Polit Behav (2015) 37:933–954
DOI 10.1007/s11109-014-9298-9

ORIGINAL PAPER

Mind the Gap: Why Large Group Deficits in Political Knowledge Emerge—And What To Do About Them

Efrén O. Pérez

Published online: 10 January 2015

Abstract Large group differences in political knowledge raise strong concerns about lasting inequities in U.S. politics. I argue such gaps emerge when factual questions operate unevenly across diverse populations, thereby inflating actual knowledge differences between groups. I illustrate this by revisiting the large knowledge deficit often observed among Latinos relative to Whites. Using a survey with traditional factual questions (e.g., what office is held by John Roberts?) and new Latino-themed items (e.g., what office is held by Marco Rubio?), I show that Whites are more likely than Latinos to correctly answer many conventional questions due to item features that are unrelated to people's level of knowledge (i.e., item bias). Latino-themed questions, however, do not display these extraneous characteristics. Consequently, Whites and Latinos have equal odds in correctly answering these latter items, provided they have the necessary knowledge level. I also show how such item quality differences matter in practical terms. Accordingly, I establish that using a scale of unbiased items reduces the Latino-White knowledge deficit from 31 to 8 %, with the latter gap more fully explained by individual differences in established correlates of political knowledge. I then show that uneven item performance distorts inferences about knowledge's influence on mass opinion. I discuss the implications of these results for political knowledge's conceptualization and measurement in an increasingly diverse U.S. polity.

Keywords Political knowledge · Content validity · Latino politics · Item response theory

Electronic supplementary material The online version of this article (doi:10.1007/s11109-014-9298-9) contains supplementary material, which is available to authorized users.

E. O. Pérez (✉)
Department of Political Science, Vanderbilt University, PMB 505, 230 Vanderbilt Place, Nashville, TN 37203, USA
e-mail: efren.o.perez@vanderbilt.edu

Introduction

Many scholars deem political knowledge vital to U.S. democracy (Zaller 1992; Verba et al. 1995; Gay 2013). As Delli Carpini and Keeter (1996, p. 1) explain, "democracy functions best when its citizens are politically informed" (cf. Bartels 1996; Althaus 1998). Yet mounting evidence suggests some groups are more informed than others, thereby raising strong concerns about political inequities. Whites appear more knowledgeable than non-Whites (Verba et al. 1995). Men seem more informed than women (Mondak and Anderson 2004). And, older people report more knowledge than the young (Delli Carpini and Keeter 1996). Scholars have explained these patterns by, *inter alia*, holding constant differences in the correlates of political knowledge (Mondak 1999) and minimizing the role of item response format in yielding such disparities (Mondak and Anderson 2004). But significant group deficits in political information persist despite such efforts, thus shaping characterizations of who knows more about politics.

I argue that large group differences in political knowledge often arise from the uneven performance of factual questions in inter-group settings. Factual questions are designed to tap people's broad sense of politics (Delli Carpini and Keeter 1996; Mondak 2001), but not all political facts are created equally. Akin to culturally biased items in education tests (e.g., the SAT) (Camilli and Shepard 1994; Hambleton et al. 1991), two people who are equally knowledgeable about politics can score differently on an item if it broaches a fact with varied salience for them. If this occurs, then any knowledge gap between groups will be artificially inflated, making some groups seem less informed than they are. Thus, I propose designing and testing factual questions that balance the goal of capturing knowledge with a need for inter-group comparability. In doing so, I build on recent work arguing for richer mixes of items in knowledge scales as a way to ameliorate information gaps between groups (Dolan 2011). Specifically, I show how scholars can statistically identify items that validly capture knowledge across diverse populations, thereby contributing to research on improved appraisals of this trait and scholars' sense of its distribution in the mass public (cf. Prior 2013; Gibson and Caldeira 2009; Prior and Lupia 2008; Lupia 2006; Sanbonmatsu 2003; Mondak and Davis 2001; Mondak 2001; Luskin 1987).

I demonstrate my framework by re-assessing the knowledge gap between Latinos and Whites. As the largest U.S. ethnic minority, Latinos are a growing segment of opinion polls, with studies often finding they are much less politically informed than Whites (Abrajano 2010; Abrajano and Alvarez 2010). Given this trend, I commissioned a national survey of Latinos and Whites asking standard knowledge questions (e.g., what office is held by John Roberts?) and new Latino-themed items (e.g., what office is held by Sonia Sotomayor?). I designed the latter to show that enhancing knowledge's appraisal among a racial minority does not entail degrading its assessment among a racial majority. Thus, all eight (8) of my items reflect the common conceptualization of political knowledge as "facts about the rules of the game, the substance of politics, and people and parties…central to…politics" (Delli Carpini and Keeter 1996, p. 294).

Consistent with prior work, I find that Latinos are much less politically informed than Whites when answering my standard knowledge items (Verba et al. 1995; Abrajano 2010). But when replying to my Latino-themed questions, this gap narrows significantly. This does not occur because the latter items favor Latinos. It happens because these items operate without bias *across* both groups. Indeed, Latinos and Whites are equally likely to answer the Latino-themed questions correctly, if they have the requisite level of political information (Embretson and Reise 2000; Hambleton et al. 1991). This implies that while Whites are more knowledgeable than Latinos, the latter are more informed than traditional knowledge items indicate.

Seizing on these measurement results, I build a knowledge scale with item bias unaddressed and another one correcting it. Across both scales, the Latino-White knowledge gap drops from 31 to 8 %. Holding constant individual differences in established correlates of knowledge (e.g., age, education) closes this smaller gap. I also show that item bias distorts inferences about knowledge's impact on mass opinion. I then discuss my results' implications for the conceptualization and measurement of political knowledge in a diversifying U.S. public.

Political Knowledge: Conceptualization and Assessment

Political knowledge is conceptualized as a single trait where a person's grasp of some aspects about politics is robustly correlated with their hold over others (cf. Zaller 1986; Price and Zaller 1993). More precisely, it is people's "range of factual information about politics stored in long-term memory," where the scope of facts covers the rules, people, and substance of politics (Delli Carpini and Keeter 1996, p. 294; Luskin 1987). Scholars often assess political knowledge with factual items like *Whose responsibility is it to determine if a law is constitutional or not…?*, with correct/incorrect answers normally combined into summated scales (cf. Luskin and Bullock 2011; Sturgis et al. 2008). Moreover, scale items are often closed- rather than open-ended (but see Prior and Lupia 2008). Since the former ask one to choose replies from alternatives rather than report an unstructured answer, they are deemed less cognitively taxing (e.g., Lupia 2006) and less prone to subjective coding when verbatim responses are quantified (e.g., Sanbonmatsu 2003). For example, Gibson and Caldeira (2009) show that only 10.5 % of respondents in the 2000 American National Election Study (ANES) correctly identified U.S. Chief Justice William Rehnquist via an open-ended item testing for knowledge about the Supreme Court—a result partly arising from the strict coding criteria for correct replies to this question. Yet these authors show that knowledge of the high court noticeably improves with use of closed-ended items.

