Public information, public learning and public opinion: democratic accountability in education policy

JOSHUA D. CLINTON
Department of Political Science, Vanderbilt University, USA
E-mail: josh.clinton@vanderbilt.edu

JASON A. GRISSOM
Peabody College of Education and Human Development, Vanderbilt University, USA
E-mail: jason.grissom@vanderbilt.edu

Abstract: Theories of political accountability assume citizens use information about the performance of government to hold public officials accountable, but whether citizens actually use information is difficult to directly examine. We take advantage of the importance of citizen-driven, performance-based accountability for education policy in Tennessee to conduct a survey experiment that identifies the effect of new information, mistaken beliefs and differing considerations on the evaluation of public officials and policy reforms using 1,500 Tennesseans. Despite an emphasis on reporting outcomes for school accountability policies in the state, mistaken beliefs are prevalent and produce overly optimistic assessments of the institutions responsible for statewide education policy. Moreover, individuals update their assessments of these institutions in an unbiased way when provided with objective performance data about overall student performance. Providing additional information about race-related performance differences does not alter this relationship, however. Finally, support for specific policies that are intended to improve student performance is unchanged by either type of performance information; opinions about policy reforms are instead most related to race and existing partisan commitments.

Keywords: accountability, education policy, public opinion, survey experiment

Introduction

Fundamental to democratic governance is the willingness and ability of citizens to hold elected officials responsible for their actions and decisions. Electoral accountability requires that citizens be responsive to new information and appropriately update their opinions on the basis of this

355

information so as to correctly reward and punish elected leaders (Key 1966; Dahl 1989). This task is not without its challenges. Not only must citizens pay attention to new information, but they must also be able to interpret the meaning and relevance of such information with respect to their existing beliefs and be willing to update those beliefs in light of this new information.

Public policy reforms are increasingly mandating the reporting of objective performance information with the expectation that greater information means more accountability and superior outcomes (James 2011). In practice, however, the connection between information and opinions may be a tenuous one. Bayesian models of political learning demonstrate that, depending on the content, strength and stability of prior beliefs and other factors, new information can produce either no learning or biased opinion change (Zechman 1979; Achen 1992; Gerber and Green 1999; Achen and Bartels 2006; Bullock 2009).[1] For example, new information consistent with one's prior beliefs will not lead to opinion updating (Hutchings 2003), and information inconsistent with prior beliefs may lead to little or no updating for recipients whose prior beliefs are strong – as may be the case for strong partisans whose opinions are based largely on ideological considerations (e.g. Berelson et al. 1954; Campbell et al. 1960) or in cases where there are disagreements (including partisan or ideological ones) about the interpretation of the information (e.g. Bartels 2002; Lenz 2009). Still other citizens may fail to appropriately update their beliefs in light of new information because they are either uninterested or do not understand the importance of the information (MacKuen 1984). It is also unclear whether information itself may change attitudes given the importance of local conditions and personal experiences (see e.g. Erickson and Stoker 2011; Egan and Mullin 2012); individual experiences may be more influential than policy-relevant information.

Testing the connection between information and beliefs or opinions is critically important for evaluating the health of a representative democracy, but it is also a difficult empirical puzzle. The difficulty arises because, in everyday life, individuals selectively expose themselves to information (e.g. Prior 2007). Even if they choose to consume news, they also choose which type of information to consume (e.g. Stroud 2009).

---

[1] Other models of political learning would also suggest that opinions are responsive to new information. For example, the "on-line" model of Lodge et al. (1995) would predict that new information is immediately used to update opinions, and prior beliefs only define the baseline evaluation that is being updated. The survey-response model of Zaller (1992) and MacGuire (1969) predicts that survey responses are responsive if only because exposure to the information will prime that consideration and cause respondents to use the information when constructing a response to the survey question.

As a result, there is a robust debate regarding the nature of accountability and whether the public holds public officials responsible for outcomes in seemingly irrational (e.g. Achen and Bartels 2002; Healy et al. 2010) or rational (e.g. Malhotra and Kuo 2008) ways (but see Ashworth 2012). We contribute to the critical task of assessing the prospects for democratic accountability by testing whether information about policy outcomes impacts citizens' evaluations of public officials and policy proposals. Given the importance of past performance for future assessments (Woon 2012), we determine whether and how citizens update their initial beliefs about policy in response to objective information about recent performance outcomes for an issue that is salient, important and consequential for the functioning of democracy: public education (Dewey 1916).

Although most of the literature on government accountability focuses on voters' responses to economic performance (e.g. Fiorina 1981; Hibbing and Alford 1981; Stein 1990; Markus 1992; Rudolph 2003), the presumed linkage between information and accountability is perhaps clearest in public education reforms in the United States. During the so-called "standards-based accountability movement" of the 1990s, many states began testing students against a set of standards for each grade and subject on an annual basis to create ratings for school performance (Hanushek and Raymond 2005). By 2001, 45 states created and published "report cards" on schools based on student test performance, and 27 of them used an explicit rating system to identify low performers (Figlio and Ladd 2008). The enactment of No Child Left Behind (NCLB) in 2002 applied such a test-based rating system to every school district in the nation.

The primary rationale for publicising school performance is that such information empowers parents (or the community) to pressure relevant decision-makers – including school staff, local school board members, state officials and others – to increase the performance of less effective schools by finding new resources or using existing resources more efficiently (Dorn 1998; Loeb and Strunk 2007).[2] This "bottom-up" pressure may take the form of informal communication, moving one's child elsewhere or, in the case of elected officials, voting for new representation (Berry and Howell 2007). The implicit theory of action assumes that citizens absorb and act upon school performance information.

Education is an appropriate and important focus for such an investigation for many reasons. Not only is it an area of policy where reforms have focused on increasing data collection and dissemination to promote

---

[2] The incentive for performance may also be elite-driven, but the fact that elite pressure may be sufficient to ensure accountability does not lessen the importance of exploring the extent to which bottom-up accountability is possible.

bottom-up accountability, but the primary policy objective in education – increasing student learning – is also clear and uncontested. There are certainly disagreements over how best to achieve increased student performance or the optimal role for the federal government, and some may care about the policy more than others [e.g. parents with children or homeowners for whom education quality is capitalised into housing values (Black 1999)], but when it comes to evaluating the performance of state governments on statewide educational performance for a given expenditure level, citizens are unlikely to disagree about the need to maximise student performance. The policy outputs of education policy are also directly measurable, comparable and widely available because of the standardised testing regime. State laws and education policies are also made by officials who are either directly elected (in the case of local school boards) or who are overseen by elected officials (in the case of the Department of Education[3]) and who are therefore, in principle, responsive to voters (Berkman and Plutzer 2005). Finally, public education is an important and essential public good. Understanding the foundations of education policy is an important undertaking for political scientists given the close linkages among education, citizenship and democracy (Guttman 1987).

Education is also an area in which few existing studies have found connections between the availability of performance information and citizens' views (e.g. Schneider et al. 2002). For example, parents give higher survey ratings to schools in New York City that have higher test score performance, attendance rates and scores on district quality reviews (Favero and Meier 2013). Chingos et al. (2012) show that survey respondents rate their schools higher when student proficiency rates are greater and that this correlation is higher among parents. They also present mixed evidence that higher state-assigned accountability grades lead to higher citizen ratings, even for schools that perform very similarly, which suggests that respondents respond to information provision. Henderson et al. (2014) similarly find that survey respondents given information about how their local districts compare to state or national averages rate their local schools less well, even when their district outperforms those averages. They also find that these newly informed respondents express stronger support for universal (though not targeted) voucher programs and charter schools, suggesting that information on performance can inform policy preferences.

Utilising a similar approach to Henderson et al. (2014), we explore the potential for citizen-driven accountability in education policy by examining

---

[3] State departments of education are overseen by elected superintendents in 14 states. In other states, including Tennessee, the department is overseen by an appointed superintendent who is in turn overseen by an elected governor or board of education.

the impact of information on citizens' evaluations of institutional performance and their education policy opinions. Given the methodological difficulties of estimating the effects of information when citizens selectively expose themselves to information, we conduct a survey experiment of 1,500 randomly selected Tennesseans. We measure their prior beliefs about statewide educational performance and the connection between their beliefs and their opinions about both public officials and proposed reforms. We then investigate whether and how they update those opinions in response to objective, non-partisan performance information that, in many cases, challenges their prior beliefs. We then characterise whether the effect of information on opinion formation differs by respondent characteristics and the type of performance information that is provided. We show that, while the information does appear to result in citizens updating their beliefs about the institutions of government responsible for education policy in meaningful ways, there is no evidence that the information also affected support for the various policy reforms that have been proposed to boost student performance and decrease the performance gap between races.

