperception.

The problem stems, in part, from the fact that the broad range of existing studies that fall under the umbrella of “political knowledge” have actually looked for information effects using different measures in different settings and with different definitions of competence. Some cast “competence” as an aggregate score on a quiz of political knowledge while others ask whether citizens’ beliefs change in light of new information. Even among the latter category, there has been more attention paid to the presence or absence of effects than to equally important questions as to why information would provoke updating in some settings and on some issues but not others. As such, the competing studies tend to accumulate in two separate piles—information does affect or information does not affect.

The three studies examined above set out to fill some of these gaps by testing different types of information within the same study, examining whether information effects in one setting hold up under different conditions, and exploring how effects vary across different types of respondents. The goal of this concluding section is use the rational learning model laid out in the introduction to help make sense of the findings in chapters two, three, and four and those in the existing literature. The hope is to refine our understanding of why information effects are so prominent in some areas of politics and public policy but absent in others.

From here, the chapter begins with a brief summary of the results across the three studies, highlighting both the evidence of political learning and evidence of its limits. When the opportunity arose, did citizens accept new information and use it to update their beliefs? It then moves on to examine an unresolved question posed in the introduction: were the information effects documented here most prominent among the most sophisticated citizens? I conclude with a discussion of how my findings fit with and inform the divergent literature on information effects in politics and policy.

Summarizing the Evidence: Rational Learning and its Limits

Figure 5.1 provides a high-level summary of the evidence by reporting effect sizes across the three empirical chapters. The bars correspond to a common measure of effect sizes—“Cohen’s d ”—which expresses group differences in terms of the standard deviation of the outcome variable across the entire sample (Cohen 1988). Unlike indicators of statistical significance, effect sizes provide a sense of how large the effect is relative to other effects in social science. The vertical lines at 0.2, 0.5, and 0.8 correspond to Cohen’s (1988) criteria for what should be considered “small”, “medium”, and “large” effects. In the case of binary variables such as the measures of accuracy and preferences in the pair-wise choice, odds ratios were computed and then converted to Cohen’s d via the method described in Chinn (2000). Bars displayed as negative numbers represent instances where the effect went in the opposite direction of what was expected.

Overall, the findings provide considerable evidence that information caused citizens to update their beliefs in rational ways, though the effects range in size from very small to medium. There were important exceptions to the pattern of positive findings, especially in chapter two and in the Food Stamp experiment presented in chapter three. These exceptions are discussed in more detail in the next section.

On the positive side of the ledger, the findings suggest that campaigns and elections do encourage citizens to retain a crucial piece of information about their incumbent senators: their party identification. In the context of an election, respondents were 9.7 percentage points more likely to recognize the partisanship of their senator (the within subjects estimate was about 7 percentage points). The between-subjects effect corresponded to a gain in accuracy of about one-

third of a standard deviation. The effect was not enormous, but it was both statistically and substantively significant.

Figure 5.1: Effect Sizes Across Studies

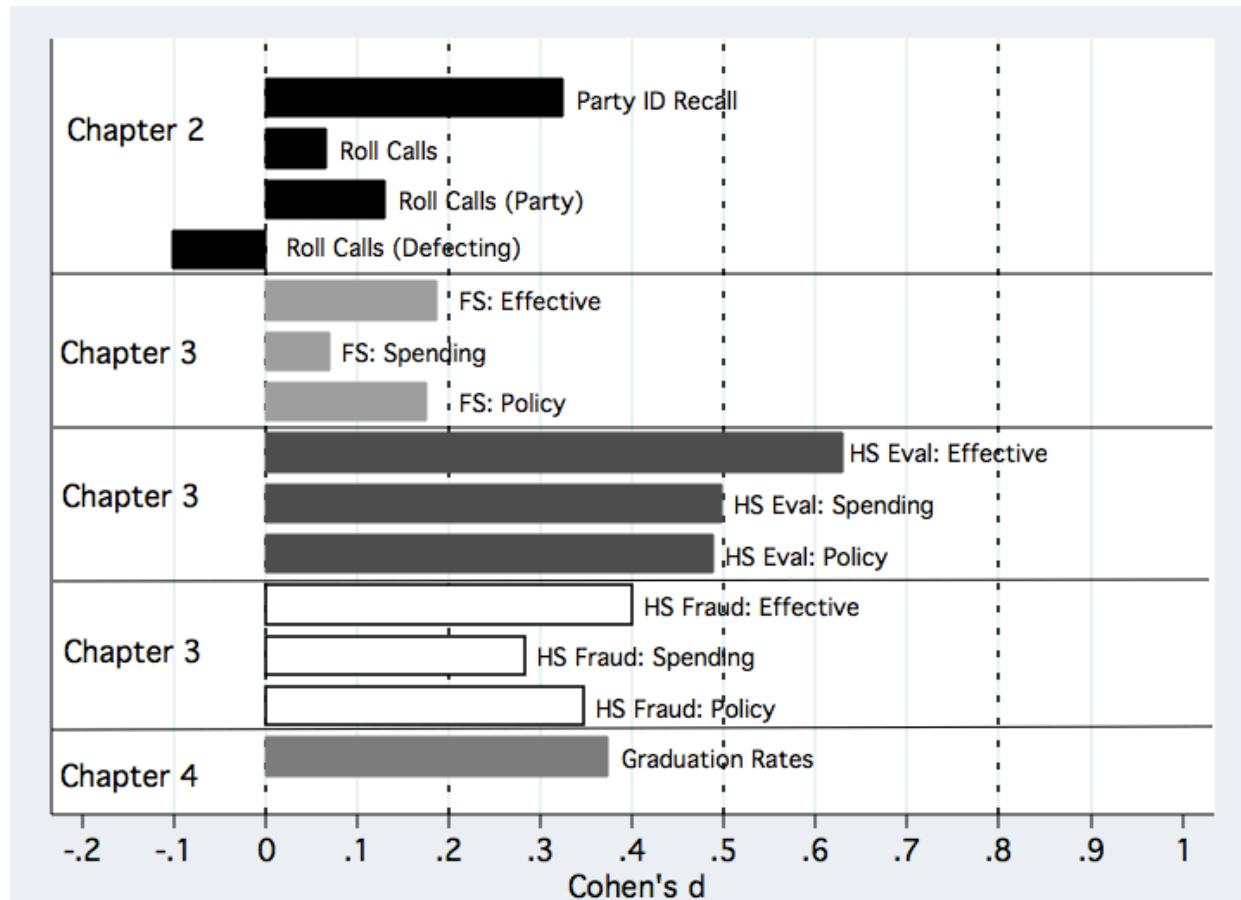


Figure Note: Bars correspond to the size of the information effects expressed relative to the standard deviation of the outcome variable. For instance, information about the Head Start study shifted spending preferences by one-half of a standard deviation. Negative effect sizes correspond to effects that went in opposite of predicted direction. FS = Food Stamps, HS = Head Start, Party = party line votes, Defecting = defecting votes.

However, there was only limited evidence that the context of an election campaign helped citizens more accurately perceive the issue positions of their senators. Study one found modest gains in overall accuracy (5.4 percentage points in the between-subjects regression model), but these gains were driven entirely by issues where the senator voted with the rest of her party. On issues where their senators defected, electoral context had a negative effect on roll call accuracy. These patterns suggest that accuracy gains were largely due to the salience of party cues rather than deeper learning about individual roll call positions. Note that that this does not mean that improved inferences on party-line votes were *entirely* a function of cue taking, only that it is difficult to distinguish between learning and cue taking on these votes.

Studies two and three provided more direct causal evidence of information effects because the surveys randomly assigned respondents to receive additional information. Those causal effects were generally significant and in the expected direction. The storyline from chapter four was straightforward: parents who received graduation rate information were

significantly more likely to choose the college with the higher completion rate. In contrast to fears that information might disproportionately benefit the most sophisticated consumers, the analysis showed that those with no college experience were the most responsive.

Chapter three presented a more complex picture of information effects. When it came to Head Start, both types of information led respondents to significantly revise their beliefs. Information about the ten-year evaluation of Head Start had particularly strong, negative effects on perceptions and preferences across all of the outcome variables. This was true for both Democrats and Republicans, and for liberals and conservatives. There was little evidence of the motivated reasoning seen in other studies. Information about the investigation of enrollment fraud in Head Start program also had significant negative effects, though they were smaller than those produced by the Head Start evaluation.

Meanwhile, positive information about the Food Stamp program's improved efficiency had more muted and uneven effects. The effects are much smaller than those in both of the Head Start studies, and there was no detectable information effect on spending preferences. However, the muted effects were not due to partisan and ideological biases. The positive effects that did surface were evident among conservatives, who had the lowest opinion of the program in the control group.

In all, there is evidence of information effects across the three empirical chapters. Typically, the information environment affected respondents' beliefs and attitudes in predictable ways, and there was little evidence of the biased updating that is so prominent in the rest of the literature. Given existing findings on the limits of information processing, the fact that information effects were found across a variety of political objects and venues was somewhat surprising.

The Limits of Rationality

However, citizens did not always respond rationally to variations in the information environment. Chapter two produced the clearest departure from a rational learning model, and it serves as a useful example of how a reliance on simple heuristics can go awry. I found that electoral context actually had a small but significant negative effect on citizen accuracy when senators took a position that was at odds with the rest of their party. This pattern is distinct from a null effect. It meant not only were respondents in the context of an election failing to learn about these votes, they were actually more likely to make mistakes.