While some scholars have sharply critiqued factual political items (Mondak 2001; Mondak and Davis 2001; Lupia 2006; Prior and Lupia 2008), they remain widely used because despite relevant criticisms, they are a direct and inexpensive way to assess political knowledge in surveys. Moreover, ample evidence confirms that factual items capture a generalized sense of knowledge (Luskin 1987; Delli Carpini and Keeter 1996; Luskin and Bullock 2011). For example, Zaller (1986)

shows that factual items on wide topics like race, economics, and foreign policy often tap a single variable in common (cf. Price and Zaller 1993; Zaller 1990).

But this prevailing tide of evidence is awash in some anomalies. Iyengar (1986) finds that factual items on diverse political topics sometimes capture domain-specific knowledge, with people being informed about some subjects but not others (cf. Hutchings 2001; McGraw and Pinney 1990). A few studies even find that domain-specific knowledge outperforms general knowledge when one's knowledge and judgment share a domain (Iyengar 1990; McGraw and Pinney 1990; Sanbonmatsu 2003). Furthermore, other scholars show that questions about particular programs and problems actually capture *policy-specific* knowledge, which is politically consequential (Gilens 2001; Kuklinski et al. 2000). "Although it has been more common to study civic facts," write Barabas and Jerit (2009, p. 73), "policy-specific knowledge appears to be more powerful than the former in shaping…political judgments."

Despite these nuances, most research finds the performance "of domain specific measures of information is…modest and uneven" (Zaller 1986, p. 2; Price and Zaller 1993; Delli Carpini and Keeter 1996). So, the modal way to gauge knowledge is by asking about facts regarding rules, policies, and figures in national politics (Luskin and Bullock 2011; Mondak 2001; Verba et al. 1995). Yet brushing aside the preceding anomalies seems counterproductive, as those studies imply that general knowledge scales might be *too* general, thereby mislabeling some people as politically uninformed (Dolan 2011). This matters for knowledge assessments in inter-group settings because general scales often reveal large information gaps. Thus, improving evaluations of political information in such contexts might be achieved, in part, by finding domain- and policy-specific items that also tap a general and shared sense of knowledge across groups.

Group Differences in Political Knowledge

Using general knowledge scales, scholars have established that many Americans know very little about politics (Delli Carpini and Keeter 1996; Luskin 1987), a trend refracted along group lines (Althaus 1998; Abrajano 2010). Consider Verba et al. (1995), who find that on a 0–8 knowledge scale, Whites answer about 1 more item correctly than Blacks. Although some group differences in knowledge are to be expected, what puzzles scholars is their size and persistence, not their presence. With a 0–5 scale, for example, Mondak (1999) shows that Blacks score lower than Whites even after controlling for several correlates of knowledge (e.g., education, political efficacy), with similar patterns arising between men and women, the wealthy and poor, and older and younger people (Delli Carpini and Keeter 1996; Mondak and Anderson 2004; Dolan 2011).

These large and durable group deficits in knowledge have sparked many attempts to explain them. One effort, alluded to above, assumes these gaps arise from omitted variable bias, where a knowledge gap between groups is driven by other differences among group members. Holding constant these attributes should thus eliminate group deficits in knowledge. Yet sizeable gaps remain in models ranging from

parsimonious to complex (Verba et al. 1995; Delli Carpini and Keeter 1996; Mondak and Anderson 2004), making this an insufficient explanation.

Other work pins knowledge gaps to how individual characteristics translate to political information within groups. Using ANES data, Dow (2009) shows that one-third of the knowledge gender gap arises from standard knowledge predictors. The remaining deficit arises from group differences in education's link to knowledge: "[a]ll else equal, women must possess associate level college degrees to obtain…the same level of political knowledge as men possessing high school diplomas" (Dow 2009, p. 132). Nevertheless, the theoretical rationale behind these patterns is not fully illuminated (e.g., why education yields unequal knowledge by gender?). Moreover, the uneven operation of knowledge predictors across groups suggests knowledge comparisons may involve "apples and oranges" (cf. Stegmueller 2011).

A third approach traces group deficits in political information to item format (Mondak 1999, 2001; Mondak and Davis 2001; Lupia 2006; Prior 2013). Mondak and Anderson (2004) argue that allowing "don't know" (DK) responses inflates knowledge gaps between groups. Since men are more likely than women to hazard answers rather than reply DK, males seem more knowledgeable than females. In fact, removing the DK option from knowledge items reduces (but does not eliminate) this gender gap. In a similar vein, Prior (2013) shows that joining visuals to knowledge questions modestly boosts information levels among people who often display lower political knowledge (e.g., women, older people). Still, the persistence of some of these gaps, despite such corrections, suggests more fundamental measurement issues.

A fourth research area finds that interview settings shape knowledge gaps. Davis and Silver (2003), for example, show that African Americans are less likely than Whites to answer knowledge items correctly if they are surveyed by a White rather than Black interviewer. This occurs because being asked by Whites to answer factual questions leads Blacks to worry about confirming a stereotype of their group as unintelligent and uninformed, with the heightened anxiety impairing question completion (i.e., stereotype threat). More recently, Prior and Lupia (2008) have revealed another situational influence on political knowledge by showing that giving people more time or money for correct replies boosts information levels by 11–24 %. This contextual malleability of knowledge reports, however, raises thorny questions about the reactivity of knowledge items and/or the stability of the trait they measure.

Equal Knowledge, Unequal Questions

I claim that knowledge gaps are often larger than they should be due to variation in how difficult factual items are for members of some groups to answer correctly. Hard items are needed to distinguish between people who know a lot about politics from those who know less. But to do this, a factual question must yield an estimate of people's knowledge that is not confounded with anything else. When items are hard due to features extraneous to knowledge—e.g., a fact's low salience among some people—two identically informed persons will have varied odds of correctly

answering the same question (Brown 2006). An example of such bias involves a Scholastic Assessment Test (SAT) item, where students completed the analogy *Runner is to marathon*: (a) *Envoy is to embassy*; (b) *Martyr is to massacre*; (c) *Oarsman is to regatta*; (d) *Referee is to tournament*; or (e) *Horse is to stable*. While the right answer is 'c', research found this item favored White students, whose cultural milieu was more likely to expose them to the meaning of *regatta*: a race with rowboats or similar vessels (Chideya 1995). Thus, even if Black and White students had identical analogical skills, the former were less likely to correctly answer this item, which made Black students seem less apt than they are.

I suggest a comparable challenge can plague knowledge items. Specifically, I argue that group deficits in observed knowledge levels are artificially enlarged when political facts have varied salience between groups. Factual items are meant to tap political knowledge, a latent trait (Luskin 1987). This means a person's observed score (y_i) on a factual item is conditional on their true knowledge score (η) and nothing else. When $F(y_i|\eta)$ holds, observed knowledge gaps reflect true knowledge differences. But if an item is biased against a group, a person's response to an item is conditional on their knowledge, plus group membership (g_i)—that is, $F(y_i|\eta, g_i)$.