The role of information in updating beliefs

Understanding how objective and verifiable information affects the evaluation of public officials and opinions about public policies is critical for evaluating the prospects for democratic accountability. A necessary condition for "bottom-up" accountability is that the opinions of enough citizens must be responsive to new information and experiences to create electoral incentives for elected officials. If citizens ignore new information, or if new information is interpreted in accordance with existing partisan or ideological leanings, then there may be no independent effect of information on opinion formation.

Deriving hypothesised effects requires modelling how individuals process new information and adjust their opinions. Our interest is in assessing how new information affects citizens' support for policies and institutions and the subsequent implications for democratic accountability. Although there are many models of cognitive processing, given our interests, we seek a suitably broad framework that can accommodate the possibility of both null and differential effects.

To this end, we use a Bayesian model of political learning (Zechman 1979; Achen 1992; Gerber and Green 1999; Bullock 2009). By reinterpreting how strong prior beliefs might be and what such beliefs entail, it is possible to accommodate the possibility of partisan bias (e.g. Campbell et al. 1960; Rahn 1993; Bartels 2002; Lenz 2009) or spur of the moment processing based upon primed considerations (e.g. McGuire 1969;

Zaller 1992). A model of Bayesian learning can generate predictions ranging from no learning to biased updating depending on the model's parameters (Achen and Bartels 2006; Bullock 2009).[4] To be clear, our intention is not to "test" this model, but rather to use it to motivate the empirical investigation that follows and to illustrate how a variety of differential effects of information may arise even while holding constant a hypothesised model of citizen learning.

We want to characterise how new information (x) changes individuals' beliefs. Suppose that individual i's opinion about an issue or public official is denoted by μ (for clarity, we drop the individual superscript unless the between-individual variation is relevant). Suppose further that prior beliefs can be thought of as being normally distributed with mean μ_P and variance σ_P^2. That is, absent new information, asking i about her opinion on an issue will result in responses centred at μ_P, but there may be variation due to transient effects [e.g. priming, the ambiguities of how the question is interpreted or other reasons why the survey response may contain error (Zaller and Feldman 1992)]. Opinions may be extremely stable (i.e. σ_P^2 is small), as might be the case if the individual is a highly educated partisan with very strong beliefs (e.g. Popkin 1994), or they may be extremely variable (i.e. σ_P^2 is large), as might be the case if the individual has no political attitudes and has never thought about the issue before being asked about it in the survey.

The effect of new information is simply the change in opinion that results. If the effect of new information can be thought of as being normally distributed with a mean of μ_I and a variance of σ_I^2, Bayesian updating requires that the opinion of individual i is a combination of prior beliefs and the new data, with the impact of each determined by the relative strength of each. Mathematically, the new (posterior) opinion is

$$\mu = \mu_P \left(\frac{1/\sigma_P^2}{(1/\sigma_P^2 + 1/\sigma_I^2)} \right) + x \left(\frac{1/\sigma_I^2}{1/\sigma_P^2 + 1/\sigma_I^2} \right)$$

with the precision of the new belief being given by $1/\sigma_P^2 + 1/\sigma_I^2$. The effect of the new information is the difference $\mu - \mu_P$.

___

[4] Of course, other models of political learning would also suggest that opinions are responsive to new information. For example, the "on-line" model of Lodge et al. (1995) would predict that new information is immediately used to update opinions and prior beliefs play no role in the process other than defining the baseline evaluation that is being updated. The survey-response model of Zaller (1992) and MacGuire (1969) would suggest that opinions would be responsive if only because exposure to the information will prime that consideration and cause respondents to use the information when constructing a response to the survey question being asked.

This seemingly sparse model reveals several possible effects of new information. One possibility is that there is no effect: $\mu - \mu_P = 0$. A null effect is possible for several reasons. First, the information that is provided may already be known or be consistent with existing opinions [e.g. individuals for whom the issue is relevant may possess correct beliefs (Hutchings 2003)]. If so, $\mu = \mu_P$, and we would obviously expect no difference.[5] Second, existing beliefs may be so strong so as to make the new information irrelevant. If $1/\sigma_P^2 > 1/\sigma_I^2$ and the difference is large, opinions may be unchanged even if $\mu \neq \mu_P$ and the difference is dramatic. This might be the case if the individual's belief is based on considerations that are unchanged by the new information or if the new information is thought to be untrustworthy. For example, parents with school-aged children may have strong beliefs about schools based largely on personal experiences. If so, new performance information may not change their opinions. Similarly, strong partisans may be less responsive to new information because their opinions are based purely on partisan considerations (e.g. Berelson et al. 1954; Campbell et al. 1960).[6]

Information updating occurs if the new information differs from their existing beliefs *and* individuals are sufficiently motivated to update their existing beliefs (MacKuen 1984; Kuklinski et al. 2001). Individuals are receptive to new information if their existing beliefs are sufficiently imprecise (i.e. σ_P^2 is large) or if the new information is precise in its implications (i.e. σ_I^2 is small). For a given piece of information, if there are no differential perceptions of the clarity of the new information (i.e. σ_I^2 is constant across individuals), differential effects between individuals can emerge if there is variation in the strength of existing priors (σ_P^2) or in how distant the prior belief is from the new information; those with stronger prior beliefs and more accurate perceptions are less sensitive to new information.

Differences in whether and how individuals update their beliefs in response to new information may result from both individual level differences (e.g. how important and interested an individual is in education policy), as well as partisan differences. For example, the importance of statewide examinations, graduation rates and other seemingly objective measures may differ by political orientation. This is because of how partisans interpret such data in light of the real or imagined political orientations of public officials

---

[5] This possibility highlights the need to control for existing beliefs when evaluating the effect of new information, because the new information will obviously only affect those for whom the information is "new".

[6] For example, in our case, Democrats may support educational institutions regardless of the performance of these institutions, or else they may be less supportive in Tennessee because the statewide policies are a result of the Republican-controlled legislature and governorship.

responsible for designing and implementing education policy. Accordingly, the contribution of the new information will differ between partisans even if they share a common goal of increased student performance given current expenditure levels.

While most education policies at least implicitly assume that there is a strong relationship between information and opinions, we do not actually know how providing objective, outcome-based information affects citizens' assessment of those institutions that are responsible for designing and implementing education policy at the state level. We also do not know how such information influences the support for various policies aimed at increasing student performance. While we may hope that assessments are responsive to information in a way that brings about the convergence of citizens' beliefs, there are also reasons to think that there may be either no effect or effects that depend on existing beliefs in ways that may prevent citizens from agreeing, even on issues such as education where everyone presumably concurs on the desirability of increasing student performance.

Experimental design

We test the straightforward hypothesis that providing citizens with objective information about the public education system's recent performance will change their opinions about (a) the performance of educationally relevant government institutions and (b) education policies and reforms, at least among the population(s) for whom the objective information is inconsistent with their prior beliefs about system performance. While others have looked at how mistaken beliefs correlate with opinions (e.g. Sides and Citrin 2007), we employ an experimental design that identifies the effect of providing information, controlling for existing beliefs. This detail is important, because it avoids the complications that may result from differences in individuals' abilities to form accurate initial beliefs.

We utilise two pieces of objective performance information: (1) student achievement on standardised math tests and (2) the extent to which student achievement on these tests varies by the race of the student (the racial "achievement gap"). Under NCLB, standardised test scores in math and reading form the basis of school accountability in every school district in the United States.[7] Within each state, common tests cover the same material for each grade level, so scores have the same meaning across schools and

---

[7] We focus on math performance to simplify data collection; in practice, math and reading performance are very highly correlated across schools. NCLB also requires subgroup-level reporting for other groups, including economically disadvantaged students and students with disabilities.

districts. NCLB requires schools to publicly report proficiency levels from these tests both for the school as a whole and for racial and ethnic subgroups, in part based on the assumption that parents, communities and other stakeholders can utilise these data to pressure schools to better serve the needs of students (see Figlio and Loeb 2011).

We embedded a survey experiment within a larger Random Digit Dial survey of 1,500 citizens of Tennessee. Tennessee provides a useful laboratory for this experiment, because their long history with school performance data means that Tennesseans are among the most familiar and, presumably, comfortable with their usage. Tennessee was a relatively early adopter of the school accountability policies that predated NCLB (Hanushek and Raymond 2005), and, prior to the mandates of NCLB, the state based its accountability policy almost purely on making information available to the public rather than on using student test data for the kinds of administrative interventions favoured in other "consequentialist" accountability states (Carnoy and Loeb 2002). The use of student data for school improvement has maintained a high profile in the state in recent years; for example, in 2011, a law passed as part of Race to the Top reforms mandated that 50% of a teacher's annual evaluation must come from the standardised test performance of his or her students (Sher 2011). In the wake of Race to the Top – with Tennessee as one of two first-round awardees in the federal grants competition, receiving more than $500 million – the state also began increasing the presence of charter schools, creating a statewide Achievement School District to turn around its lowest-performing schools, moving towards implementation of the Common Core State Standards and investing heavily in innovation in science and mathematics instruction, among other reforms, ensuring the salience of education reform among the public at the time of the study.