What explains this? One interpretation is that the increased salience of party cues in an electoral environment makes voters more willing to make inferences about the incumbents' issue positions based on those cues. This strategy paid dividends when senators toed the party line, resulting in higher rates of accuracy in the electoral context. But when senators defected, respondents in an electoral context who provided an answer (excluding those who were unsure) actually did *worse* than they would have if they had simply guessed; 37 percent got the placement correct. In fairness, respondents evaluating an off-year incumbent also did poorly on these roll calls (43 percent), but the electors did worse. In other words, when party cues were made salient by the election environment, voters used them indiscriminately, failing to recognize when those cues were predictive of the senators' position and when they were not. As psychologists and some political scientists have often argued, adherence to simple heuristics can often lead to significant errors in judgment (Kuklinski and Quirk 2000).

Before moving on, it is worth putting this finding in context. First, making accurate, cue-based inferences is not as simple a task as it seems. Effective use of party cues requires that citizens learn two pieces of information: the party of their senator, and the position of the parties

on the issue. Study one showed how the former piece of information was more readily recalled in the context of an election. When it comes to the latter piece of information, Conover and Feldman (1989) show that voters learn about the positions of the parties over the course of an election cycle, while Levendusky (2009) documents how partisans have gained a clearer picture of parties' positions as elite polarization has increased (see also Lenz 2009). In other words, cue taking itself requires a fair amount of learning on the part of citizens. Even if the effect on citizen accuracy is restricted to party line votes, the increase—about six to seven percentage points—actually reflects the results of a nontrivial information-processing task.

Second, as I argued in chapter two, while cautions about overreliance on cognitive shortcuts are appropriate, it is important to explore how often such mistakes are likely to occur in the real world. If heuristics produce errors only rarely, then the overall gains in competence may well outweigh the occasional mistake. In the context of study one, for instance, while it is true that cues seem to have let voters down on a subset of roll call votes, the truth is that those votes are not all that common. Party-based inferences are far more likely to be correct than not in an era of increasing polarization and partisan sorting. A look at the party unity scores for 2006 illustrates the point: just six senators had unity scores under 75 percent (Poole 2013). Sixty-six senators had unity scores over 90 percent. Conditional on voters knowing the position of the parties on the issue, party-based inferences are much more likely to be accurate than inaccurate. It is also true that defecting issue positions are unlikely to divide the incumbent from the challenger in a general election, meaning the information is of less utility in comparing candidates.

These qualifications aside, the findings from chapter two are an important deviation from the overall pattern. But documentation of this effect is also one of chapter two's contributions. It provides political scientists with a concrete example of how the availability heuristic can lead citizens astray. Too often, the political science literature on heuristics has touted their value and downplayed findings from psychology that reveal their shortcomings. Chapter two provided a window into how their utility depends on the circumstances.

One other area is worth highlighting. Study two found suggestive evidence of a “negativity bias” in the effect of policy performance information. The idea that individuals weigh losses more heavily than equivalent gains is a cornerstone of behavioral economics (Tversky and Kahneman 1991) and psychologists have documented a “negativity bias” across many different domains (Rozin and Royzman 2001). While it is not possible to measure whether policy performance messages were equally extreme, it is striking how much more potent the negative performance information about Head Start was when compared to the positive information about Food Stamps (the effect sizes were between two and three times as large). Though negativity bias is a common finding, it can represent a deviation from a rational updating process.

However, it is worth considering two additional explanations. First, it is possible that voters pay more attention to effectiveness (whether the program accomplishes public goals) than efficiency (whether the program wastes public money). Measures of effectiveness may have more diagnostic value than measures of efficiency when it comes to policy preferences. Even if a program is efficient in handing out benefits to the right people, it may not improve the public problem it is designed to solve. The difference in effects across the two Head Start messages may also lend support to this interpretation, as evidence of enrollment fraud (partly a question of efficiency) had significantly smaller effects than evidence of program effectiveness.

It is also possible that the different results reflect differences in both the location and the precision of prior beliefs. For one thing, the information about Head Start effectiveness was

probably more surprising to respondents; fully 63 percent of the sample believed Head Start students “do better” than their peers in elementary school. Just 41 percent thought Food Stamp error rates had gotten worse over the past decade (30 percent did not know). The relative precision of prior attitudes toward each program is difficult to document. But existing research on voters’ perceptions of “deservingness” in welfare programs suggests strong biases against adult beneficiaries (Gilens 1999) and more support for programs targeting poor children (Hecklo 1997). Testing these alternative hypotheses is beyond the scope of the data, but I return to a discussion of diagnostic value and prior beliefs below.

Who Updates?

One commonly voiced concern in the literature on political knowledge and quasi-markets is the idea that the most sophisticated citizens are also the best equipped to learn from new information. Information effects might therefore widen existing gaps in knowledge, political power, and access to services, as those who know the most learn even more. This fear is especially relevant to social policy markets, where it can lead to a scenario in which the most advantaged will capture all of the spaces at the best providers, leaving low-income families to choose from among the worst ones (Steuerle 2000).

Here I examine this concern in the context of chapters two, three, and four. Recall that chapter four turned up little evidence that effects of graduation rate information were concentrated among the most educated. In fact, the effects were largest amongst the two least educated groups. After disaggregating the effects in chapters two and three according to the same three-tier education variable used in chapter four, I also find little evidence that the effects were concentrated among the best educated. Again, it was often the opposite. Figure 5.2 displays the information effects documented in chapters two and three disaggregated according to a three-category attainment variable: those with a high school diploma or below (denoted with a “1” in the figure), those with some college or a two-year degree (denoted with a “2”), and those with a bachelor’s degree or above (denoted with a “3”). It is important to note that these effects are expressed in terms of the outcome variable in question and are not standardized (as in Figure 5.1). The pattern of interest is the variation across the groups within the individual bar charts, not the comparison of effects across the charts.

The two panels of Figure 5.2 show that in most instances, the information effects were either comparable across groups or stronger among those with less education. In only one study were information effects most evident among those with a BA or above (attitudes toward the Food Stamp program). In chapter two, for instance, the context of an election campaign boosted the probability that respondents with a high school diploma would recall the partisanship of their senator by 12 percentage points; the comparable effect for those with a BA or above was 6 percentage points (this difference in effects was statistically significant at the 0.05 level). Interestingly, the negative effect on defecting votes was largest among the most educated (5 percentage points).

This last finding is an important one for the literature on heuristics. As was discussed above, political scientists have traditionally paid less attention to the notion that simple cues may be more reliable in some settings than others. But the patterns in chapter two raise interesting questions as to whether voters can tell when heuristics are more and less useful. Figure 5.2 suggests that even the best-educated voters may have trouble distinguishing when party cues are diagnostic and when they are not. They may even be more prone to these errors. I return to this topic below.

Figure 5.2: Effects By Educational Attainment

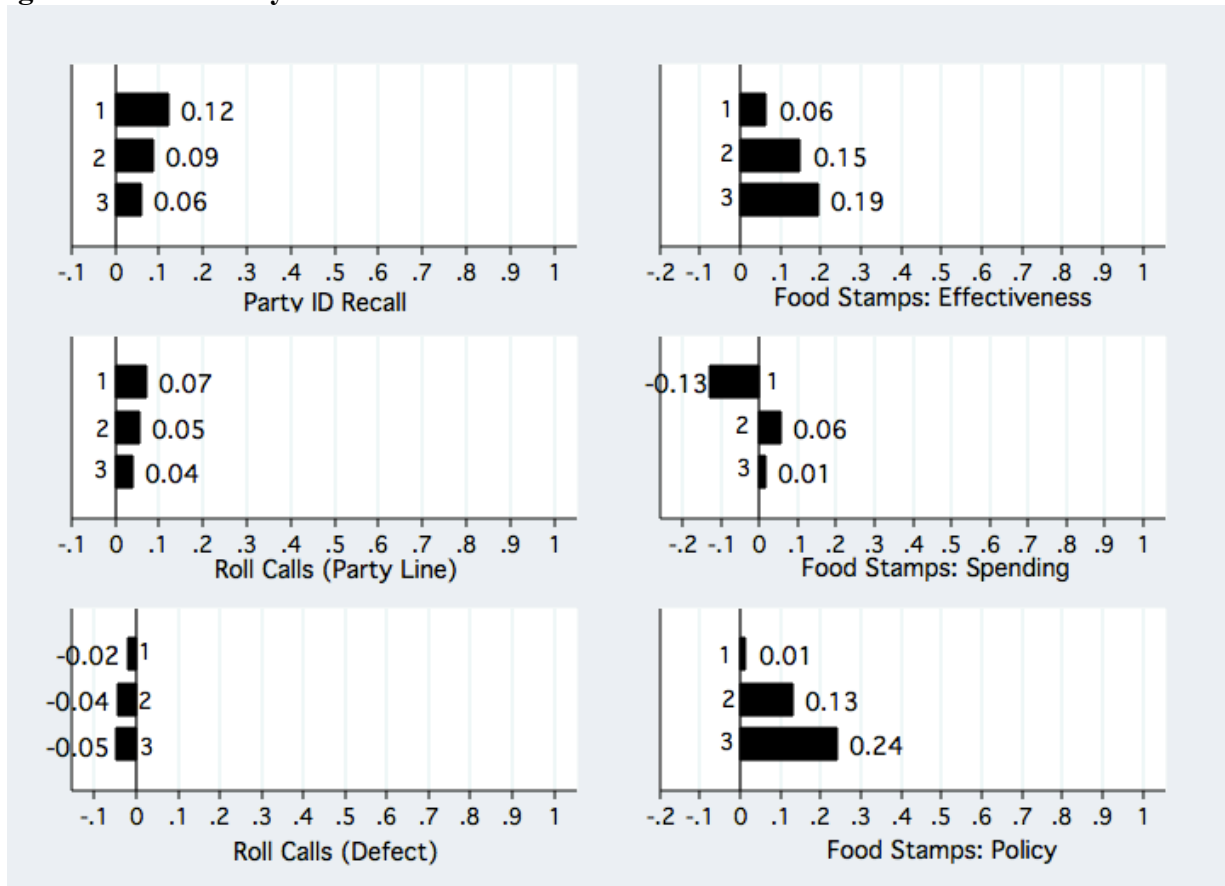


Figure Note: Bars correspond to the raw differences between informed and uninformed respondents in each educational group (1 = high school diploma or below; 2 = some college or a two-year degree; 3 = bachelor's degree or above). The differences are expressed in terms of the outcome variable in question.