I trace this challenge to lingering uncertainty about the *content validity* of knowledge scales. A scale has content validity if its items are representative of a given domain (DeVellis 2003; Zeller and Carmines 1980). Since the universe of what people *could* know about politics is vast and unknown, scholars have delimited knowledge to "facts about the rules of the game, the substance of politics, and people and parties…central to…politics" (Delli Carpini and Keeter 1996, p. 294). This has yielded factual items that scale well. But such items might also tap aspects of politics that some groups attend to more, thus letting subtle biases creep in (Dolan 2011; Lupia 2006). Consider the question *How much of a majority is required for the U.S. Senate and House to override a presidential veto?* On average, this item will be within reach of a White person with at least a high school education, since it broaches a civics fact often covered in U.S. secondary schools. Yet the same item will be harder for an equally educated Latino who was partially or completely schooled outside the U.S., a likely prospect given the large number of foreign-born Latinos (Alvarez and Abrajano 2010). While this item will show that the White individual is more informed than the Latino person, the gap is inflated by the item's greater salience among the former, making Latinos seem much less politically informed than Whites.

Hence, I contend that factual items should be harmonized across groups before making strong claims about knowledge gaps between groups. This entails fielding richer question sets (Dolan 2011; Lupia 2006), while determining which items actually perform impartially across groups. The goal is not to replace knowledge items that are "stacked against" one group answering correctly for items that are "stacked in favor" of another. Rather, it is to design items that adhere to a common formulation of political knowledge *and* perform equally across groups. If this standard is met, then members of each group will have equal odds of answering a knowledge question correctly, provided they have the same knowledge level.

One way to meet these goals is by developing new domain-specific (cf. Dolan 2011; Iyengar 1986) and policy-specific items (cf. Barabas and Jerit 2009; Gilens

2001) that fall under knowledge's traditional conceptualization (Delli Carpini and Keeter 1996). Fielding wider sets of knowledge questions enables researchers to statistically identify items that perform equally across groups, thereby producing scales that capture real political information differences. For example, scholars often ask what office is held by John Boehner, the speaker of the U.S. House of Representatives, since he is a visible figure in national politics (cf. Mondak 2001). But using this logic, why not also query people with a domain-specific item that follows knowledge's traditional definition, such as what office is held by Marco Rubio, Florida's Hispanic Republican Senator? Like John Boehner, Marco Rubio has been prominent in national political discourse (e.g., as a contender for the 2016 Republican presidential nomination) (Grunwald 2013; Metzler 2013). In addition, if policy-specific knowledge is consequential (Gilens 2001), then why not ask about information that is widely salient? Taking a cross-cutting domain like immigration (Tichenor 2002), for instance, one might ask which government branch recently changed policy toward undocumented immigrants brought to the U.S. as minors (i.e., Executive)?

Of course, the kinds of items I am suggesting potentially capture a subtype of general knowledge. Consider that while John Boehner and Marco Rubio are both visible politicians, a House Speaker is different than a Florida Senator. This challenge is typically faced by studies measuring knowledge with new items that seem to fall beyond the range of recommended ones (Delli Carpini and Keeter 1996, p. 305), such as those testing for foreign leader identification (e.g., What position is currently held by Nicolas Sarkozy?) (Prior 2013). Thus, I treat as an empirical question whether any new items I propose tap the same knowledge type as traditional items.

To this end, I evaluate an enriched bank of factual questions designed to measure political knowledge among Latinos and Whites. Prior research shows that Latinos are much less politically informed than Whites (Abrajano 2010; Abrajano and Alvarez 2010). For example, Verba et al. (1995) find that on a 0–8 scale, Whites and Latinos averaged 4.1 and 2.7 correct items, respectively. Such gaps often remain despite controlling for covariates like education and political efficacy (e.g., Mondak and Anderson 2004). Whites and Latinos also vary widely by demographic and political attributes, which casts doubt on knowledge item performance across such diverse groups. For example, Latinos are younger and less educated than Whites: attributes which undercut knowledge (Delli Carpini and Keeter 1996). Latinos are also more likely than Whites to be foreign-born or have foreign-born parents, which affects one's political socialization (Branton 2007). It thus remains unclear "whether generic measures…are adequate for studying the political knowledge of… Hispanics" (Abrajano and Alvarez 2010, p. 117).

Data

My analysis draws on a 10-minute online survey of Latino ($n = 505$) and White ($n = 559$) adults. The study was run by GfK (formerly Knowledge Networks), which maintains a nationally representative Internet panel. Panelists are recruited via probabilistic sampling and furnished with web access and computer hardware if

needed.[1] GfK fielded this study from March 28 to April 12, 2013. Latino respondents interviewed in English or Spanish.

Besides five (5) traditional knowledge items, the survey asked three (3) new Latino-themed questions.[2] Item order was randomized across respondents, with the order of response options randomized within items. Table A (supplementary appendix—SA) lists the item wordings. All eight items focus on national political figures, institutions, or rules (Delli Carpini and Keeter 1996), but the last three touch on Latino themes, which lets me examine how a knowledge deficit between groups is shaped by traditional items and how it can be improved via new questions. The Latino-themed items ask what office is held by Marco Rubio, Florida's Republican Cuban American Senator; what office is held by Sonia Sotomayor: the U.S. Supreme Court's first Hispanic justice; and what institution granted a stay of deportation to undocumented immigrants who entered the U.S. as children—an executive order from President Obama.[3]

This mix of eight items includes general knowledge items (e.g., *What job does John Boehner hold?*), items I consider domain-specific (e.g., *What job does Marco Rubio hold?*), and one item I deem policy-specific (*What institution granted a stay of deportation to some undocumented immigrants?*). Since I designed some of my Latino items to parallel traditional ones, half of my items ask about political figures (e.g., it is hard to design a rule-based item that is Latino-specific). Although space constraints prevented me from fielding a wider question set, these items are sufficient to identify which capture knowledge impartially across both groups.[4]

One last feature of these items bears mentioning: they include a "don't know" (DK) option. Factual questions have often been scored by pooling DKs with incorrect answers. Yet work in the early 2000s began advising against using a DK option (cf. Mondak and Davis 2001; Mondak 2001), in part, on the grounds that DKs contain non-trivial amounts of hidden knowledge (e.g., being somewhat sure, but not certain of a correct answer). More recently, Luskin and Bullock (2011) have shown that while discouraging DKs boosts the number of correct answers on closed-

[1] Specifically, GfK draws survey samples from its online respondent panel, whose members are recruited through an address-based sample (ABS) frame, which probabilistically samples addresses from the U.S. Postal Service's Delivery Sequence File. Randomly sampled addresses are invited to join GfK's panel, which yields coverage of about 97 % of U.S. households. GfK supplements the number of Latino panelists recruited this way via random digit dialing (RDD) in high Latino-density census blocks (author's personal communication with GfK).

[2] I wished to limit how much subjectivity seeped into my item pool, so I did not field any open-ended items because these are prone to idiosyncratic coding criteria (e.g., Gibson and Caldeira 2009).