The survey experiment contained five randomly assigned conditions comprising two experiments. As Table 1 summarises, in four of the conditions (conditions 2–5), respondents' beliefs about school performance in Tennessee were measured by asking about student performance on standardised math tests. Half were asked a single question about student performance on end-of-year math exams (the "performance experiment" in conditions 2 and 3), and half were asked this question and a question about the race-related gap in student performance on these tests (the "achievement gap experiment" in conditions 4 and 5).[8] In two of the four

---

[8] Respondents were asked to assess the performance of Tennessee's public schools via the question, "Based on state standards, what percent of elementary and middle school students in Tennessee do you think performed at grade level or better on Tennessee's end-of-grade math tests?" Respondents were provided with five possible ranges within which to answer (0–19%, 20–39%, 40–59%, 60–79%, 80–100%). To assess the possible achievement gap between black

Table 1. Experimental design

	Performance Experiment		Achievement Gap Experiment	
Condition 1: No Prime or Correction	Condition 2: Performance Prime (Control)	Condition 3: Performance Prime and Correction (Treatment)	Condition 4: Performance + Achievement Gap Prime (Control)	Condition 5: Performance + Achievement Gap Prime and Correction (Treatment)
	Performance question	Performance question Correct answer provided	Performance question Achievement gap question	Performance question Achievement gap question
				Correct answer provided (Both)
Institutional evaluation questions	Institutional evaluation questions	Institutional evaluation questions	Institutional evaluation questions	Institutional evaluation questions
Policy preference questions	Policy preference questions	Policy preference questions	Policy preference questions	Policy preference questions
$n = 150$	$n = 330$	$n = 345$	$n = 330$	$n = 345$

conditions, respondents received the treatment of being told the correct answer(s) after they expressed their beliefs concerning performance (condition 3) or both overall performance and the achievement gap between black and white students (condition 5). To identify the impact of the information itself rather than the impact of the framing of student performance (see, e.g. Chong and Druckman 2007), the actual performance was reported without commentary.[9]

After the intervention, respondents in conditions 2 through 5 were asked (1) to rate the performance of various public institutions involved in setting education policy and (2) what they thought about various educational reforms that have been proposed. Respondents in the remaining group (condition 1) were asked the same battery of evaluations, but they were not primed to consider performance issues beforehand. This group lies outside both the performance and achievement gap experiments.

Our design identifies several effects of interest. Because asking citizens about student performance may prime considerations that are not commonly used when citizens articulate preferences for education policies or the public officials responsible for education policy (see, e.g. the theories of McGuire 1969; Zaller 1992 and Zaller and Feldman 1992), the experimental manipulation may itself affect the evaluation by priming the respondent to think in terms of performance or equity considerations when answering the questions. We can identify the possible priming effect by comparing the responses of condition 1 to condition 2 (and also conditions

and white students, we asked: "Now thinking about the performance of white and black elementary and middle school students, historically white students have performed more strongly on state math tests than black students. What do you think was the difference between the percentage of white students and black students who performed at grade level or better on Tennessee's end-of-year math tests?" Respondents were provided with nine possible ranges within which to answer (no difference in performance, 0–5% more white students than black students were at grade level, 6–10%, 11–15%, 16–20%, 21–25%, 26–30%, 31–35%, more than 35% more white students than black students were at grade level).

[9] After their responses were recorded but before being asked the assessment and policy preference item, the interviewer provided the *actual* performance level to treatment 3 by reading: "You thought that *<fill in based on their answer>* of elementary and middle school students in Tennessee last year performed at grade level or better on Tennessee's end-of-grade math tests. The actual percentage of students performing at grade level or better on Tennessee's end-of-grade math tests was 34%". For treatment 5, the respondents received: "You thought that *<fill in with performance response>* of elementary and middle school students in Tennessee last year performed at grade level or better on Tennessee's end-of-grade math tests. The actual percentage of students performing at grade level or better on Tennessee's end-of-grade math tests was 34%. You also thought that *<fill in with gap response>*. Actually, the gap between the performance of white students at grade level and the percentage of black students at grade level on the Tennessee state math tests was 22 percentage points". The design is similar to the one employed by Guillermo et al. (2011) to study the effect of knowing the income distribution on preferences towards redistribution.

1 to 4). To identify how the performance information we provide affects the opinions of otherwise similar individuals, we compare individuals' responses in conditions 2 and 3. The difference in evaluations and opinions reveals whether individuals with otherwise identical characteristics and beliefs about statewide student performance differ as a result of being exposed to the objective performance information. Because we can condition on prior beliefs, we can identify the effect of the information we provide holding initial beliefs fixed. We also examine if the effect varies depending on how important educational issues are to the respondent, because the importance of an issue is presumably related to the strength of prior beliefs or the motivation to update beliefs. Comparing the differences in conditions 2 and 4 reveals how additionally priming racial gap considerations – and, more specifically, the racial disparity in educational performance – affects opinions. Do opinions change if respondents are thinking not only in terms of overall performance, but also in terms of the relative performance of students by race?

Replicating the comparison for conditions 2 and 3 using conditions 4 and 5 reveals how providing information about student performance and student performance by race affects evaluations. Not only is the comparison between the corrected and uncorrected individuals of interest, but it is also of interest to see how the overall effect of providing these two pieces of information compares with the effect revealed when comparing conditions 2 and 3.

The accuracy of prior beliefs about student performance and racial gaps

The effect of information presumably depends on both the accuracy and strength of existing beliefs. The first task in identifying the effect of information on citizens' evaluations of public officials and public policy proposals therefore involves assessing the strength and accuracy of existing beliefs (Delli Carpini and Keeter 1996). Figure 1 graphs the distribution of beliefs regarding the percentage of elementary and middle school students who are performing at grade level or better on Tennessee's end-of-grade math tests (left) and the difference between the percentage of white students and black students performing at grade level or better on Tennessee's end-of-year math tests (right).

Figure 1 facilitates several conclusions. First, as is the case for other issues (e.g. Gilens 2001; Kuklinski et al. 2001), very few citizens hold accurate beliefs. Despite the amount of attention paid to the issue and the number of policies in Tennessee that use student performance data, only 20% chose the response category containing the true level of student performance (34%), and only 8% chose the category containing the true gap in student
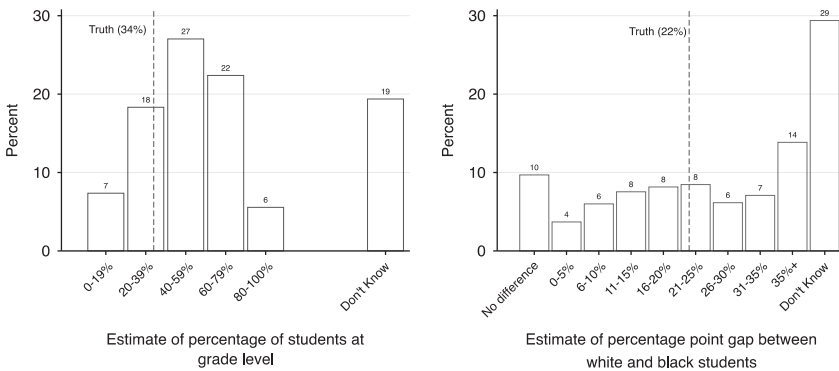
Figure 1   Distribution of citizens' beliefs.
Note: The figures provide the distribution of responses using the 1,328 respondents in conditions 2–5 who were asked the overall performance question (left) and the 650 respondents in conditions 4–5 who were asked about possible racial disparities in student performance (right). The vertical line denotes the true percentage in each instance.

performance (22%).[10] In fact, the nearly uniform distribution of responses to the achievement question suggests that the 71% of respondents who chose a response other than "Don't know" were simply guessing. Second, Figure 1 reveals that citizens are less likely to possess correct beliefs about the race-related gap.

Third, citizens' beliefs about statewide overall performance are too optimistic; 54% of the respondents think student performance is better than it actually is. However, a larger percentage of respondents also think the racial gap in student performance is larger than it actually is than smaller than it actually is (36% and 27%, respectively).[11] The inaccuracy of beliefs does not vary according to the importance of education to the respondent.

[10] This number was based on the fraction of students in kindergarten through eighth grade who scored at the level of "proficient" or better on the 2009–2010 round of statewide standardised tests. State-level data from 2010–2011, the most recently completed academic year, were not yet publicly available at the time of the survey, so respondents could not know statewide results from those tests. When those results were released in January 2012, they showed that, for the 2010–2011 school year, 41% of tested elementary and middle school students attained grade-level proficiency or better in mathematics, a 7 percentage-point increase.