Figure 5.2: Effects By Educational Attainment (Continued)

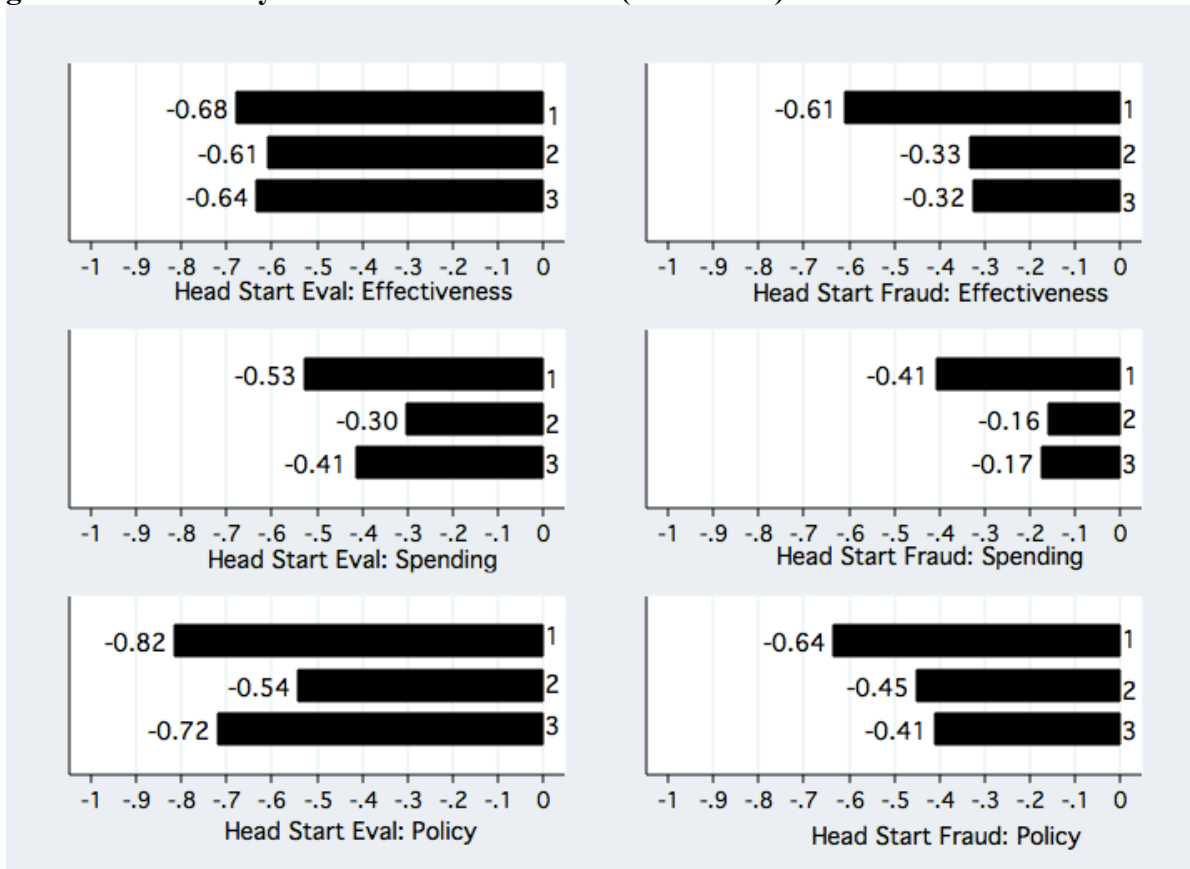


Figure Note: Bars correspond to the raw differences between informed and uninformed respondents in each educational group (1 = high school diploma or below; 2 = some college or a two-year degree; 3 = bachelor’s degree or above). The differences are expressed in terms of the outcome variable in question.

When it came to chapter three, the information about the Head Start evaluation had large and comparable effects across the three education groups. The information about enrollment fraud seemed to resonate with citizens with a high school diploma, though the differences across groups were not statistically significant. But the analysis of attitudes toward the Food Stamp program produced a slightly different pattern. Information effects were most evident among those with a BA or above and nearly non-existent among those without one. On two out of the three outcome variables, information had a significant effect on the most educated respondents, while the effect among the least educated was not significant (though these differences did not reach statistical significance).

Further analysis lent support to this pattern. I divided the sample according to their accuracy on four factual items (whether they could identify the party identification of each of their senators, their member of Congress, and their governor) and re-estimated treatment effects for two groups: those who got all four questions correct and everyone else. The most informed significantly revised their attitudes upward in response to information about Food Stamp error rates on each outcome. These effects were significantly different from respondents with fewer correct answers ($p < 0.1$). Interestingly, this did not lead beliefs to diverge. Instead, information actually caused the beliefs of the most informed and the rest of the sample to converge.

Attitudinal differences between these two groups that were significant in the control group were no longer evident in the treatment group.

In sum, I found little evidence that learning was isolated to the most sophisticated respondents. On all but three outcome variables, the effects were comparable or larger among those with a high school diploma or less.

Reconciling the Divergent Literature On Information Effects

Most of these findings fall on the positive side of the ledger. But rather than simply put them in one of the piles and call it a day, the remainder of this chapter makes an effort to explain why the literature on information effects has often come to such different conclusions. To do so, I revisit the rational learning model described in the introduction and hold it up to both the current findings and those in the existing literature. It seems that a good deal of the divergence can be explained by variation across different studies on key aspects of the rational learning model.

What is the Probability that Citizens Receive the Message?

Recall that the first step in political learning is receiving the message, and the probability of reception is some unknown function of individual characteristics and contextual factors. Chapter two argued that this probability varies systematically with the environment in which voters operate. Put more simply, when the environment makes particular pieces of information salient, voters are more likely to learn them. But the results also revealed that researchers must be careful in assuming that these opportunities are a substitute for more systematic search and analysis. In some cases, the political environment can provide cues that lead voters to hold incorrect posterior beliefs. These findings inform both the traditional scholarship on political knowledge and the alternate takes. I discuss each in turn.

Like others before, chapter two suggests that the traditional approach to measuring political knowledge may miss important evidence of competence. Typically, conclusions about the level of political knowledge come from survey items that measure general factual knowledge about governmental processes, public figures, and public policies. It is akin to a low-stakes quiz, where the number of questions answered accurately is used to sort who is well informed from who is not. Quizzing citizens in this way assumes two things: that the facts on the quiz are of equal importance to citizenship and that citizens have equal opportunity to learn them. So the name and party identification of ones' elected representative is given the same weight as knowing who the Chief Justice of the Supreme Court is, despite the fact that citizens have a regular opportunity to vote for one but not for the other. And to the extent knowledge varies, it is due to individual-level differences in education, interest, and experience.

But this is an admittedly strict view of what citizens need to know to exert control over their government. It ignores the fact that the opportunities to learn about particular topics vary dramatically across issues and over time, and that those opportunities tend to be correlated with the tasks of citizenship. Think of it this way: if the modal score on a ten question quiz of political knowledge was three, does it not matter which three questions most citizens got right? If the three correct answers were related to important imminent political decisions—which candidate to support, whether life has gotten better or worse under an incumbent, or the performance of the local public schools—that changes things. On the other hand, if citizens still cannot learn when the environment provides a clear opportunity to do so, the story may be even worse than we thought. Measures that ignore both the relative importance of different topics to the decisions at hand and the relative supply of information set an artificially high bar.

Chapter two's look at the 2006 Senate election showed how acknowledging opportunities and importance can change the resulting picture of citizen competence. Pooling the responses of voters who were evaluating a Senator up for reelection with those who were not obscured important differences in levels of citizen competence, particularly among the least educated. Take respondents at the bottom of the education distribution, typically thought of as the least prepared to participate in politics. In chapter two, those with a high school diploma or less were able to identify the partisanship of their senator 68 percent of the time in the aggregate, a rate 35 percent better than guessing. When we look only at evaluations of senators who were up for reelection, that rate increases to 76 percent, or 52 percent better than guessing. On party-line roll call votes, these same respondents were seven percentage points more accurate. Would either of these figures reverse the verdict that the less educated are not well informed? Not at all, but they would qualify it. After all, should we not be most concerned that voters who have a chance to vote are able to learn the party identification and issue positions of their incumbent senator?