[3] Respondents did not systematically use the Internet to answer my items. GfK informed me that respondents typically answer a survey item within 20 s. The median response times (in seconds) for seven (7) of the knowledge items fall below this threshold among Latinos (L) and Whites (W). Only *Stay of deportation* slightly surpasses this threshold (L: 27 s, W: 22 s).

[4] One might wonder why I do not use other common items (e.g., Which party had the most members in the House of Representatives in Washington before the elections last month?). There is actually little guidance on how many and which items to include in knowledge scales. Delli Carpini and Keeter (1996) suggest a 5-item index with items varying by what they ask and how hard they are to answer. Yet there is often a weak tie between the types and number of items that they recommend and those used by other prominent researchers (cf. Mondak 2001; Prior and Lupia 2008; Luskin and Bullock 2011). For this reason, I rely on the reasoning described in the text.

ended items, such gains are mostly due to lucky guessing, which implies that "don't know" really means "don't know."[5] Given this development, my items retain a DK option, which I score in the traditional way. This has the added benefit of letting me explore whether any improved item performance is related, in some measure, to DK reductions.

Item Bias Influences Question Quality and Performance

Table 1 reports the reliabilities and mean group scores for three (3) political knowledge scales. Scale 1 is comprised by my five (5) traditional knowledge items (Table A, SA). Scale 2 expands this list by adding the three (3) Latino-focused items. And, scale 3 is comprised exclusively by the three (3) Latino-themed questions. On all scales, higher values reflect more items answered correctly—i.e., greater political knowledge.

All three scales are reasonably reliable, with (α) between 0.61 and 0.80.[6] Moreover, the knowledge gap between groups varies by index. On scale 1, Latinos average 2.12 correct answers while Whites average 3.65, for a Latino deficit of 1.53 correct items. Standardizing this difference yields a Cohen's $d = 0.88$, which is a large gap ($d \approx 0.20$, 0.50, and 0.80 are small, medium, and large, respectively) (Cohen 1988). Scale 2 shrinks this difference. Here Latinos correctly answer 3.71 items on average, while Whites answer 5.30 correct items, for a Latino deficit of 1.59 items (Cohen's $d = 0.64$). No reliable gap emerges on scale 3 (Cohen's $d = 0.05$).

These results suggest some of my items reduce the knowledge deficit between both groups. But absent finer information, I cannot tell why certain items decrease this gap. In particular, it is important to establish that any narrowing of this difference is due to items that operate impartially across both groups, rather than items that are easier for Latinos but harder for Whites. Thus, I submit my items to a confirmatory factor analysis (CFA), where I estimate a one-factor model within each group. The parameters of interest here are (λ), an item's loading or discrimination; and (τ), an item's threshold or difficulty. Loadings (λ) disclose how well an item distinguishes between people at a given knowledge level (θ), while thresholds (τ) reveal how much knowledge (θ) is needed for a 0.50 probability of a correct response.[7] I use these parameters to create Item Characteristic Curves (ICCs) linking knowledge to each question's loading and threshold, which allows me to visually inspect item performance across both groups.

I report the raw results from my factor analysis in Table C (SA). Here I simply note that this measurement model displays great fit according to several criteria

[5] This last part of the sentence is a paraphrase of the title to Luskin and Bullock's (2011) article.

[6] For scale 1, $\alpha = 0.71$ (Whites) and $\alpha = 0.72$ (Latinos). For scale 2, $\alpha = 0.80$ (Whites) and $\alpha = 0.77$ (Latinos). And, for scale 3, $\alpha = 0.63$ (Whites) and $\alpha = 0.60$ (Latinos). But note that α is the lower bound of a scale's reliability—and it is sensitive to the number of scale items (Sijtsma 2009). At any rate, classical test theory suggests that random measurement error in my knowledge items will average out to zero (or close to it) when I scale them (Spector 1992).

[7] I discuss in more detail the mechanics behind my factor analysis in section B (SA).

Table 1 Observed scale reliability and raw knowledge gaps, Latinos and Whites

	Alpha (α)	Pooled mean	Mean: Latinos (L)	Mean: Whites (W)	Raw difference: (L − W)	Difference as Cohen's *d*
Scale 1 (5 traditional items only)	0.77	2.92	2.12	3.65	−1.53[a]	0.88
Scale 2 (5 traditional + 3 Latino items)	0.80	4.54	3.71	5.30	−1.59[a]	0.64
Scale 3 (3 Latino items only)	0.61	1.62	1.59	1.65	−0.05	0.05

Scale 1 runs from 0 to 5 items correct. Scale 2 runs from 0 to 8 correct items. Scale 3 runs from 0 to 3 correct items. Means and differences in means are in item units. Cohen's *d* is a standardized mean difference produced by dividing a raw mean difference by its pooled standard deviation. Cohen's *d* values around 0.20, 0.50, and 0.80 are considered small, medium, and large, respectively

[a] Difference reliable at 5 % level or better (two-tailed)