[11] If we examine the distribution of knowledge among the 424 respondents in conditions 4 and 5 who answered both the performance and the racial disparity question with a response other than "don't know", the responses exhibit a slightly negative correlation (Spearman's $r = -0.15$, p<0.01)—the higher a respondent thought the overall performance of the state was, the smaller the racial difference in performance was thought to be. Interestingly, black and white respondents provided qualitatively similar assessments of the black–white gap ($\chi^2_9 = 5.8$, p = 0.76).

In an additional analysis (not shown), we measured salience – and, presumably, strength of prior beliefs regarding education issues – by using whether the respondent believes education should be the top priority of the Tennessee government,[12] whether the individual has children that attend public school and whether the respondent owns their home or has a mortgage (see Figlio and Lucas 2004). Proportions with correct performance information were similar and statistically indistinguishable across these conditions.

The fact that many citizens are unaware of actual student performance despite the fact that it is a centrepiece of accountability-based educational reforms is consistent with the pervasive lack of information that the public has routinely exhibited on political issues (e.g. Campbell et al. 1960; Delli Carpini and Keeter 1996). However, what matters are the consequences of the misinformation (e.g. Bartels 1996) and whether citizens are willing to update these beliefs and the opinions for which these beliefs are relevant. It is to this analysis that we now turn.

Estimating the effect of objective information about performance outcomes

Having shown that citizens often misperceive – by large margins – the performance of the public education system on commonly used metrics, we now assess: (1) whether misperceptions about performance and opinions are indeed linked, (2) whether correcting citizens' misperceptions via the provision of performance information leads to changes in policy opinions and (3) whether opinions are differentially responsive to different types of performance information (i.e. overall performance versus black–white gaps). We investigate these questions for opinions based on the performance of three education institutions (Tennessee schools as a whole, the Tennessee Department of Education and the local school board) and by which education reforms should be pursued.

The effect of performance information on evaluations of educational institutions

We begin by assessing whether the act of simply *asking* about student performance primes considerations and influences evaluations, even in the absence of new information. There is no evidence that priming affects institutional evaluations; given our sample size, we have sufficient power to

---

[12] This question was asked well before the experimental treatments so as to not confound the effects or be confounded by the treatments.

detect differences of 0.2 on the five-point scale we use.[13] Comparing the evaluations for respondents in conditions 1 and 2 via *t*-tests reveals that the smallest p-value that was obtained was 0.55.[14] Comparing the average responses for conditions 1 and 4 also produces null results even though the respondents in condition 4 are asked to think about both overall performance and race-related differences.[15]

To better test the association between performance information and evaluations of education institutions, we estimate a series of ordered probits of the form:

$$\Pr\left(Y_{ig} = \{A, B, C, D, F\}\right) = \beta_0 + \beta_1 \text{Belief}_i + \beta_2 T_i + X_i \delta + \varepsilon_i \quad (1)$$

The dependent variable in equation (1) is the grade respondent *i* gives to institution *g*, using a survey question that asks every respondent to assign a grade of *A, B, C, D* or *F* to "public schools in Tennessee," "the Tennessee Department of Education" (TDOE) and "your local public school board." While the assumptions of the ordered probit are most appropriate given the ordinal nature of the evaluations, the supplemental appendix shows that translating the grades into their GPA equivalents (i.e. $A = 4, B = 3, C = 2, D = 1, F = 0$) and using OLS produces identical substantive results.

We first compare conditions 2 and 3. Subjects in these two conditions were asked to assess the performance of Tennessee schools after providing their estimate of the percentage of elementary and middle school students performing at grade level or better according to state standardised tests using a series of 20 percentage point ranges (i.e. 0–19%, 20–39% and so forth). We use these responses to measure prior beliefs using a series of indicator variables (i.e. the interval containing subject *i*'s estimate, including "I don't know," is set to 1, and all other intervals are set to 0, with the 0–19% interval being the omitted category).

Because all respondents in condition 3 were told the correct answer, the treatment variable T_i is set to 1 if subject *i* is in condition 3 and 0 if in condition 2. The row-vector X_i for individual *i* contains the control variables used to improve the estimates' precision. These include indicators for: female, black, Democrat, Republican, having a college education, having children in school and owning a home, as well as a linear (three-item) ideology

---

[13] For a sample size of 330, and for a mean and standard deviation of 3.40 and 1.28 (calculated using the untreated control group's evaluation of local schools), for $\alpha = 0.05$ we have a power of 0.81 to detect a difference of 0.2 on the five-point scale.

[14] The difference in average evaluations for "schools in Tennessee" is –0.06 (with a standard error of 0.10), for "the Tennessee Department of Education" 0.04 (0.11), and for "your local school board" –0.05 (0.12).

[15] The smallest p-value for the hypothesis test of no difference is 0.14, followed by 0.42 and 0.66.

scale, age, age squared, the number of years residing in Tennessee and a six-category measure of respondent income. In equation (1), β_2 estimates the average effect of being told the true performance level, and the coefficient vector β_1 measures the association between the prior beliefs of i and the grade i assigns to institution g.

Table 2 presents the results. The odd-numbered columns display the results of estimating equation (1).[16] Several important conclusions are evident. First, as expected, respondents' evaluations of both Tennessee schools in general and TDOE are increasing in their beliefs about statewide student math performance; the better an individual thinks student performance is, the higher the grade that was given. Those citizens who overestimated student performance the most (i.e. a performance guess of 80–100% performing at grade level) also gave the highest average grade to the educational institution. However, this pattern is least true for evaluations of the local school board (Model 5). Consistent with the murkier connection between statewide performance and the efficacy of one's local board, beliefs about statewide performance are largely uncorrelated with local board evaluations.

Second, the average effect of receiving the informational update containing the true student performance level is negative and statistically distinguishable from zero at conventional levels for all three institutions. Moreover, the magnitude of the effect is sensibly ordered: the effects are largest for the evaluation of Tennessee schools (Model 1) and the TDOE (Model 3), which are most responsible and relevant for statewide performance, but there is little effect of learning about statewide performance on citizens' evaluations of local school boards (Model 5).[17]

---

[16] Control variables are omitted from the tables for brevity, though see Table A.1 in the supplemental appendix for full results. Most are statistically indistinguishable from zero across models, with a few notable exceptions. First, being a Democrat positively predicts the evaluation of all three institutions, as does having a child in school. African Americans and homeowners evaluate the local school board significantly worse, though neither race nor homeownership is significant for the other two dependent variables.

[17] We also investigated whether the impact of the treatment varied by the performance of the respondent's school district. For this analysis, we first had respondents' phone numbers geocoded to their billing addresses and used geographic information systems technology to place those addresses in school districts. Next, we obtained 2011 district report card data from the TDOE web site. These data included district mean student achievement information for K-8 math and reading and high school end-of-course tests, plus ACT scores, high school graduation rates and suspension rates. Factor analysis of these measures showed 1 underlying latent trait, which we take to be district performance. Factor scores from this analysis were then assigned to each respondent as a measure of the performance of the district in which he or she resided. We then re-estimated the odd-numbered models in Table 2, interacting the treatment variable with the district performance measure. This interaction was not significant for any of the models, nor was an indicator for being below the mean. Thus, it does not appear that receiving performance

Table 2. The effect of prior beliefs about performance on evaluations of education institutions

Grade for	Tennessee Schools		Tennessee Department of Education		Local School Board	
	(1)	(2)	(3)	(4)	(5)	(6)
Performance guess						
20–39%	0.147 (0.204)	0.157 (0.279)	0.151 (0.209)	0.222 (0.287)	0.156 (0.207)	– 0.191 (0.280)
40–59%	0.244 (0.196)	0.313 (0.275)	0.339* (0.201)	0.380 (0.282)	0.187 (0.198)	0.027 (0.274)
60–79%	0.470** (0.200)	0.879*** (0.284)	0.652*** (0.205)	0.868*** (0.292)	0.557*** (0.203)	0.177 (0.284)
80–100%	0.908*** (0.280)	1.403*** (0.378)	1.353*** (0.281)	1.837*** (0.386)	0.312 (0.274)	0.341 (0.369)
Don't know	0.238 (0.223)	0.201 (0.308)	0.479** (0.232)	0.327 (0.318)	0.443** (0.224)	0.021 (0.306)
Received performance update	–0.55*** (0.098)	–0.299 (0.344)	–0.37*** (0.098)	–0.224 (0.353)	–0.249** (0.097)	–0.822** (0.354)
Received performance update × performance guess						
20–39%		0.021 (0.403)		–0.127 (0.415)		0.773* (0.415)
40–59%		–0.143 (0.388)		–0.081 (0.397)		0.383 (0.396)
60–79%		–0.754* (0.396)		–0.391 (0.404)		0.779* (0.405)
80–100%		–1.127** (0.559)		–1.020* (0.554)		–0.081 (0.549)
Don't know		0.117 (0.435)		0.368 (0.446)		0.909** (0.443)
Observations	511	511	499	499	500	500
Pseudo R^2	0.061	0.072	0.075	0.082	0.041	0.047

Ordered probit coefficients shown.
Models also condition on control variables.
Standard errors in parentheses.
*p < 0.10, **p < 0.05, ***p < 0.01.