But study one also laid out important limits to these optimistic conclusions. In particular, it showed that electoral environments could raise the probability of having accurate and inaccurate beliefs simultaneously. Yes, the environment provides voters with useful cues that they can use to update other beliefs. But political scientists have to be clear about where and when these shortcuts are useful and when they are not. As we saw, partisan projection was made more acute by the availability of party cues in an election, and projection can lead to representational disconnects. The fact that voters are even more prone to these errors when the stakes are highest (in the middle of an election) suggests the limits of the information environment.

In reality, citizen competence is not only a function of using heuristics, but of knowing when not to use them. On this front, we know far less than we should about whether citizens are able to recognize when cues are reliable and when they are not. The results above, disaggregated by educational attainment, suggest that the most educated are not any better at avoiding errors of projection, and may even be worse. These findings barely scratch the surface of this important question, and political scientists should take it on in future research.

Is the Message Precise? Variations in Diagnostic Value

The availability of new information is important, but so is its precision. Recall that the ratio between the precision of the message and the precision of the priors determines the influence of new information on posterior beliefs. These dimensions are difficult to measure in the real world and existing literature provides little help. But I argued that notion of “diagnostic value” offered by Kuklinski et al. (2001) could be a valuable proxy for precision: whether the message provides citizens with a clear signal about how to evaluate objects and the information necessary to do so.

I hypothesized that messages about the performance of public programs and institutions would be of high diagnostic value to citizens. I also hypothesized that negative information might be more potent than positive. These hypotheses were largely borne out: citizens seized on messages that sent clear signals about the performance of colleges in their area and public programs like Head Start. In the case of Head Start, measures of effectiveness seemed to have larger effects than measures of efficiency. These results seem largely consistent with Gilens' (2001) findings that data about the declining crime rate leads to reduced support for prison spending and Jerit's (2009) study of predictive appeals. In each case, respondents used information about policy outcomes—real or predicted—to update their beliefs.

Other studies have been less clear in acknowledging how their informational interventions might vary in terms of their diagnostic value. Take the Kuklinski et al. (2000) analysis of welfare facts and attitudes. The treated group received a set of items that conveyed facts about welfare—the percent of the budget spent on the program, the average annual payment, the percent of the population on welfare, and the proportions of welfare recipients who were African American, had a high school degree, or were on the welfare rolls for more than eight years. A second group was not given the facts but asked to estimate the answer, and a third was a pure control group. The study found no evidence of an information effect.

But this raises two questions. First, were these facts of equal diagnostic value? Second, if not, did people's misperceptions on the most precise facts run in an anti-welfare direction, such that additional information should increase support for welfare? On the first question, the authors are largely silent. But it seems plausible that some data are more diagnostic than others. For instance, information about the proportion of recipients who are on welfare rolls for eight or more years could be seen as an indicator of success or failure in helping low-income people get back on their feet. In contrast, information about the race and educational attainment of recipients seems likely to tap group identities like class and race.

With respect to the second question, Kuklinski et al. (2000) use the estimates made by the second treatment group to argue that misperceptions generally run in an anti-welfare direction. But this was only true on some of the factual items. Specifically, respondents were more likely to *underestimate* the proportion of recipients who had been on welfare for eight years and the proportion with less than a high school education. When citizens learned that their estimates of the former were lower than they should have been, it seems unlikely that this would lead them to feel more warmly toward the program. In other words, the study's "corrections" might not have all run in a pro-welfare direction, and if the negative corrections had more diagnostic value than the positive, they could have canceled out any positive effect.

Though it was beyond the scope of chapter three, some data from the Food Stamp experiment bear on this particular question. As part of that survey, a second treatment group received information about the racial breakdown of Food Stamp recipients, drawn from the U.S. Department of Agriculture's annual report on the program. How did the effect of these data compare to the effect of error rate information? The findings were somewhat mixed, but they do not suggest that racial cues outweighed performance information. Error rate information had a significantly larger effect than racial cues on people's beliefs about program effectiveness ($p < 0.1$). On the other two outcomes, the effects were not significantly different from one another. This type of comparison is too rare; future work should make predictions about differences in diagnostic value and test them.

A final point on this concept: diagnostic value is unlikely to be constant across individuals or settings. Liberals often value equity goals over efficiency goals, while fiscal conservatives may rank them in the opposite order. Likewise, as chapter four showed in the market for higher education, parents weigh the various attributes of colleges very differently. Parents of lower-socioeconomic status placed more weight on finishing college on time, and the effects reflected that prior belief. For these parents, graduation rate information was probably high in diagnostic value.

To refine our understanding of information effects, scholars should move beyond binary tests of whether or not information affects attitudes and instead explore how different messages might have different levels of diagnostic value. Chapters three and four provided a brief snapshot

of how different messages can produce very different effects, but future research should embrace such variation.

How Precise Are Prior Beliefs?

Finally, the character of respondents' prior beliefs has significant implications for how individuals will respond to informational stimuli. Both the location and precision of the prior are of interest here. First, holding the precision constant, citizens will be more reactive when they learn of a large gap between their prior beliefs and the message. But it is also true that the more precise the priors, the less reactive citizens are likely to be to new information. On the most high-profile issues or events, information may not even be "new" if citizens have already encountered it and built it into their judgments. In the opposite scenario—under- or uninformed priors—new messages may be so strong as to create a preference where there was none before.

When it comes to prior beliefs, existing research on policy information is under-specified. Most studies assume that information should affect attitudes regardless of the issue and the likely contours of prior beliefs. But these contours vary considerably across different issues and objects. At the risk of generalizing, studies of policy-specific information tend to examine issues of two types: those where individuals are likely to have strong prior beliefs and those where attitudes and preferences are probably weaker. Not surprisingly, the studies in the former category have generally failed to find evidence of information effects, while those in the latter category (including chapters three and four) have found the opposite. It may seem like an elementary point, but it is one that is too often ignored. Not all prior beliefs are created equal.

Let us start with the studies in the former category, which includes analyses of the Iraq War (Gaines et al. 2007; Nyhan and Reifler 2010), the Bush tax cuts (Nyhan and Reifler 2010), abortion (Bullock 2007), Guantanamo Bay (Bullock 2007), welfare in the late 1990s (Kuklinski, et al. 2000), and immigration in the mid-2000s (Sides and Citrin 2007). Each of these issues is fairly polarized and salient, and most are identified with a president or political party. In other words, they are issues on which partisans are likely to have strong prior beliefs and relatively high levels of knowledge.

For example, Nyhan and Reifler (2010) attempt to correct individuals' factual beliefs about the justification for the Iraq War, the effects of the Bush tax cuts, and stem cell research. To their credit, the authors do hypothesize that ideology will mediate information effects. But does a failure to find information effects actually run counter to a rational learning model? Should we expect supporters of president Bush to change their beliefs about the presence of weapons of mass destruction (WMD) or the effects of the tax cuts when provided with information to the contrary? The implicit counterfactual ignores the strong anchoring role played by conservatives' prior beliefs. Take Iraq: by 2005 (the time of the survey), the debate about WMD was already two years old, Iraq was constantly in the news, and opinions on the war had polarized. Conservatives had likely already heard the findings on WMD as well as counterarguments from the president himself. Not only were prior beliefs strong, the information itself was probably not "new" in any real sense. The rational learning model would predict minimal information effects in this scenario: priors were strong and informed (however misguided).¹²

¹² Likewise, Kuklinski et al. (2000) describe how citizens with the most dramatic misperceptions also expressed the highest level of confidence in their estimates. To the extent this confidence is a proxy for subjective precision, it follows that the new information may have had the lowest weight relative to prior beliefs among those respondents who should have been most surprised by it.

Contrast these prior beliefs with those that are likely to exist on issues like foreign aid (Gilens 2001), prison spending (Gilens 2001), education policy (Howell and West 2009), Social Security (Jerit 2009), and social policies like Head Start. These issues are only occasionally on the national agenda and rarely drawn into the partisan fray. On these issues, citizens are simply less likely to receive repeated messages of any kind, let alone those that clarify and reinforce the stances of the two parties. The same is true of consumer preferences on things like schools, colleges, or health care providers. These experience goods are difficult to evaluate from the outside, making it hard for even the savviest consumer to be well informed. Instead, preferences often reflect word of mouth, proximity to home, and sports allegiances. The results in chapters three and four suggest that on attitudes and preferences like these, prior beliefs are quite malleable.

Put simply, expectations about the likelihood of updating would clearly differ between an issue like the Iraq War and one like prison spending. Does this mean that the latter category of issues is less important to the study of American politics? On the contrary. Limiting our lens to the most polarized issues of the day would be a mistake for two reasons. First, today's elected officials inherit existing policies and spending patterns and must work within those confines. In an era of budget deficits, new policies and spending commitments often require reforming or cutting existing ones, but this entails political risk. To lay the groundwork for change, politicians strategically use messages about the limitations of existing policy. Arguments about Social Security solvency and privatization, Medicare reform, and the lackluster results of Head Start are all examples. Whether these reform efforts will have public support will depend, in part, on whether these arguments resonate with voters.