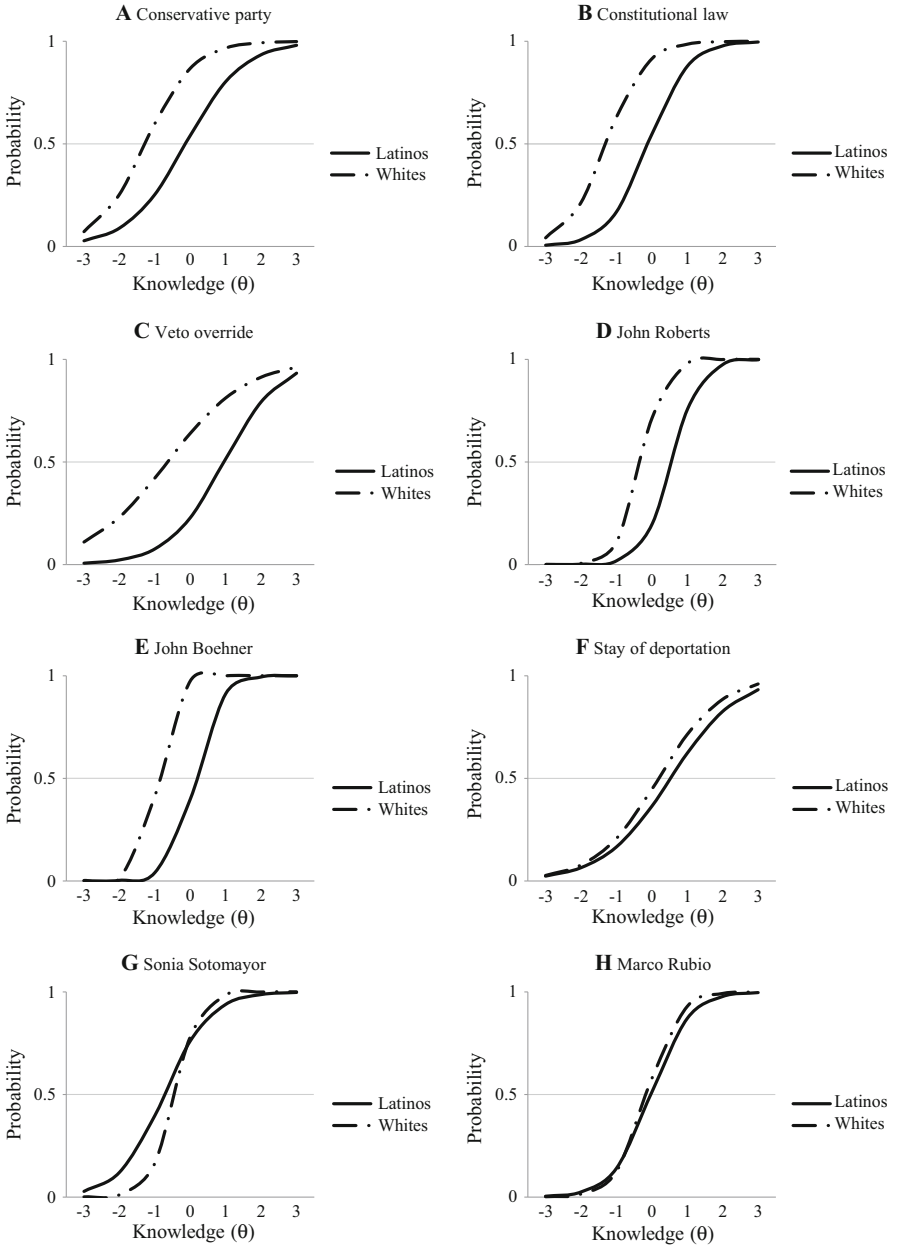

Fig. 1 Item characteristic curves for political knowledge items: Latinos and Whites

(e.g., CFI and TLI near their 1.00 maximum; RMSEA well below the 0.10 cutoff) (Brown 2006). This model's quality is further affirmed by an absence of large residuals and substantively large item loadings (all standardized $\lambda > 0.40$, as recommended;

see Brown 2006). Consider the item *Conservative party*. Among Latinos, a unit shift in latent political knowledge yields a 0.591 standard deviation shift in the response variable underlying this item. For Whites, the corresponding shift is 0.657. Similarly robust relationships characterize the other items. Finally, Table D (SA) reveals that all eight items capture general political information and not two variants of domain-specific knowledge (i.e., "traditional" versus "Latino" knowledge).

Using the thresholds and loadings from this factor analysis, I produce item characteristic curves (ICCs) for each item, which I display in Fig. 1, panels A through H. Thresholds with positive (negative) values indicate harder (easier) items, as it takes more (less) knowledge (θ) to answer such items. Loadings with higher positive values indicate an item more effectively distinguishes between individuals at a given threshold, as evidenced by an ICC's steeper slope.[8]

While the items discriminate similarly well among respondents, most items vary sharply by their thresholds. Indeed, for most items the ICC for Latinos is to the right of Whites, which suggests such items are more difficult for Latinos to correctly answer. For example, among Latinos, the item *John Roberts* yields a threshold of 0.547, a relatively hard item. Yet among Whites, the same item produces a threshold of -0.311, a relatively easier item. Comparable patterns emerge for *Veto override* and *John Boehner*. Furthermore, even when items that are relatively easy (hard) for both groups and discriminate well at their respective levels, the absolute differences in these parameters seem large, as illustrated by *Constitutional law*.

The exceptions to these patterns appear to be the Latino-themed items, which have ICCs that are remarkably similar across groups, thus suggesting impartial performance across both groups. In fact, unlike most of the traditional items, none of the Latino-themed questions seems to be appreciably harder for Whites than Latinos. This pattern does not arise because these items are more accessible to Latinos. On two of these questions, there is no reliable difference in the proportion of DKs between both groups (*Stay of deportation*: Whites = 30 %, Latinos = 29 %, $p < 0.91$; *Sonia Sotomayor*: Whites = 28 %, Latinos = 26 %, $p < 0.37$). On the third of these items, there is a reliable difference in DKs, but it is Whites—not Latinos—who report the lower DK proportion (*Marco Rubio*: Whites = 35 %, Latinos = 44 %, $p < 0.01$).[9]

I formally assess the visual patterns in Fig. 1 by statistically testing which items operate uniformly across both groups. I do this by comparing two measurement models. The first is the one that yielded the input for the ICCs I just analyzed (Table C, SA). There the item loadings and thresholds were freely estimated within each group. I compare this unrestricted model to a series of restricted ones that fix to equality these parameters for each group on an item-by-item basis.[10] If an equality

[8] Following Brown (2006), I transform the raw results in table C (SA) into discriminations/difficulties, which underlie Fig. 1's ICCs. Table E (SA) reports the raw CFA loadings/thresholds and their IRT analogs. As expected, these quantities are similar to each other.

[9] In fact, the Latino-White DK gap is smaller on the Latino-themed items than on the traditional ones. Specifically, the mean DK gap for the traditional items is 19 %. For the Latino items, the mean DK gap is 6 % (Table F, SA).

[10] This is akin to a *partial invariance* analysis (Byrne et al. 1989). A *full invariance* test is unnecessary here since Fig. 1 already shows that many items do not perform uniformly across Latinos and Whites.

constraint fails to produce a deteriorated model fit, then I have evidence that a given item performs equivalently across groups. I further explain the mechanics of this analysis and its raw output in Table G (SA). In the interest of space, I report here that out of eight items, four (4) fail to display statistical evidence of unequal performance across Latinos and Whites: *Conservative party*, *Stay of deportation*, *Sonia Sotomayor*, and *Marco Rubio*.

Item Bias Distorts Knowledge Levels

I find that many knowledge questions operate unevenly across Latinos and Whites. But what do scholars gain (lose) from this insight? I answer this by examining each group's mean score on two knowledge scales. The first scale leaves item bias unaddressed. As an index of five (5) traditional factual items, it reflects the typical approach to measuring knowledge. The second scale uses the four (4) items that operate impartially across Latinos and Whites: *Conservative party*, *Stay of deportation*, *Sonia Sotomayor*, and *Marco Rubio*.[11]

Table 2 shows Latinos underperform by 1.53 correct questions on the un-harmonized scale of traditional items, which is a large gap (Cohen's $d = 0.88$). On the harmonized scale, however, this deficit dwindles to 0.33 correct items (Cohen's $d = 0.25$).[12] In other words, Latinos still possess less political knowledge than Whites, but the gap between them is less a chasm and more a fissure. In fact, the smaller knowledge deficit between both groups is more fully explained by individual differences in demographic, social, and political characteristics among Latinos and Whites. Table 3 reports two analyses where each knowledge scale is regressed on key correlates, including education and political interest. These models also control for individual differences in generational status, citizenship status, and interview language to grapple with the heterogeneity introduced by Latinos in my sample. All variables here have a 0–1 range.