These results assume that the effect of information does not depend on prior beliefs. If the performance update treatment affects institutional evaluations via the adjustment of respondents' posterior beliefs, however, the treatment effect should be the greatest among those respondents who most overestimate student performance. To test this hypothesis, we test for a possible interaction between prior beliefs and treatment status using the specification of equation (2):

$$\Pr\left(Y_{ig} = \{A, B, C, D, F\}\right) = \beta_0 + \beta_1 \text{Belief}_i + \beta_2 T_i + \beta_2 (T_i \times \text{Belief}_i) + X_i \delta + \varepsilon_i \quad (2)$$

The even-numbered columns of Table 2 report the results of estimating equation (2) by ordered probit. Several important refinements emerge. First, columns 2 and 4 reveal that the negative effect of the performance update is driven by the substantially lower performance evaluations given to Tennessee schools and TDOE by those respondents who most overestimate student performance. Figure 2 graphs the substantive magnitude of this effect on the probability of assigning a grade of A. Margins are shown separately for those in the treatment and control groups, with the vertical bracketed lines corresponding to 95% confidence intervals. (The effects are substantively similar using the probability of assigning a B or higher, and the supplemental appendix replicates the results using OLS.)

Figures 2a and 2b reveal that the provision of objective performance information similarly affects citizens' evaluations of Tennessee schools and TDOE. In both cases, the line denoting the opinions of the control group shows that, in the absence of information, the respondents who most severely overestimate student performance are also much more inclined to assign a high grade to institutional performance. In contrast, those given the information update in the treatment group have a roughly equally low probability of assigning the highest grade regardless of their prior belief. Because respondents are equally likely to assign a grade of "A" regardless of their initial beliefs even after they are told the actual level of performance, this pattern is consistent with the evaluations being based on the provided information rather than prior beliefs or individual experience.

Figure 2c provides an important contrast by showing that the information has almost no effect on citizens' evaluations of local school boards. This null effect is consistent with the observation that evaluations of local school boards should not be affected by the information provided by the experimental condition, because information about statewide performance is irrelevant for assessing the performance of the school board.

information affects institutional evaluations differently in relatively high- and low-performing districts.
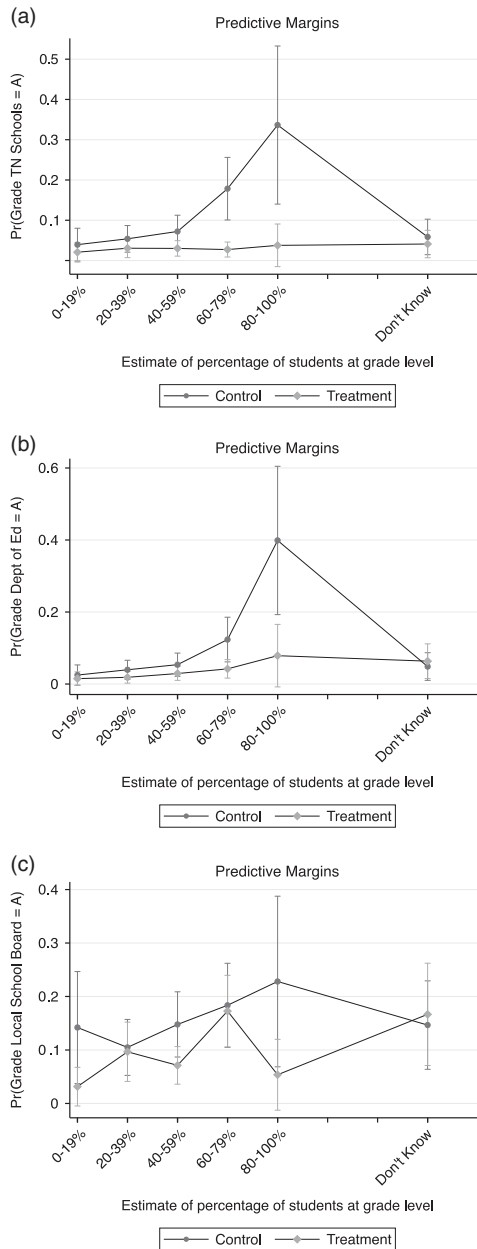
Figure 2 The effect of prior beliefs on the effect of information. (a) Probability of assigning Tennessee schools a grade of A; (b) probability of assigning the Tennessee Department of Education a grade of A; (c) probability of assigning the Local School Board a grade of A.

We also tested whether the effect of information on evaluations varied according to the likely importance of the issue of education to the respondent. We did so using measures of issue salience based on having a child in school, owning a home and naming education as the top priority for state government. We might expect salience to matter, for example, if it predicts stronger prior beliefs that are more difficult to update. For instance, parents can access numerous sources of information about school performance on a variety of dimensions – including their day-to-day interactions with their children's schools, which may make them less responsive to test score information. We re-estimated equation (1) including interactions between each measure of issue salience and the performance assessment indicators and the treatment indicator. Table A.1 in the supplemental appendix online reveals no evidence that there are differences in the treatment effect related to the saliency of the issue.[18]

The effect of performance information on support for policy reforms

Ostensibly, citizens' support for a policy initiative is driven in part by perceptions that it can address deficiencies in the status quo. If so, does learning about the status quo change opinions about specific public policies? More specifically, does learning that student performance is lower than expected increase citizens' support for education reforms?

We answer this question for six policies common to current education reform debates: test-based performance pay for teachers, NCLB, governmental provision of pre-kindergarten programs, public vouchers for private school attendance, charter schooling and differential pay to incentivise teachers to work in low-income schools. Following the three institutional evaluation questions, each respondent was asked whether or not they support each of these six policies using the questions listed in supplemental appendix B. The effect of information was identified by re-estimating equations (1) and (2) to predict the support for each policy separately using probit models.

Figure 3 summarises the effect of information on policy opinions by graphing the predicted probabilities of supporting each reform for otherwise typical and identical individuals in the treatment and control groups and allowing the effect of information to vary by prior beliefs about student performance [equation (2)].[19] The six panels reveal that the effect of information is quite different from the effects evident in prior sections.

---

[18] There is some evidence in Table A.1 that homeowners are not as negatively impacted by the performance update as non-homeowners when evaluating TDOE.

[19] Estimating the specification of equation (1) reveals substantively identical effects. Given the number of tables needed to display the 12 regression specifications, the results are available from the authors upon request.
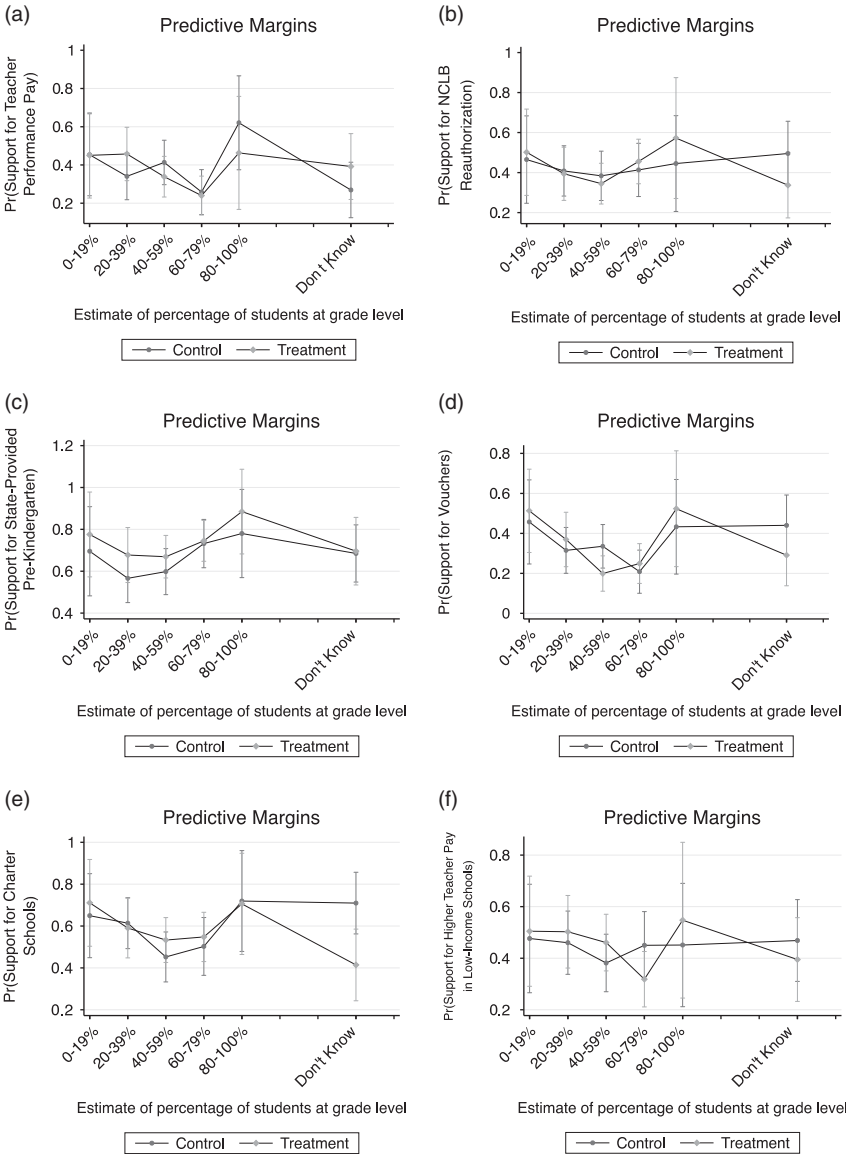
Figure 3 Beliefs about student performance and support for education policy reform with and without informational updating. (a) Teacher performance pay; (b) No Child Left Behind; (c) state-provided pre-kindergarten; (d) private school vouchers; (e) charter schools; (f) higher pay for teachers in low-income schools.