Second, issues that are not polarized today are likely to evolve over time. As partisan elites become increasingly polarized across multiple issue dimensions—what Layman and Carsey (2002) call “conflict extension”—partisans in the mass public learn about these conflicts and often change their attitudes to bring them in line with the party. Therefore, even if issues are not currently part of these partisan conflicts, they can be drawn into them as elites take sides and voters sort into opposing camps. Once the conflict is extended and partisans sort, these prior beliefs are likely to become stable and resistant to change. But the “action,” in terms of political learning, occurred prior to this equilibrium. As such, limiting our attention to the most polarized and salient issues of the day can lead us to ignore some of the most important processes of information-driven attitude change—those that lead to sorting and polarization in the first place.

We know that partisans have increasingly sorted according to their partisanship (Levendusky 2009). But questions about the actual mechanisms that drive sorting are largely unresolved. Are cues about party positions enough to drive attitudinal change? Or do partisan elites also use policy information to convince the rank and file to change their minds? Are some of these arguments more effective than others? These are areas for future research.

Zooming Back Out: Revisiting Transparency and Citizen Competence

Collectively, the chapters in this study suggest that efforts to increase transparency could help citizens to make better choices about politics and policy, but such benefits are not guaranteed. Simply increasing the supply of information is unlikely to pay dividends. Most Americans have limited time and motivation when it comes to learning about the political choices they face or the options they have under different social policies. But they have also proven capable of learning about these things when the environment actively provides an opportunity to do so. Whether governments can create and structure such opportunities so as to maximize the likelihood of informed choice is an outstanding question. And while that question

has become a major preoccupation of public policy researchers and behavioral economists (the nudge brigade) it has not yet attracted as many political scientists.

Instead, they have spent considerable time documenting the ways in which voters fall short of the ideal model of citizenship laid out by political theorists. Studies that start by comparing what citizens actually know about politics to a preconceived notion of what an ideal citizen should know are destined to disappoint. These expectations simply do not jibe with well-known limits of human cognition and information processing. In many ways, these studies resemble analogous ones from economics and psychology that document just how far human judgment deviates from the rational actor model.

But scholarship in economics has since moved on from these comparisons to ask whether and under what conditions individuals can learn to make informed decisions in spite of their well-documented shortcomings. This lens shifts the focus from individual deficiencies to the structures, institutions, and interventions that enable learning. At the core of this movement is a sense that policies around health and well-being could work better than they do, and that governments have both the capacity and the responsibility to make it happen. Whether or not the discipline realizes it, some political science scholarship has inched the field toward a similar perspective: the research on deliberative polling, get-out-the-vote field experiments, and even small-scale experimental surveys have provided key insights to political campaigns and party organizations (Issenberg 2012). This research has proactively sought to identify or create decision-making environments that help citizens learn from new information and participate meaningfully in the political process.

But there are opportunities for a greater contribution. Because of its long fascination with political learning and the political economies of different public policies, political science has much to offer nascent efforts to improve transparency. My hope is that studies like this one help to advance both the political science literature and policymakers' efforts to identify the opportunities for and obstacles to cultivating a more informed citizenry.

References

- Achen, Christopher H. 1978. "Measuring Representation." *American Journal of Political Science* 22 (August): 475-510.
- Administration for Children and Families. 2010. *Head Start Impact Study: Final Report*. (Washington, DC: Department of Health and Human Services).
- Althaus, Scott L. 1998. "Information Effects in Collective Preferences." *American Political Science Review* 92 (September): 545-58.
- Alvarez, R. Michael and Paul Gronke. 1996. "Constituents and Legislators: Learning About the Persian Gulf War Resolution." *Legislative Studies Quarterly* 21 (February): 105-127.
- Ansolabehere, Stephen. 2007. The 2006 Cooperative Congressional Election Study Common Content. Dataset.
- Ansolabehere, Stephen and Shanto Iyengar. 1995. *Going Negative: How Attack Ads Shrink and Polarize the Electorate*. New York: Free Press.
- Ansolabehere, Stephen and Philip Jones. 2010. "Constituents' Responses to Congressional Roll-Call Voting." *American Journal of Political Science* 54 (July): 583-597.
- Ansolabehere, Stephen, James M. Snyder Jr., and Charles Stewart III. 2001. "Candidate Positioning in US House Elections." *American Journal of Political Science* 45 (January): 136-159.
- Arnold, R. Douglas. 1992. *The Logic of Congressional Action*. Princeton, NJ: Princeton University Press.
- Bafumi, Joseph and Michael Herron. 2010. "Leapfrog Representation and Extremism: A Study of American Voters and their Members in Congress." *American Political Science Review* 104 (August): 519-542.
- Barabas, Jason, and Jennifer Jerit. 2009. "Estimating the Causal Effects of Media Coverage on Policy-Specific Knowledge." *American Journal of Political Science* 53 (January): 73-89.
- Baron, Jon and Ron Haskins. 2011. "Building the Connection between Policy and Evidence: The Obama Evidence-based Initiatives." Washington, DC: Brookings Institution Social Genome Project. No. 19.
- Bartels, Larry M. 1993. "Messages Received: The Political Impact of Media Exposure." *American Political Science Review* 87 (June): 267-285.
- Bartels, Larry M. 1996. "Uninformed Votes: Information Effects in Presidential Elections." *American Journal of Political Science* 40 (February): 194-230.
- Bartels, Larry M. 2002. "Beyond the Running Tally: Partisan Bias in Political Perceptions." *Political Behavior* 24 (June): 117-150.

- Bartels, Larry M. 2009. *Unequal Democracy: The Political Economy of the New Gilded Age*. Princeton, NJ: Princeton University Press.
- Berry, Christopher R. and William G. Howell. 2007. "Accountability and Local Elections: Rethinking Retrospective Voting." *Journal of Politics* 69 (August): 844–858.
- Bettinger, Eric P., Bridget T. Long, Philip Oreopoulos, and Lisa Sanbonmatsu. 2012. "The Role of Application Assistance and Information in College Decisions: Results from the H&R Block Fafsa Experiment." *The Quarterly Journal of Economics* 127 (April): 1205-1242.
- Bickers, Kenneth N., Lapo Salucci, and Robert M. Stein. 2006. "Assessing the Microfoundations of the Tiebout Model." *Urban Affairs Review* 42 (September): 57-80.
- Bickers, Kenneth N. and Robert M. Stein. 1998. "The Micro Foundations of the Tiebout Model." *Urban Affairs Review* 34 (September): 76-93.
- Bok, Derek. 2003. *Universities in the Marketplace: The Commercialization of Higher Education*. Princeton, NJ: Princeton University Press.
- Bornstein, David. 2012. "The Dawn of the Evidence-based Budget." *New York Times* Opinionator [Blog]. May 30.
- Bowen, William, Matt Chingos, and Michael McPherson. 2009. *Crossing the Finish Line: Completing College at America's Public Universities*. Princeton, NJ: Princeton University Press.
- Brass, Clinton T., Blas Nunez-Neto, and Erin D. Williams. 2006. "Congress and Program Evaluation: An Overview of Randomized Control Trials (RCT's) and Related Issues." Washington, DC: Congressional Research Service.
- Buckley, Jack and Mark Schneider. 2007. *Charter Schools: Hope or Hype?* Princeton, NJ: Princeton University Press.
- Bullock, John G. 2007. *Experiments on partisanship and public opinion: Party cues, false beliefs, and Bayesian updating* (Doctoral dissertation, Stanford University).
- Bullock, John G. 2009. "Partisan Bias and the Bayesian Ideal in the Study of Public Opinion." *Journal of Politics* 71 (July): 1109-1124.
- Bullock, John G. 2011. "Elite Influence on Public Opinion in an Informed Electorate." *American Political Science Review* 105 (August): 496-515.
- Campbell, Angus, Philip E. Converse, Warren E. Miller, and E. Donald Stokes. 1960. *The American Voter*. Chicago: University of Chicago Press.
- Canes-Wrone, Brandice, David W. Brady, and John F. Cogan. 2002. "Out of Step, Out of Office: Electoral Accountability and House Members' Voting." *American Political Science Review* 96 (March): 127-140.