On the un-harmonized scale, individual differences in political knowledge correspond with individual differences in several variables, including education, income, and political interest. Yet after holding constant differences in these and other correlates, Latinos still display about 7 % less knowledge than Whites. In contrast, there is no significant knowledge difference between Latinos and Whites when using the harmonized scale.[13] That is, being Latino is no longer a reliable predictor of political knowledge. Instead, any differences on the harmonized

[11] In Table 1, the un-harmonized scale of five (5) traditional items displays a reliability of $\alpha = 0.77$. But since most of those items perform unequally across Latinos and Whites, this estimate is artifactual. In contrast, the harmonized scale of four (4) items displays a reliability of $\alpha = 0.62$, which is reasonably high given the use of fewer items.

[12] Un-harmonized scale: $\alpha = 0.66$ (Whites) and $\alpha = 0.67$ (Latinos). Harmonized index: $\alpha = 0.64$ (Whites) and $\alpha = 0.60$ (Latinos). Since the un-harmonized scale has one more item than the harmonized index, I estimate five (5) separate alphas, each time dropping one of the scale items. The reported un-harmonized alphas are an average of these.

[13] It is unlikely that this short scale's lower reliability ($\alpha = 0.62$) yields the null association between being Latino and knowledge. Table H (SA) reports a set of analyses supporting this inference.

Table 2 Raw knowledge gaps between latinos and whites: un-harmonized and harmonized knowledge scales

	Pooled mean	Mean: Latinos (L)	Mean: Whites (W)	Raw difference (L − W)	Difference as Cohen's *d*	Difference as percentage (%)
Un-harmonized scale (5 traditional items)	2.92	2.12	3.65	−1.53[a]	0.88	31
Harmonized scale (1 traditional + 3 Latino items)	2.29	2.11	2.45	−0.33[a]	0.25	8

The un-harmonized scale runs from 0 to 5 items. The harmonized scale runs from 0 to 4 items. Means and differences in means are in item units. Cohen's *d* is a standardized mean difference produced by dividing a raw mean difference by its pooled standard deviation. Cohen's *d* values around 0.20, 0.50, and 0.80 are considered small, medium, and large, respectively

[a] Difference is reliable at 5 % level or better (two-tailed)

Table 3 Multivariate models of Latino–White differences in political knowledge

	Political knowledge (un-harmonized scale)	Political knowledge (harmonized scale)
Latino	−0.07* (0.02)	0.03 (0.02)
Income	0.12* (0.04)	0.19* (0.04)
Education	0.23* (0.03)	0.20* (0.03)
Age	0.16* (0.04)	0.24* (0.04)
Female	−0.06* (0.02)	−0.04* (0.02)
Political interest	0.29* (0.03)	0.31* (0.03)
Political efficacy	0.05^ (0.03)	0.08* (0.03)
Second generation	−0.02 (0.04)	−0.07^ (0.04)
Third generation	−0.00 (0.03)	−0.07^ (0.04)
U.S. citizen	0.07* (0.03)	0.04 (0.04)
Spanish interview	−0.10* (0.04)	0.07^ (0.04)
Constant	0.20* (0.05)	0.08 (0.05)
Adj. R^2	0.48	0.29
N	974	974

Entries are unstandardized OLS coefficients with standard errors in parentheses. All variables run from 0 to 1

* $p < .05$, two-tailed test; ^ $p < .10$, two-tailed test

knowledge scale are largely explained by "non-racial" variables, including age, education, and income. Since these factors are considered established correlates of political knowledge, their robust associations with the harmonized scale help to further validate this index.[14]

Beyond these conventional predictors of political knowledge, my analysis of the harmonized scale also reveals a marginally reliable association between interviewing in Spanish and political information, which cautions against viewing U.S. political knowledge as more prevalent among those who interview in English, the nation's dominant tongue. Indeed, the harmonized scale shows that Spanish interviewees might be slightly more attentive to national politics, and thus, more politically informed than English interviewees.[15] This comports with the results for second and third generation individuals, who are marginally less informed than their first generation counterparts. Finally, I find no reliable knowledge gap between U.S. citizens and non-citizens, which suggests that people who are not formal members of the polity are just as politically informed as those who have an official tie to the nation. These patterns align with prior

[14] The top five states with the highest Latino populations are California, Texas, Florida, New York, and Illinois (Brown and Lopez 2013). 49 % of my respondents are from those states (CA: 19 %; TX: 14 %; FL: 7 %; NY: 5 %; IL: 4 %). This implies that residence in those states, especially among Latinos, might affect knowledge levels on the harmonized scale, thus biasing the results in Table 3 (column 2), which omit this covariate. Table I (SA) reports a set of analyses showing that residence in those states is generally unrelated to knowledge levels.

[15] It is plausible that the association between knowledge and interview language is moderated by being Latino, yet collinearity prevents me from conducting this analysis since all Spanish interviewees in my sample are Latino.

work demonstrating an under-appreciated level of political engagement among less integrated segments of the polity, such as foreign-born, Spanish-dominant individuals (cf. Pantoja et al. 2001; Pantoja and Segura 2003; Barreto 2005).

Item Bias Affects Knowledge's Impact on Mass Opinion

Political knowledge is often deemed a workhorse variable in U.S. politics because it is associated with normatively appealing attitudes and behaviors (Delli Carpini and Keeter 1996). Below, I provide an example of how failure to address uneven item performance can also affect the substance of these types of conclusions. Specifically, I examine the relationship between knowledge levels and support for stricter immigration policy among Latinos and Whites, conditional on established correlates of policy preferences. Prior correlational analyses find large opinion differences between Whites and non-Whites on race-related issues, with non-Whites often displaying relatively more tolerant views (cf. Kinder and Winter 2001; Peffley and Hurwitz 2010). Yet despite this racial divide, Delli Carpini and Keeter (1996) find that higher knowledge levels often push the opinions of Whites *and* non-Whites in a more liberalized direction. Applied to immigration politics, these insights lead us to anticipate two patterns. First, since immigration reforms are likely to benefit many Latino families (Abrajano and Alvarez 2010), opposition to immigration should be weaker among members of this ethnic group relative to Whites. Second, net of this ethnic difference, greater knowledge should correspond with lower opposition to immigration among Latinos and Whites since, according to Delli Carpini and Keeter (1996), knowledge promotes more liberalized opinion on race-related issues.

To this end, I regressed an illegal immigration policy item on my un-harmonized (harmonized) scale of political knowledge, Latino status, and the interaction between these two variables, plus controls for partisanship, ideology, education, age, gender, employment status, and residence in a border state. The policy item reads, "Lawmakers in our nation's capital should make it harder for illegal immigrants to become U.S. citizens," with responses running on a four point scale from strongly disagree to strongly agree. Table J (SA) contains the raw results. Here, I visually display the relevant changes in the probability of strong agreement with this item based on each knowledge scale for both groups.[16]

Figure 2 (panel A) reveals mixed evidence for the set of anticipated patterns. On the one hand, in an absolute sense, Latinos do in fact express weaker opposition to immigration than Whites. This is seen in the near absence of overlap between Latinos' and Whites' probability of strongly agreeing with stricter illegal immigration policy across levels of political knowledge. On the other hand, greater levels of political knowledge reduce opposition to immigration among Whites, but

[16] I do this while setting my other covariates at their means (medians). This model closely follows those estimated by Delli Carpini and Keeter (1996) for other groups. See table J (SA) for details on variable codings.