First, there is not a decreasing (or increasing) association between a respondent's prior beliefs about student performance in the status quo – shown on the x-axis – and the likelihood of supporting *any* of the examined reforms. There is slight evidence of a U-shaped relationship (i.e. respondents in the 40–59% or 60–79% categories have statistically significantly lower likelihoods of support than categories on the ends) for three of the six policies (performance pay, vouchers and charter schools), but the reason for this relationship is unclear.

Second, providing correct information about the status quo has no effect on the probability that a citizen supports any of the policies. Regardless of prior beliefs, receiving information about actual student performance has no impact whatsoever on their support for some of the reforms that have been proposed to increase student performance.

What may explain the pervasive null effect of information on the support for proposed reforms? One possibility is simply that respondents do not believe that any of these reform strategies are likely to produce changes in student achievement, so they do not update their beliefs in response to this information. Alternatively, the Bayesian learning model described in the first section reveals that performance information and its updating will not substantially affect opinion formation if prior beliefs about policies are strong. Although citizens possess inaccurate beliefs about student performance and are seemingly willing to update them when called upon to evaluate educational institutions, it is possible that opinions about public policies are more strongly held because they are closely tied to the individuals' ideology and partisanship. If partisan or ideological beliefs drive policy preferences, or if there are partisan and ideological disagreements about the efficacy of the various policies (and their costs), correcting mistaken beliefs about student performance may be insufficient to change opinions about the policies themselves (Rahn 1993). In other words, these results would be consistent with citizens' policy evaluations being primarily driven by party and ideology (Campbell et al. 1960; Jacoby 1988; Green et al. 2002; Highton and Kam 2011).

Examining the covariates of these models reveals evidence consistent with this explanation. Two factors – political ideology and race – are the only consistent predictors of respondents' policy opinions across the various specifications. The joint test that the ideology and party variables are statistically indistinguishable from zero can be rejected at the 0.05 level in four of the six models (performance pay and NCLB are the exceptions). Perhaps because of the racial gap in student performance, black respondents are more supportive of five of the six proposed reforms (differentiated pay for teachers in low-income schools is the exception).

The effect of information about the racial achievement gap

So far, we have considered the effect of providing performance information about the overall level of student performance. In reality, citizens may also be responsive to other types of performance information, and different citizens may respond differently depending on the type of information. In education policy, for example, a significant amount of research and public debate focuses on the achievement gaps between students with different backgrounds, particularly with respect to race (Hochschild and Shen 2012). Closing the achievement gap between white students and black students – estimated to be a standard deviation or more on standardised tests (Fryer and Levitt 2004) – is a demonstrably important goal in education and a central aim of many education reform efforts. In evaluating the performance of education institutions or making decisions about their support for particular education policy changes, do citizens' perceptions about the relative performance of white and black students inform their evaluations?

To characterise the effect of the information, we use conditions 4 and 5 in our survey experiment (see Table 1). Subjects in these conditions were asked to estimate the percentage of elementary and middle school students testing at grade level in math (performance prompt) and to estimate the difference in this percentage for white and black students (achievement gap prompt). Respondents in condition 4 serve as the control group. In condition 5 (treatment), respondents were given an information update containing the true percentages for overall performance (34%) and the percentage gap between white and black students (22 percentage points).

Table 3 reports the associations between prior beliefs about overall performance and race-related performance differences and institutional evaluations. The coefficients for prior beliefs about overall performance reported in the top half of Table 3 reveal that institutional evaluations increase when respondents' overestimate overall student performance. This pattern is consistent with the results of Table 2. However, in contrast to the prior results, the coefficients for prior beliefs regarding the race-related achievement gap show no clear pattern, and most are indistinguishable from the reference group of those who believe that there is no gap.[20] All else equal, individuals who think that there is no difference in student achievement and those who think that the performance gap is more than 35% provide the same grade to educational institutions. This finding suggests that, while evaluations of education officials depend on beliefs about overall student performance, evaluations do not depend on beliefs regarding race-related differences in student performance.

---

[20] A comparable set of probit models run for the six policy opinion variables also show scant evidence of a pattern between prior beliefs about equity and opinions.

Table 3. *Prior beliefs about achievement gaps do not predict institutional evaluations*

Grade for	Tennessee Schools (1)	Tennesse Department of Education (2)	Local School Board (3)
Performance guess			
20–39%	–0.04 (0.21)	0.12 (0.21)	0.32 (0.21)
40–59%	0.10 (0.20)	–0.01 (0.20)	0.48** (0.20)
60–79%	0.29 (0.21)	0.17 (0.21)	0.72*** (0.21)
80–100%	1.19*** (0.29)	0.88*** (0.29)	0.90*** (0.29)
Don't know	0.30 (0.23)	0.29 (0.23)	0.50** (0.23)
Achievement gap guess			
1–5%	–0.01 (0.30)	0.32 (0.32)	0.19 (0.30)
6–10%	–0.10 (0.25)	0.01 (0.26)	–0.17 (0.25)
11–15%	–0.11 (0.22)	0.02 (0.23)	–0.20 (0.22)
16–20%	–0.27 (0.23)	–0.32 (0.23)	–0.03 (0.22)
21–25%	–0.02 (0.22)	0.01 (0.23)	0.08 (0.22)
26–30%	–0.05 (0.24)	–0.26 (0.25)	0.03 (0.24)
31–35%	–0.25 (0.23)	–0.30 (0.24)	–0.41* (0.23)
35% +	–0.10 (0.20)	–0.30 (0.21)	–0.16 (0.20)
Don't know	–0.48** (0.19)	–0.34* (0.19)	–0.18 (0.19)
Received performance and equity update	–0.20** (0.10)	–0.04 (0.10)	0.03 (0.10)
Observations	509	504	507
Pseudo R^2	0.051	0.042	0.037

Ordered probit coefficients shown.
'No difference' is the omitted category for the equity guess.
Models also condition on control variables.
Standard errors in parentheses.
*$p < 0.10$, **$p < 0.05$, ***$p < 0.01$.

To identify whether priming or updating information about race-based achievement differences affects citizens' opinions, we pool data from all subjects in conditions 2 through 5 and examine their institutional evaluations. We estimate a version of equation (1) that controls for prior beliefs about student performance and includes indicators for the condition to which the respondent was assigned. An indicator's coefficient tells us the average response change that is attributable to random assignment to that condition relative to the excluded category (condition 2, the control group for the performance experiment).[21]

[21] Determining if information about overall performance and equity differentially affect opinion formation at different points in the distribution of prior beliefs is more complicated

Table 4. Equity prime, equity information update and institutional evaluations

Grade for	Tennessee Schools (1)	Tennessee Department of Education (2)	Local School Board (3)
Condition 3: received overall performance prime and update	−0.54*** (0.10)	−0.33*** (0.10)	−0.26*** (0.10)
Condition 4: received performance and equity prime	−0.11 (0.09)	−0.17* (0.09)	−0.17* (0.09)
Condition 5: received performance and equity prime and update	−0.30*** (0.09)	−0.24** (0.10)	−0.14 (0.09)
p-value from test of equality of coefficients for conditions 4 and 5	0.04	0.42	0.70
p-value from test of equality of coefficients for conditions 3 and 5	0.01	0.35	0.18
Observations	1031	1013	1018
Pseudo R^2	0.047	0.048	0.028

Ordered probit coefficients shown.
Models run on pooled sample from Conditions 2 through 5.
Models also condition on performance guess and control variables.
Standard errors in parentheses.
*$p < 0.10$, **$p < 0.05$, ***$p < 0.01$.