- Carsey, Thomas M. and Geoffrey C. Layman. 2006. "Changing Sides or Changing Minds? Party Identification and Policy Preferences in the American Electorate." *American Journal of Political Science* 50 (April): 464-477.
- Chaiken, Shelly. 1980. "Heuristic Versus Systematic Information Processing and the Use of Source Versus Message Cues in Persuasion." *Journal of Personality and Social Psychology* 39 (November): 752-766.
- Chaiken, Shelly, and Durairaj Maheswaran. 1994. "Heuristic Processing Can Bias Systematic Processing: Effects of Source Credibility, Argument Ambiguity, and Task Importance on Attitude Judgment." *Journal of Personality and Social Psychology* 66 (March): 460-473.
- Chingos, Matthew M., Michael Henderson, and Martin R. West. 2012. "Citizen Perceptions of Government Service Quality: Evidence from Public Schools." *Quarterly Journal of Political Science* 7 (October): 411-445.
- Chinn, Susan. 2000. A simple method for converting an odds ratio to effect size for use in meta-analysis. *Statistics in Medicine* 19 (November): 3127-3131.
- Chubb, John E. and Terry M. Moe. 1990. *Politics, Markets, and America's Schools*. Washington, DC: Brookings Institution Press.
- Clark, Robert L., Jennifer A. Maki, and Melinda S. Morrill. 2013. "Can Simple Informational Nudges Increase Employee Participation in a 401 (k) Plan?" *Southern Economic Journal*, forthcoming.
- Cohen, Geoffrey L. 2003. "Party over Policy: The Dominating Impact of Group Influence on Political Beliefs." *Journal of Personality and Social Psychology* 85 (November): 808-822.
- Cohen, Jack. 1988. *Statistical Power Analysis for the Behavioral Sciences*. (New York: Lawrence Erlbaum).
- College Board Advocacy and Policy Center and the College Foundation of North Carolina. 2013. "College is Affordable: a Pilot Study." Washington, DC: The College Board Advocacy and Policy Center.
- Conover, Pamela J., and Stanley Feldman. 1989. "Candidate perception in an ambiguous world: Campaigns, cues, and inference processes." *American Journal of Political Science* 33 (November): 912-940.
- Converse, Philip E. 1964. "The Nature of Belief Systems in Mass Publics." In *Ideology and Discontent*, ed. David E. Apter. New York: Free Press, 206-261.
- Converse, Philip E. 2000. "Assessing the Capacity of Mass Electorates." *Annual Review of Political Science* 3 (June): 331-353.
- Delli Carpini, Michael. 2005. "An Overview of the State of Citizens' Knowledge About Politics." Annenberg School for Communication Department Papers. Philadelphia: University of Pennsylvania.

- Delli Carpini, Michael and Scott Keeter. 1996. *What Americans Know About Politics and Why It Matters*. New Haven, CT: Yale University Press.
- Downs, Anthony. 1957. "An Economic Theory of Political Action in a Democracy." *The Journal of Political Economy* 65 (April): 135-150.
- Erikson, R.S. 1978. "Constituency Opinion and Congressional Behavior: A Reexamination of the Miller-Stokes Representation Data." *American Journal of Political Science* 22 (August): 511-535.
- Erikson, Robert S., Michael B. MacKuen, and James A. Stimson. 2002. *The Macro Polity*. Cambridge: Cambridge University Press.
- Fenno, Richard F. 1978. *Homestyle: House Members in Their Districts*. New York: Little, Brown.
- Ferejohn, John. 1999. "Accountability and Authority: Toward a Theory of Political Accountability." In *Democracy, Accountability and Representation*, eds. Adam Przeworski, Susan C. Stokes, and Bernard Manin. New York: Cambridge University Press, 131-153.
- Fiorina, Morris P. 1981. *Retrospective Voting in American Elections*. New Haven: Yale University Press.
- Fiorina, Morris P. 1978. "Economic Retrospective Voting in American National Elections: A Micro-Analysis." *American Journal of Political Science* 22 (May): 426-433.
- Franklin, Charles H. 1991. "Eschewing Obfuscation? Campaigns and the Perception of U.S. Senate Incumbents." *American Political Science Review* 85 (December): 1193-1214.
- Friedman, Milton. 1955. "The Role of Government in Education." In *Economics and the Public Interest*, ed. Robert A. Solo. New Brunswick, NJ: Rutgers University Press, 123-144.
- Friesen, Jane, Mohsen Javdani, Justin Smith, and Simon Woodcock. 2012. "How Do School 'Report Cards' Affect School Choice Decisions?" *Canadian Journal of Economics* 45 (May): 784-807.
- Gaines, Brian J., James H. Kuklinski, Paul J. Quirk, Buddy Peyton, and Jay Verkuilen. 2007. "Same facts, different interpretations: Partisan motivation and opinion on Iraq." *Journal of Politics* 69 (November): 957-974.
- Gerber, Alan and Donald P. Green. 1999. "Misperceptions about Perceptual Bias." *Annual Review of Political Science* 2 (June): 189-210.
- Gerber, Alan S., Donald P. Green, and Christopher W. Larimer. 2008. "Social Pressure and Voter Turnout: Evidence from a Large Scale Field Experiment." *American Political Science Review* 102 (February): 33-48.
- Gerber, Alan S. and Gregory A. Huber. 2010. "Partisanship, Political Control, and Economic Assessments." *American Journal of Political Science* 54 (February): 153-173.

- Gilens, Martin. 1996. "'Race Coding' and White Opposition to Welfare." *American Political Science Review* 60 (September): 593-604.
- Gilens, Martin. 1999. *Why Americans Hate Welfare: Race, Media, and the Politics of Antipoverty Policy*. Chicago: University of Chicago Press.
- Gilens, Martin. 2001. "Political Ignorance and Collective Policy Preferences." *American Political Science Review* 95 (June): 379-396.
- Gilmour, John B. and David E. Lewis. 2006. "Does performance budgeting work? An Examination of the Office of Management and Budget's PART Scores." *Public Administration Review* 66 (September-October): 742-752.
- Government Accountability Office. 2010a. "For-Profit Colleges: Undercover Testing Finds Colleges Encouraged Fraud and Engaged in Deceptive and Questionable Marketing Practices." Washington, DC: General Accountability Office, GAO-10- 948T.
- Government Accountability Office. 2010b. "Head Start: Undercover Testing Finds Fraud and Abuse at Selected Head Start Centers." May 18. Washington, DC: Government Accountability Office.
- Grim, Ryan. 2005. "Part-Ing Shots: The Bush Administration's Curiously One-Sided Good-government Tool." *The American Prospect*, December 20.
- Grodsky, Eric and Melanie T. Jones. 2004. "Real and Imagined Barriers to College Entry: Perceptions of Cost." *Social Science Research* 36 (June): 745-766.
- Hauptman, Arthur. 2000. "Vouchers and American Higher Education." In *Vouchers and the Provision of Public Services*, eds. C. Eugene Steuerle, Van Doorn Ooms, George Peterson, & Robert D. Reischauer. Washington, DC: Brookings Institution Press, 336-367.
- Heclo, Hugh. 1997. "Values Underpinning Poverty Programs for Children." *The Future of Children* 7 (Summer/Fall): 141-148.
- Hirschman, Albert O. 1970. *Exit, Voice, and Loyalty: Responses to Decline in Firms, Organizations, and States*. Cambridge: Harvard University Press.
- Horn, Laura J., Xianglei Chen, and Chris Chapman. 2003. "Getting Ready to Pay for College: What Students and Their Parents Know About the Cost of College Tuition and What They Are Doing to Find Out." Washington, D.C.: National Center for Education Statistics (Report No. 2003-030).
- Hossler, Don, Jack Schmit, and Nick Vesper. 1999. *Going to College. How Social, Economic, and Educational Factors Influence the Decisions Students Make*. Baltimore: Johns Hopkins University Press.
- Howell, William. 2006. "Switching Schools? A Closer Look at Parents' Initial Interest in and Knowledge About the Choice Provisions of No Child Left Behind." *Peabody Journal of Education*, 81 (1): 140-179.

- Howell, William and West, Martin. 2009. "Educating the Public." *Education Next* 9 (Summer): 41-47.
- Hoxby, Caroline M. and Christopher Avery. 2012. "The Missing" One-Offs": The Hidden Supply of High-Achieving, Low Income Students." National Bureau of Economic Research, Working Paper No. 18586.
- Hoxby, Caroline M. and Sarah Turner. 2012. "Expanding College Opportunities for Low-Income, High-Achieving Students." National Bureau of Economic Research Working Paper, No. 18728.
- Hutchings, Vincent. 2003. *Public Opinion and Democratic Accountability: How Citizens Learn about Politics*. Princeton, N.J.: Princeton University Press.
- Issenberg, Sasha. 2012. *The Victory Lab: The Secret Science of Winning Campaigns*. New York: Crown.
- Ito, Tiffany J., Jeff T. Larsen, N. Kyle Smith, and John T. Cacioppo. 1998. "Negative Information Weighs More Heavily on the Brain: The Negativity Bias in Evaluative Categorizations." *Journal of Personality and Social Psychology* 75 (October): 887-900.
- Iyengar, Sheena S. and Mark R. Lepper. 2000. "When Choice is Demotivating: Can One Desire Too Much of a Good Thing?" *Journal of personality and social psychology* 79 (December): 995-1005.
- Jacob, Brian, Brian McCall, and Kevin M. Stange. 2013. "College as Country Club: Do Colleges Cater to Students' Preferences for Consumption?" National Bureau of Economic Research Working Paper, No. 18745.
- Jerit, Jennifer. 2009. "How Predictive Appeals Affect Policy Opinions." *American Journal of Political Science* 53 (March): 411-426.
- Jerit, Jennifer and Jason Barabas. 2006. "Bankrupt Rhetoric: How Misleading Information Affects Knowledge About Social Security." *Public Opinion Quarterly* 70 (Fall): 278-303.
- Jerit, Jennifer, Jason Barabas, and Toby Bolsen. 2006. "Citizens, Knowledge, and the Information Environment." *American Journal of Political Science* 50 (April): 266-282.
- Jolls, Christine, Cass R. Sunstein, and Richard Thaler. 1998. A Behavioral Approach to Law and Economics." *Stanford Law Review*: 1471-1550.
- Kahn, Kim F. and Patrick J. Kenney. 1997. "A Model of Candidate Evaluations in Senate Elections: The Impact of Campaign Intensity." *Journal of Politics* 59 (November): 1173-1205.
- Kahneman, Daniel and Alan B. Krueger. 2006. "Developments in the measurement of subjective well-being." *The Journal of Economic Perspectives* 20 (Winter): 3-24.
- Kahneman, Daniel, Paul Slovic, and Amos Tversky (eds.). 1982. *Judgment Under Uncertainty: Heuristics and Biases*. London: Cambridge University Press.