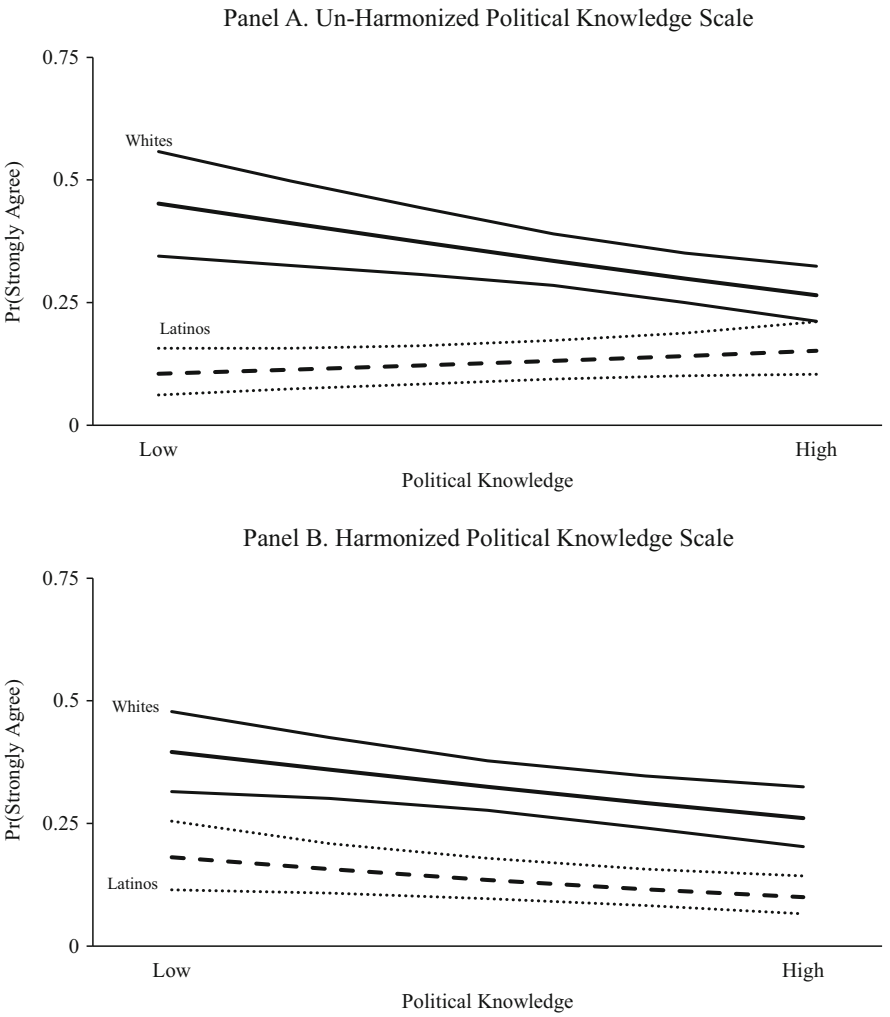

Fig. 2 Probability of "strongly agree" with stricter illegal immigrant policy by knowledge level among Latinos and Whites (with 95 % confidence intervals)

not Latinos. In fact, among Latinos, there is a mild but discernible increase in opposition to immigration across political knowledge levels.

More precisely, Whites with the lowest knowledge level show a 0.45 (95 % CI 0.34–0.56) probability of strongly agreeing that lawmakers should make it harder for illegal immigrants to obtain citizenship. Among Whites with the highest knowledge level, this probability drops significantly to 0.26 (95 % CI 0.21–0.32)—a large shift of 0.19 points. This predicted fall in support is absent among Latinos. While their absolute level of support is expected to be lower than Whites (which it is), prior work suggests political knowledge should further reduce this already weaker support (Delli Carpini and Keeter 1996). It does not. Moving from the

lowest to highest level of knowledge on the un-harmonized scale actually produces a 0.04 increase. Specifically, higher knowledge levels bump the probability of strong agreement among Latinos from 0.11 (95 % CI 0.06–0.16) to 0.15 (95 % CI 0.10–0.21). Although this shift is small and unreliable, the more important point is that it trends in the direction opposite from what prior work on varied groups and topics predicts (Delli Carpini and Keeter 1996). What is more, knowledge levels are associated with less opposition to immigration among Whites, but unassociated with it among Latinos—a group that is likely to directly experience the repercussions of a policy like this, given that a majority of undocumented immigrants are from Latin America (Hoefer et al. 2012). Thus, using an un-harmonized scale prevents me from uncovering a tried and true pattern where higher knowledge levels further whittle Latino opposition to this exclusionary immigration policy.

In contrast, panel B (Fig. 2) reveals that my expectations are generally confirmed when I use the harmonized knowledge scale. First, Latinos are significantly less likely than Whites to support this proposed illegal immigrant policy at all knowledge levels, as expected. Second, members of each group display the anticipated negative association between political knowledge and support for this policy, although at varied significance levels. For Whites with the lowest knowledge level, the probability of strong agreement with this proposal is 0.40 (95 % CI 0.31–0.48). Among Whites with the highest knowledge level, this probability drops to 0.26 (95 % CI 0.20–0.33), a large 0.19 decrease that is marginally reliable at the 10 % level.[17] Turning to Latinos with the lowest knowledge level, the probability of strong agreement is 0.18 (95 % CI 0.11–0.25). This probability falls to 0.10 (95 % CI 0.07–0.14) among Latinos with the highest knowledge level—a still substantial 0.08 drop that is just shy of significance at the 10 % level.[18] This pattern affirms prior work (Delli Carpini and Keeter 1996). Latinos already express a lower level of support for this immigration policy than Whites, yet higher knowledge levels pare down this support further. Thus, the harmonized items let me detect evidence that comports with the theorized relation between knowledge and immigration policy support.

These patterns also emerge if I operationalize knowledge with my one unbiased question that is not Latino-themed (i.e., *Conservative party*), though the trends unsurprisingly fall short of statistical significance given the low reliability entailed by measuring knowledge with one item. Specifically, Whites with the lowest knowledge level show a 0.39 (95 % CI 0.31–0.48) probability of strongly agreeing with the illegal immigration item. For Whites with the highest knowledge level, this probability drops to 0.29 (95 % CI 0.23–0.34). In turn, Latinos with the lowest knowledge level display a 0.15 (95 % CI 0.09–0.21) probability of strongly

[17] For Whites with the lowest knowledge level, the probability of strong agreement is 0.40 (90 % CI 0.33–0.47); for Whites with the highest knowledge level, the probability is 0.26 (90 % CI 0.21–0.32).