There are several comparisons of interest. A significant coefficient on condition 4 (the control group for the equity experiment) would suggest that receiving the equity *prime* (i.e. the question about race-related differences, but not the update) in addition to the overall performance prime affects opinions, because the only difference between conditions 2 and 4 is that condition 4 respondents were also asked to think about the black–white test score gap (but neither group was given updated information). A significant difference between the coefficients for conditions 4 and 5 would suggest that receiving the actual updated information about student performance changes one's opinion relative to simply being asked about performance and equity (without being told the actual performance). Lastly, a significant difference between conditions 3 and 5 (the two treatment groups) would suggest that receiving information about the achievement gap

because respondents are simultaneously updated on two dimensions at once. Recoding respondents as having prior beliefs that underestimated, overestimated or correctly assessed the true performance of the Tennessee education system on both the overall performance and the equity dimension and estimating models predicting institutional evaluations along with treatment interactions reveals no statistically significant results.

in addition to receiving information about overall performance changes the average response.

Table 4 summarises the main results (Table A.2 in the supplemental appendix contains the full results). Interestingly, respondents in condition 4 and who were primed to think about the achievement gap (but not updated) gave more negative institutional evaluations in two of the three models; the coefficient is negative but not significant at conventional levels in the third. Comparing the coefficients for conditions 4 and 5 reveals that the coefficients are statistically distinguishable only for evaluations of all Tennessee schools. The effect of being informed about actual student performance only matters for evaluations of Tennessee schools. This result is surprising given the effect of condition 3 in Table 4, which shows that receiving only the overall performance update negatively affects institutional evaluations.

Moreover, in only one model – again for all Tennessee schools – does receiving an informational update about overall performance and equity result in a different response from being updated about overall performance alone. Somewhat surprisingly, the coefficient for condition 5 is *smaller*, meaning that receiving additional information about the race-related achievement gap reduces the impact of being updated about overall performance.

We replicated this analysis for the six policy reforms. Supplemental appendix Table A.3 shows that (1) prior beliefs about the achievement gap are not clearly related to support for policy reforms and (2) receiving updated information has no effect.[22] Instead, opinions about educational reforms are best explained by partisan self-identification.

Overall, the effect of being updated on both overall student performance and the racial achievement gap is largely consistent with the effects of being informed only about overall performance. We cannot determine from these data why we find no effect for achievement gap information. One possibility is that citizens' opinions regarding education policy are not influenced by achievement gap concerns. Another is that receiving information on

---

[22] The results show some evidence that achievement gap information may be important for opinion formation for some policies. First, whereas receiving information about overall performance only (condition 3) is unrelated to support for any of the policies, for two of them—charter schools and vouchers—respondents are more likely to express support when given the update containing information about both performance and equity (the coefficients for conditions 4 and 5 are statistically different from one another). Moreover, for two of the policies—vouchers and charter schools—we can reject the null hypothesis that conditions 3 and 5 are the same. In both cases, the coefficient for condition 5 is more positive, suggesting that receiving information about equity (and performance) increases support for the policy relative to receiving only information about performance, which further suggests that respondents may view vouchers and charters as potentially effective interventions for reducing achievement gaps.

race-based gaps primes respondents to consider aspects of student achievement that are heavily correlated with family characteristics and thus beyond the control of schools. Alternatively, it could simply be that providing two pieces of information prevented respondents from being able to cognitively process the second piece.

Discussion and conclusion

Assessing citizens' responsiveness to new information is critical for determining the prospects for democratic accountability. Unless citizens change their opinions and beliefs in response to new information, it is hard to imagine how votes cast at the ballot box could reflect an informed assessment of public officials' performance and create the correct incentives for elected officials (Achen and Bartels 2002). Forming accurate beliefs, however, requires that citizens be able to appropriately update their existing beliefs in response to new information. If citizens update their beliefs in biased ways based on prior beliefs and partisan leanings, or if they fail to update their beliefs, the prospects for democratic accountability may be dim. This is a particularly disconcerting possibility in areas like education where the stakes are high and many reforms either implicitly or explicitly depend on public pressure to improve performance.

Our results suggest mixed implications from the perspective of democratic accountability. Despite Tennessee's long-standing emphasis on reporting educational outcomes and the large number of existing policies utilising such assessment information – which make it unlike many other states – most citizens overestimate student performance and therefore hold overly favourable assessments of the institutions responsible for education policy. However, consistent with prior work (e.g. James 2011; Chingos et al. 2012), we also find that citizens' assessments of educational institutions respond in seemingly rational ways to performance-related information. Not only are assessments driven by the level of student performance rather than ideological predispositions, but the institutions that are most closely associated with statewide performance are most affected by information about statewide performance.

To be clear, the fact that citizens' evaluations of educational institutions are responsive to performance-based information is not necessarily evidence of democratic accountability in action, because this is only the first step of what is required for policy performance-based accountability. Our experimental design allows us to cleanly identify the effect of learning about various dimensions of student performance and the possibility of priming, but we cannot determine whether the effects are transient or long lasting (but see, e.g. Chong and Druckman 2010). Moreover, even if citizens are completely informed and unsatisfied with educational performance, they

may be reluctant to punish elected officials and create required electoral consequences for a lack of performance in secondary education if they are content with other issues (James 2011). Creating the incentives necessary for democratic accountability in education policy may be difficult given the many possible dimensions of interest to citizens. Future research might link experimental data like ours to voter files or later survey responses about voting behaviour to assess whether information provision affects subsequent turnout or voting choices as a means of delving further into these processes. Of course, voting behaviour is but one mechanism through which citizens can express dissatisfaction with public institutions, and future research might examine the connections between performance information and other mechanisms, such as advocacy or residential or school mobility, as well.

The fact that citizens' opinions about particular policies designed to improve educational performance are unresponsive to learning about student performance under the status quo is also potentially sobering for the prospects of citizen-led policy change. Unlike the results in Henderson et al. (2014), learning that educational performance is worse than expected does not cause any change in the support for various policies aimed to increase student performance, perhaps reflecting the conflicting state of both the evidence base and policy discourse surrounding the effectiveness of these reform initiatives. Instead, citizens' opinions about education policies are primarily driven by ideological and partisan affiliations. These affiliations may represent prior beliefs that are too strong to be affected by the information we provide, or perhaps citizens base their policy opinions on beliefs or information unrelated to average student performance. In either case, consistent with the conclusions of other work highlighting the importance of elite messaging for public opinion change (e.g. Ladd and Lenz 2009; Noel 2013), changing the public's support for particular reforms would appear to depend on the actions of partisan and ideological leaders. Clearly, investigation of the factors underlying education policy opinion formation would be another fruitful avenue for research.

Acknowledgements

Supplementary material

To view supplementary material for this article, please visit http://dx.doi.org/10.1017/S0143814X14000312

References

Achen C. H. (1992) Social Psychology, Demographic Variables, and Linear Regression: Breaking the Iron Triangle in Voting Research. *Political Behavior* 14: 195–211.

Achen C. H. and Bartels L. M. (2002) Blind Retrospection: Electoral Responses to Droughts, Flu, and Shark Attacks. Presented at the Annual Meeting of the American Political Science Association, Boston, September 2002.

—— (2006) It Feels Like We're Thinking: The Rationalizing Voter and Electoral Democracy. Presented at the Annual Meeting of the American Political Science Association, Philadelphia, August 2006.

Ashworth S. (2012) Electoral Accountability: Recent Theoretical and Empirical Work. *Annual Review of Political Science* 15: 183–201.

Bartels L. M. (1996) Uninformed Votes: Information Effects in Presidential Elections. *American Journal of Political Science* 40: 194–230.

—— (2002) Beyond the Running Tally: Partisan Bias in Political Perceptions. *Political Behavior* 24(2): 117–150.

Berelson B., Lazarsfeld P. F. and McPhee W. N. (1954) *Voting: A Study of Opinion Formation and Presidential Campaigns*. Chicago: University of Chicago Press.

Berkman M. B. and Plutzer E. (2005) *Ten Thousand Democracies*. Washington, DC: Georgetown University Press.

Berry C. R. and Howell W. G. (2007) Accountability and Local Elections: Rethinking Retrospective Voting. *Journal of Politics* 69(3): 844–858.

Black S. E. (1999) Do Better Schools Matter? Parental Valuation of Elementary Education. *Quarterly Journal of Economics* 114(2): 577–599.

Bullock J. G. (2009) Partisan Bias and the Bayesian Ideal in the Study of Public Opinion. *Journal of Politics* 71(3): 1109–1124.

Campbell A., Converse P., Miller W. and Stokes D. (1960) *The American Voter*. Chicago, IL: University of Chicago Press.