- Kahneman, Daniel and Amos Tversky. 1979. "Prospect Theory: An Analysis of Decision Making Under Risk." *Econometrica: Journal of the Econometric Society* 47 (March): 263-291.
- Kaiser Family Foundation. 2013. "Kaiser Health Tracking Poll March 2013." Menlo Park, CA: Kaiser Family Foundation.
- Key, Valdimer O. 1966. *The Responsible Electorate: Rationality in Presidential Voting*. Cambridge, MA: Belknap Press.
- Kingdon, John. 1989. *Congressmen's Voting Decisions, 3rd Edition*. Ann Arbor, MI: University of Michigan Press.
- Kling, Jeffrey R., Sendhil Mullainathan, Eldar Shafir, Lee C. Vermeulen, and Marian V. Wrobel. 2012. "Comparison friction: Experimental evidence from Medicare drug plans." *The Quarterly Journal of Economics* 127 (1): 199-235.
- Knoll, Melissa AZ. 2010. "The Role of Behavioral Economics and Behavioral Decision Making in Americans' Retirement Decisions." *Social Security Bulletin* 70 (4): 1-23.
- Krosnick, John. 1990. "Government Policy and Citizen Passion: A Study of Issue Publics in Contemporary America." *Political Behavior* 12 (March): 59-92.
- Kuklinski, James H. and Paul J. Quirk. 2000. "Reconsidering the rational public: Cognition, heuristics, and mass opinion." In *Elements of reason: Cognition, choice, and the bounds of rationality*, eds. Arthur Lupia, Matthew D. McCubbins, and Samuel L. Popkin. Cambridge: Cambridge University Press, 153-82.
- Kuklinski, James H., Paul J. Quirk, Jennifer Jerit, David Schweider, and Robert F. Rich. 2000. "Misinformation and the Currency of Democratic Citizenship." *Journal of Politics* 62 (August): 790- 816.
- Kuklinski, James H., Paul J. Quirk, Jennifer Jerit, David Schweider, and Robert F. Rich. (2001). "The Political Environment and Citizen Competence." *American Journal of Political Science* 45 (April): 410-424.
- Lau, Richard R. 1985. "Two Explanations for Negativity Effects in Political Behavior." *American Journal of Political Science* 29 (February): 119-138.
- Lau, Richard and David P. Redlawsk. 2007. *How Voters Decide: Information Processing During Election Campaigns*. Cambridge: Cambridge University Press.
- Lau, Richard R., Lee Sigelman, and Ivy B. Rovner. 2007. "The Effects of Negative Political Campaigns: A Meta-Analytic Reassessment." *The Journal of Politics* 69 (November): 1176-1209.
- Layman, Geoffrey C. and Thomas M. Carsey. 2002. "Party Polarization and 'Conflict Extension' in the American Electorate." *American Journal of Political Science* 46 (October): 786-802.

- Le Grand, Julian. 1991. "Quasi-markets and social policy." *The Economic Journal* 101 (September): 1256-1267.
- Le Grand, Julian. 1997. "Knights, Knaves or Pawns? Human Behaviour and Social Policy." *Journal of Social Policy* 26 (2): 149-169.
- Lenz, Gabriel S. 2009. "Learning and opinion change, not priming: Reconsidering the priming hypothesis." *American Journal of Political Science* 53 (October): 821-837.
- Levendusky, Matthew. 2009. *The Partisan Sort: How Liberals Became Democrats and Conservatives Became Republicans*. Chicago: University of Chicago Press.
- Liebman, Jeffrey B. 2013. "Building on Recent Advances in Evidence-based Policymaking." Washington DC: The Brookings Institution.
- Lodge, Milton and Charles Taber. 2000. "Three Steps Toward a Theory of Motivated Political Reasoning." In *Elements of Reason: Cognition, Choice, and the Bounds of Rationality*, eds. Arthur Lupia, Matthew D. McCubbins, and Samuel L. Popkin. Cambridge: Cambridge University Press, 183-213.
- Long, Bridget T. 2004. "How Have College Decisions Changed Over Time? An Application of the Conditional Logistic Choice Model." *Journal of Econometrics* 121 (July-August): 271-296.
- Loveless, Tom. 1997. "The Structure of Public Confidence in Education." *American Journal of Education* 105 (February): 127-159.
- Lowery, David, and William E. Lyons. 1989. "The Impact of Jurisdictional Boundaries: An Individual-Level Test of the Tiebout Model." *Journal of Politics* 51 (February): 73-97.
- Lowery, David, William E. Lyons, and Ruth H. DeHoog. 1995. "The Empirical Evidence for Citizen Information and a Local Market for Public Goods." *American Political Science Review* 89 (September): 705-707.
- Lupia, Arthur. 1994. "Shortcuts Versus Encyclopedias: Information and Voting Behavior in California Insurance Reform Elections." *American Political Science Review* 88 (March): 63-76.
- Lupia, Arthur and Matthew McCubbins. 1998. *The democratic dilemma: Can citizens learn what they need to know?* Cambridge: Cambridge University Press.
- Luskin, Robert C. 1987. "Measuring political sophistication." *American Journal of Political Science* 31 (November): 856-899.
- Luskin, Robert C. and John G. Bullock. 2011. "'Don't Know' Means 'Don't Know': DK Responses and the Public's Level of Political Knowledge." *Journal of Politics* 73 (April): 547-557.

- McConnell, Sheena, Elizabeth Stuart, Kenneth Fortson, Paul Decker, Irma Perez-Johnson, Barbara Harris, and Jeffrey Salzman. 2006. "Managing Customers' Training Choices: Findings from the Individual Training Account Experiment." Washington, DC: Mathematica Policy Research (No. 5222).
- McFadden, Daniel. 2006. "Free Markets and Fettered Consumers." *The American Economic Review* 96 (March): 3-29.
- Miller, Warren E., and Donald Stokes. 1963. "Constituency Influence in Congress." *American Political Science Review* 57 (March): 45-56.
- Nyhan, Brendan and Jason Reifler. 2010. "When Corrections Fail: The Persistence of Political Misperceptions." *Political Behavior* 32 (June): 303-330.
- Obama, Barack. 2009. "Memorandum for the Heads of Executive Departments and Agencies." January 21. Washington, DC: The White House.
- Obama, Barack. 2012. "Fact Sheet: President Obama's Blueprint for Keeping College Affordable and Within Reach for All Americans." Washington, DC: The White House.
- Office of Management and Budget. 2012. "Memorandum for the Heads of Executive Departments and Agencies." May 18. Washington, DC: Office of Management and Budget.
- Osborne, David and Ted Gaebler. 1992. *Reinventing Government: How the Entrepreneurial Spirit is Transforming Government*. Reading, MA: Addison Wesley Public Comp.
- Oreopoulos, Philip and Ryan Dunn. 2013. "Information and College Access: Evidence from a Randomized Field Experiment." *The Scandinavian Journal of Economics* 115 (January): 3-26.
- Page, Benjamin I. and Robert Y. Shapiro. 1992. *The rational public: Fifty years of trends in Americans' policy preferences*. Chicago: University of Chicago Press.
- Partin, Randall W. 2001. "Campaign Intensity and Voter Information: A Look at Gubernatorial Contests." *American Politics Research* 29 (March): 115-140.
- Pax World Funds and National Head Start Association. 2003. "Public Attitudes Toward Head Start Program Survey." iPoll Databank, Roper Center for Public Opinion Research.
- Percy, Stephen L., Brett W. Hawkins, and Peter E. Maier. 1995. "Revisiting Tiebout: Moving Rationales and Interjurisdictional Relocation." *Publius: The Journal of Federalism* 25 (4): 1-17.
- Perna, Laura W. and Marvin A. Titus. 2005. "The Relationship Between Parental Involvement and Social Capital on College Enrollment: An Examination of Racial/ethnic Group Differences." *The Journal of Higher Education* 76 (September-October): 485-518.
- Poole, Keith. 2013. "Party Unity Scores for Democrat and Republican Members of Congresses 35-112 (1857-2012)." Vote View (voteview.com/Party_Unity.htm).

- Popkin, Samuel L. 1991. *The reasoning voter: Communication and persuasion in presidential campaigns*. Chicago: University of Chicago Press.
- Price, Vincent and John Zaller. 1993. "Who Gets the News? Alternative Measures of News Reception and Their Implications for Research." *Public Opinion Quarterly* 57 (Summer): 133-164.
- Pryor, John H., Sylvia Hurtado, Victor B. Saenz, Jose L. Santos, and William S. Korn. 2011. "Survey of the American Freshman." Los Angeles: Higher Education Research Institute.
- Rahn, Wendy. 1993. "The Role of Partisan Stereotypes in Information Processing about Political Candidates." *American Journal of Political Science* 37 (May): 472- 496.
- Redlawsk, David P. 2002. "Hot Cognition or Cool Consideration: Testing the Effects of Motivated Reasoning on Political Decision-making." *Journal of Politics* 64 (November): 1021-1044.
- Rivlin, Alice. 1971. *Systematic Thinking for Social Action*. Washington, DC: Brookings Institution Press.
- Rogers-Dillon, Robin. 2004. *The Welfare Experiments*. Stanford, CA: Stanford University Press.
- Rozin, Paul and Edward B. Royzman. 2001. "Negativity Bias, Negativity Dominance, and Contagion." *Personality and Social Psychology Review* 5 (November): 296-320.
- Schlozman, Kay L., Sidney Verba, and Henry E. Brady. 2012. *The Unheavenly Chorus: Unequal Political Voice and the Broken Promise of American Democracy*. Princeton, NJ: Princeton University Press.
- Schneider, Mark, Paul Teske, Melissa Marschall, and Christine Roch. 1998. "Shopping for Schools: In the Land of the Blind, the One-eyed Parent May Be Enough." *American Journal of Political Science* 42 (July): 769-793.
- Schneider, Mark, Melissa Marschall, Christine Roch, and Paul Teske. 1999. "Heuristics, Low Information Rationality, and Choosing Public Goods Broken Windows as Shortcuts to Information about School Performance." *Urban Affairs Review* 34 (May): 729-741.
- Schneider, Mark, Paul E. Teske, and Melissa Marschall. 2000. *Choosing schools: Consumer choice and the quality of American schools*. Princeton, NJ: Princeton University Press.
- Sides, John and Jack Citrin. 2007. "How Large the Huddled Masses? The Causes and Consequences of Public Misperceptions About Immigrant Populations." Paper presented at the Annual Meeting of the Midwest Political Science Association.
- Simon, Herbert A. 1997. *Administrative Behavior: A Study of Decision-making Processes in Administrative Organizations* (4th ed.). New York: Free Press.
- Smith, Tom W. 1987. "That Which We Call Welfare By Any Other Name Would Smell Sweeter: An Analysis of the Impact of Question Wording on Response Patterns." *Public Opinion Quarterly* 51 (Spring): 75-83.

- Sniderman, Paul M., Richard A. Brody, and Phillip E. Tetlock (Eds.). 1991. *Reasoning and choice: Explorations in political psychology*. Cambridge: Cambridge University Press.
- Somin, Ilya. 2013. *Democracy and Political Ignorance: Why Smaller Government is Smarter*. Stanford, CA: Stanford University Press.
- Soroka, Stuart N. 2006. "Good News and Bad News: Asymmetric Responses to Economic Information." *Journal of Politics* 68 (May): 372-385.
- Steinhauer, Jennifer. 2011. "Cuts to Head Start Show Challenge of Fiscal Restraint." *New York Times* March 11. A14.
- Steuerle, C. Eugene. 2000. "Common Issues for Voucher Programs." In *Vouchers and the provision of public services*, eds. C. Eugene Steuerle, Van D. Ooms, George Peterson, and Robert Reischauer. Washington, DC: Brookings Institution, 3-39.
- Stimson, James A., Michael B. MacKuen, and Robert S. Erikson. 1995. "Dynamic representation." *American Political Science Review* 89 (September): 543-565.
- Taber, Charles S. and Milton Lodge. 2006. "Motivated skepticism in the evaluation of political beliefs." *American Journal of Political Science* 50 (July): 755-769.
- Taylor, Paul, Kim Parker, Richard Fry, D'Vera Cohn, Wendy Wang, Gabriel Velasco, and Daniel Dockterman. 2011. "Is College Worth It?: College Presidents, Public Assess Value, Quality and Mission of Higher Education." Washington, DC: Pew Research Center.
- Tedin, Kent L. and Gregory R. Weiher. 2011. General Social Capital, Education-Related Social Capital, and Choosing Charter Schools. *Policy Studies Journal* 39 (November): 609-629.
- Teske, Paul, Mark Schneider, Michael Mintrom, and Samuel Best. 1993. "Establishing the Micro Foundations of a Macro Theory: Information, Movers, and the Competitive Local Market for Public Goods." *American Political Science Review* 87 (September): 702-13.
- Teske, Paul, Jody Fitzpatrick, and Gabriel Kaplan. 2006. "The Information Gap?" *Review of Policy Research* 23 (September): 969-981.
- Thaler, Richard and Cass R. Sunstein. 2009. *Nudge: Improving Decisions About Health, Wealth, and Happiness*. New York: Penguin.
- Tiebout, Charles. 1956. "A Pure Theory of Local Expenditures." *Journal of Political Economy* 64 (October): 416-424.
- Tomz, Michael and Robert P. Van Houweling. 2009. "The Electoral Implications of Candidate Ambiguity." *American Political Science Review* 103 (February): 83-98.
- Tornatzky, Louis G., Richard Cutler, and Jongho Lee. 2002. "College Knowledge." Claremont, CA: Tomas Rivera policy Institute and SallieMae Fund.
- Tversky, Amos, and Daniel Kahneman. 1974. "Judgment Under Uncertainty: Heuristics and Biases." *Science* 185 (April): 1124-1131.

- Tversky, Amos and Daniel Kahneman. 1991. "Loss Aversion in Riskless Choice: A Reference-dependent Model." *The Quarterly Journal of Economics* 106 (November): 1039-1061.
- U.S. Census. 2012. "Table A-2. Percent of People 25 Years and Over Who Have Completed High School or College, by Race, Hispanic Origin and Sex: Selected Years 1940 to 2012." Washington, DC: Current Population Survey.
- Vargas, Joel H. 2004. *College Knowledge: Addressing Information Barriers to College*. Boston, MA: The Education Resources Institute.
- Weiher, Gregory, and Kent Tedin. 2002. "Does Choice Lead to Racially Distinctive Schools? Charter Schools and Household Preferences." *Journal of Policy Analysis and Management* 21 (Winter): 79-92.
- Whitehurst, Grover J. 2010. "Is Head Start Working for American Students?" Washington, DC: Brookings Institution *Up Front* Blog, January 21.
- Wilson, Matthew and Paul Gronke. 2000. "Concordance and Projection in Citizen Perceptions of Congressional Roll-Call Voting." *Legislative Studies Quarterly* 25 (August): 445-467.
- Zaller, John. 1992. *The Nature and Origins of Mass Opinion*. Cambridge: Cambridge University Press.

Appendix

Table A.1: Citizen Accuracy on Roll Call Votes (Excluding Bill Nelson’s Constituents)

	Correct	Incorrect	Unsure	Difference
Non-Election	45.83 (45.6 to 46.0)	11.77 (11.6 to 11.9)	42.4 (41.9 to 42.9)	4.7 (0.2 to 9.2)
Election	50.56 (50.3 to 50.8)	12.53 (12.3 to 12.7)	36.91 (36.6 to 37.2)	

Table A.2: Regression Model of Citizen Accuracy on Roll Call Votes (Excluding Bill Nelson’s Constituents)

	All Votes	Non-Party Line
Election	0.06***	-0.04***
(s.e.)	0.002	0.025
Education	0.037***	0.01***
	0.001	0.003
Pol. Knowledge	0.31***	0.11***
	0.005	0.014
Co-Partisan	-0.06	0.05***
	0.003	0.012
Ideology	-0.01	-0.04***
	0.004	0.039
Extremity	0.14*	0.02
	0.014	0.170
Party Unity	0.37***	-0.003***
	0.093	0.002
Tenure	-0.00	0.00***
	0.001	0.002
Constant	-0.24*	0.31***
	0.074	0.104
R-squared	0.08	0.03
N	464,378	53,428

Table Note: Cells contain unstandardized coefficients and standard errors from a linear probability model where the dependent variable was scored 1 if the respondent placed their Senator correctly and 0 if the respondent was incorrect or unsure. Respondents are pooled across Senators and issues. Standard errors were clustered at the level of the individual respondent to reflect pooling. Symbols indicate statistical significance: * p<0.05, ** p<0.01, *** = p<.001.

Table A.3: Regression Results Clustering Standard Errors By Senator

	All Votes		Party Line		Non-Party Line	
Election	0.05 **		0.07 ***		-0.04 +	
(s.e.)	0.017		0.015		0.022	
Education	0.04 ***		0.04 ***		0.01 ***	
	0.002		0.002		0.003	
Pol. Knowledge	0.31 ***		0.34 ***		0.11 ***	
	0.009		0.010		0.012	
Co-Partisan	-0.06 ***		-0.07 ***		0.04 ***	
	0.008		0.008		0.011	
Ideology	-0.00		0.00		-0.03	
	0.018		0.016		0.037	
Extremity	0.12 *		0.08		0.02	
	0.053		0.049		0.172	
Party Unity	0.00 *		0.00 *		-0.00	
	0.001		0.001		0.002	
Tenure	0.00		-0.00		0.00	
	0.001		0.001		0.002	
Constant	-0.29		-0.14		0.30 *	
	0.093		0.114		0.117	
R-squared	0.08		0.09		0.03	
N	480,281		419,570		60,234	

Table Note: Cells contain unstandardized coefficients and standard errors from a linear probability model where the dependent variable was scored 1 if the respondent placed their Senator correctly and 0 if the respondent was incorrect or unsure. Respondents are pooled across Senators and issues. Standard errors were clustered at the level of the senator to reflect pooling. Symbols indicate statistical significance: * p<0.05, ** p<0.01, *** = p<.001.