Carnoy M. and Loeb S. (2002) Does External Accountability Affect Student Outcomes? A Cross-State Analysis. *Educational Evaluation and Policy Analysis* 24(4): 305–331.

Chingos M. M., Henderson M. and West M. R. (2012) Citizen Perceptions of Government Service Quality: Evidence from Public Schools, Washington, DC. *Quarterly Journal of Political Science* 7(4): 411–445.

Chong D. and Druckman J. N. (2007) Framing Theory. *Annual Review of Political Science* 10(1): 103–26.

—— (2010) Dynamic Public Opinion: Communication Effects over Time. *American Political Science Review* 104(4): 663–680.

Dahl R. A. (1989) *Democracy and Its Critics*. New Haven: Yale University Press.

Delli Carpini M. X. and Scott K. (1996) *What Americans Know about Politics and Why It Matters*. New Haven, CT: Yale University Press.

Dewey J. (1916) *Democracy and Education*. New York, NY: Macmillan.

Dorn S. (1998) The Political Legacy of School Accountability Systems. *Education Policy Analysis Archives* 6(1): 1–33.

Egan P. J. and Mullin M. (2012) Turning Personal Experience into Political Attitudes: The Effect of Local Weather on Americans' Perceptions about Global Warming. *Journal of Politics* 74(3): 796–809.

Erikson R. S. and Stoker L. (2011) Caught in the Draft: The Effects of Vietnam Draft Lottery Status on Political Attitudes. *American Political Science Review* 105(2): 221–237.

Favero N. and Meier K. J. (2013) Evaluating Urban Public Schools: Parents, Teachers, and State Assessments. *Public Administration Review* 73(3): 401–412.

Figlio D. and Loeb S. (2011) School Accountability. In Hanushek E. A., Machin S. and Woessmann L. (eds.) *Handbooks in Economics*, Vol. 3. The Netherlands: North-Holland, 383–421.

Figlio D. N. and Ladd H. F. (2008) School Accountability and Student Achievement. In Ladd H. F. and Edward B. F. (eds.), *Handbook of Research in Education Finance and Policy*. New York and London: Routledge, 166–182.

Figlio D. N., Maurice E. L. (2004) What's in a Grade? School Report Cards and the Housing Market. *American Economic Review* 94(3): 591–604.

Fiorina Morris P. (1981) *Retrospective Voting in American National Elections*. New Haven: Yale University Press.

Fryer R. G. Jr. and Levitt S. D. (2004) Understanding the Black-White Test Score Gap in the First Two Years of School. *Review of Economics and Statistics* 86(2): 447–464.

Gerber A., Green D. (1999) Misperceptions about Perceptual Bias. *Annual Review of Political Science* 2: 189–210.

Gilens M. (2001) Political Ignorance and Collective Policy Preferences. *American Political Science Review* 95: 379–396.

Green D., Palmquist B. and Schickler E. (2002) *Partisan Hearts and Minds: Political Parties and the Social Identities of Voters*. New Haven, CT: Yale University Press.

Guttman A. (1987) *Democratic Education*. Princeton: Princeton University Press.

Guillermo C., Truglia R. P. and Tetax M. (2011) Biased Perceptions of Income Distribution and Preferences for Redistribution: Evidence from a Survey Experiment. Working Paper of the Centro de Estudios Distributivos Laborales y Sociales [Center for Social, Labor and Distributional Studies] (CEDLAS), National University of La Plata.

Hanushek E. A. and Raymond M. E. (2005) Does School Accountability Lead to Improved Student Performance? *Journal of Policy Analysis and Management* 24(2): 297–327.

Healy A. J., Malhotra Neil and Mo C. H. (2010) Irrelevant Events Affect Voters' Evaluations of Government Performance. *Proceeding of the National Academy of Sciences* 29(107): 12804–12809.

Henderson M. B., Howell W. G. and Peterson P. E. (2014) Information Fuels Support for School Reform. *Education Next* 14(2): 27–35.

Hibbing J. R. and Alford J. R. (1981) The Electoral Impact of Economic Conditions: Who is Held Responsible? *American Journal of Political Science* 25(3): 423–439.

Highton B. and Kam C. D. (2011) The Long-Term Dynamics of Partisanship and Issue Orientations. *Journal of Politics* 73(1): 202–215.

Hochschild J. L. and Shen F. X. (2012) Race, Ethnicity, and Education Policy. *Oxford Handbook of Racial and Ethnic Politics in America*. New York: Oxford University Press.

Hutchings V. L. (2003) *Public Opinion and Democratic Accountability: How Citizens Learn about Politics*. Princeton: Princeton University Press.

Jacoby W. G (1988) The Impact of Party Identification on Issue Attitudes. *American Journal of Political Science* 32(3): 643–661.

James O. (2011) Performance Measures and Democracy: Information Effects on Citizens in Field and Laboratory Experiments. *Journal of Public Administration Research & Theory* 21(3): 399–418.

Key V.O. (1966) *The Responsible Electorate: Rationality in Presidential Voting 1936–1960*. Cambridge: Harvard University Press.

Kuklinski J. H., Quirk P. J., Jerit J. and Rich R. F. (2001) The Political Environment and Citizen Competence. *American Journal of Political Science* 45(2): 410–424.

Ladd J. M. and Lenz G. S. (2009) Exploiting a Rare Communication Shift to Document the Persuasive Power of the News Media. *American Journal of Political Science* 53(2): 394–410.

Lenz G. S. (2009) Learning and Opinion Change, Not Priming: Reconsidering the Priming Hypothesis. *American Journal of Political Science* 53(4): 821–837.

Lodge M., Steenbergen M. R. and Brau S. (1995) The Responsive Voter: Campaign Information and the Dynamics of Candidate Evaluation. *American Political Science Review* 89(2): 309–326.

Loeb S. and Strunk K. O. (2007) Accountability and Local Control: Response to Incentives with and without Authority over Resource Generation and Allocation. *Education Finance and Policy* 2(1): 10–39.

MacKuen M. (1984) Exposure to Information, Belief Integration, and Individual Responsiveness to Agenda Change. *American Political Science Review* 78(2): 372–391.

Malhotra N. and Kuo A. G. (2008) Attributing Blame: The Public's Response to Hurricane Katrina. *The Journal of Politics* 70(1): 120–135.

Markus G. B. (1992) The Impact of Personal and National Economic Conditions on Presidential Voting, 1956-1988. *American Journal of Political Science* 36(3): 829–834.

McGuire W. J. (1969) The Nature of Attitudes and Attitude Change. In Lindzey G. and Aronson E. (eds.), *Handbook of Social Psychology*. Reading, MA: Addison-Wesley, 136–314.

Noel H. (2013) *Political Ideologies and Political Parties in America*. New York: Cambridge University Press.

Popkin S. L. (1994) *The Reasoning Voter: Communication and Persusasion in Presidential Campaigns*. 2nd ed. Chicago: University of Chicago Press.

Prior Markus (2007) *Post-Broadcast Democracy*. NY: Cambridge University Press.

Rahn W. M. (1993) The Role of Partisan Stereotypes in Information Processing about Political Candidates. *American Journal of Political Science* 37(2): 472–496.

Rudolph T. J. (2003) Who's Responsible for the Economy? The Formation and Consequences of Responsibility Attributions. *American Journal of Political Science* 47(4): 698–713.

Schneider M., Teske P. and Marschall M. (2002) *Choosing Schools: Consumer Choice and the Quality of American Schools*. Princeton, NJ: Princeton University Press.

Sher A. (2011) Governor Haslam, Tennessee Education Commissioner Defend Teacher Evaluation System. *Chattanooga Times Free Press*, http://timesfreepress.com/news/2011/sep/29/governor-education-commissioner-defend-teacher/?news (accessed 14 May 2012).

Sides J. and Citrin J. (2007) European Opinion About Immigration: The Role of Identities, Interests and Information. *British Journal of Political Science* 37: 477–504.

Stein R. M. (1990) Economic Voting for Governor and U.S. Senator: The Electoral Consequences of Federalism. *Journal of Politics* 52(1): 29–53.

Stroud N. J. (2009) Media Use and Political Predispositions: Revisiting the Concept of Selective Exposure. *Political Behavior* 30(3): 341–366.

Woon J. (2012) Democratic Accountability and Retrospective Voting: A Laboratory Experiment. *American Journal of Political Science* 56(4): 913–930.

Zaller J. R. (1992) *The Nature and Origins of Mass Opinion*. New York: Cambridge University Press.

Zaller J. and Feldman S. (1992) A Simple Theory of the Survey Response: Answering Questions Versus Revealing Preferences. *American Journal of Political Science* 36: 579–616.

Zechman M. J. (1979) Dynamic Models of the Voter's Decision Calculus: Incorporating Retrospective Considerations into Rational Choice Models of Individual Voting Behavior. *Public Choice* 34(3–4): 297–315.