[18] For Latinos with the lowest knowledge level, the probability of strong agreement with this item is 0.18 (90 % CI 0.13–0.24). For Latinos with the highest knowledge level, the probability is 0.10 (90 % CI 0.07–0.14).

agreeing with this item. Among Latinos with the highest knowledge level, this probability drops to 0.11 (95 % CI 0.08–0.16).[19]

Summary and Implications

My analysis reveals that large group deficits in political knowledge are artificially inflated by uneven item performance in diverse populations. I illustrated this by re-examining knowledge levels among U.S. Latinos and Whites, where scholars have observed a large information gap. I showed that knowledge items that are relatively easy for Whites to correctly answer are often more difficult for Latinos due to item features that are unrelated to people's level of political information. This has the perverse effect of widening any actual difference in knowledge between these groups. Moreover, failure to address uneven item performance across groups mischaracterizes how much members of diverse groups actually know about politics, while distorting the statistical relations between knowledge and other political variables.

My findings train attention on the content validity of knowledge scales. While many scholars have settled on using factual items to assess political knowledge, less agreement exists on the depth of questions that should be used to measure it. This is easy to understand. It is difficult, if not impossible, to assess content validity when the entire item universe is unknown (Delli Carpini and Keeter 1996). Yet my results suggest it might be too much to ask for a complete tally of all possible knowledge items before content validity is more rigorously addressed. By expanding and validating the bank of available knowledge items, scholars can yield a finer-grained understanding about *which* items work best and among *whom*. I found that in my enriched question battery, domain- and policy-specific items worked more effectively across Latinos and Whites than most traditional knowledge questions (except *Conservative party*). This comports with prior work (cf. Dolan 2011; Barabas and Jerit 2009), but further research with more extensive question batteries can help determine (1) whether other traditional items work impartially across these and other groups; and (2) what the optimal balance is between traditional, domain-specific, and policy-specific items in political knowledge scales.

My findings also stress the need for more research on the origins and prevalence of knowledge gaps between Whites and non-Whites. Here it might be especially productive to create synergy between prior work and my results. For example, Davis and Silver (2003) find that stereotype threat leads African Americans to report less political knowledge to White rather than Black interviewers. Yet this study does not consider how this gap is shaped by item bias among Whites and Blacks. My analysis illustrates the role of unequal items among Whites and Latinos, but does not address how the interview setting affects the reporting of political facts. Hence, future

[19] To bolster my case that un-harmonized scales distort conclusions about knowledge's impact on political judgments, I conduct a second analysis where—based on prior work (e.g., Delli Carpini and Keeter 1996; Zaller 1992)—I examine the extent to which knowledge moderates the link between one's partisanship and rating of a co-partisan (i.e., President Obama). I report the raw results in table K, graph the relevant quantities in figure L, and discuss these findings in section M (see SA).

studies might consider building harmonized knowledge scales across Blacks, Latinos, Asians, and Whites, while varying the race of interviewer and mode of interview. This combination would enable scholars to gain a better handle on the interplay between measurement and survey context in the analysis of political knowledge in the American mass public.

While my analysis focused on the knowledge gap between Latinos and Whites, the approach I used and the insights it revealed travel beyond this case. When designing factual knowledge items, scholars tend to examine how well a set of items reflect this variable, with affirming evidence arriving in the form of high inter-item correlations, high (α) coefficients, or some combination of these criteria (cf. Luskin 1987; Delli Carpini and Keeter 1996; Mondak and Anderson 2004). Yet this approach privileges the role of item characteristics in generating estimates of political knowledge. In contrast, my framework distinguishes between item characteristics (i.e., loadings) and respondent characteristics (i.e., thresholds). By doing so, it underlines the role of respondent heterogeneity in affecting measurement quality. This is a lesson with growing import as America's ethnic and racial landscape continues to change.

Acknowledgments I express my deepest thanks to Larry Bartels, Joshua Clinton, John Geer, Marc Hetherington, Cindy Kam, Steven Rogers, and Liz Zechmeister for providing me with invaluable advice on the design and analysis of this study. Any remaining errors are mine alone.

Conflict of interest The author declares he has no conflict of interest.

Ethical Standards This study complies with relevant U.S. laws.

References

Abrajano, M. A. (2010). *Campaigning to the new American electorate: advertising to Latino voters*. Palo Alto: Stanford University Press.

Abrajano, M. A., & Alvarez, R. M. (2010). *New faces, new voices: the Hispanic electorate in America*. Princeton: Princeton University Press.

Althaus, S. L. (1998). Information effects in collective preferences. *American Political Science Review, 92*(3), 545–558.

Barabas, J., & Jerit, J. (2009). Estimating the causal effects of media coverage on policy-specific knowledge. *American Journal of Political Science, 53*(1), 73–89.

Barreto, M. A. (2005). Latino immigrants at the polls: foreign-born voter turnout in the 2002 election. *Political Research Quarterly, 58*(1), 79–86.

Bartels, L. M. (1996). Uninformed votes: information effects in presidential elections. *American Journal of Political Science, 40*(1), 194–230.

Branton, R. (2007). Latino attitudes toward various areas of public policy: the importance of acculturation. *Political Research Quarterly, 60*(2), 293–303.

Brown, T. A. (2006). *Confirmatory factor analysis for applied research*. New York: Guilford Press.

Brown, A., & Lopez, M. H. (2013). *Mapping the Latino population, by state, county and city*. Washington, D.C.: Pew Research Center.

Byrne, B., Shavelson, R. J., & Muthén, B. (1989). Testing for the equivalence of factor covariance and mean structures: the issue of partial measurement invariance. *Psychological Bulletin, 105*(3), 456–466.

Camilli, G., & Shepard, L. A. (1994). *Methods for identifying biased test items*. Thousand Oaks: Sage Publications.

Chideya, F. (1995). *Don't believe the hype: fighting cultural misinformation about African-Americans*. New York: Penguin Books.
Cohen, J. (1988). *Statistical power analysis for the behavioral sciences*. Hillsdale: Lawrence Erlbaum.
Davis, D. W., & Silver, B. D. (2003). Stereotype threat and race of interviewer effects in a survey on political knowledge. *American Journal of Political Science, 47*(1), 33–45.
Delli Carpini, M. X., & Keeter, S. (1996). *What Americans need to know about politics and why it matters*. New Haven: Yale University Press.
DeVellis, R. F. (2003). *Scale development: theory and applications*. Thousand Oaks: Sage.
Dolan, K. (2011). Do women and men know different things? Measuring gender differences in political knowledge. *The Journal of Politics, 73*(1), 97–107.
Dow, J. K. (2009). Gender differences in political knowledge: distinguishing characteristics-based and returns-based differences. *Political Behavior, 31*(1), 117–136.
Embretson, S. E., & Reise, S. P. (2000). *Item response theory for psychologists*. Mahwah: Lawrence Erlbaum Associates Inc.
Gay, C. (2013). Knowledge matters: policy cross-pressures and Black partisanship. *Political Behavior,*. doi:10.1007/s11109-013-9227-3.
Gibson, J. L., & Caldeira, G. A. (2009). Knowing the Supreme Court? A reconsideration of public ignorance of the High Court. *The Journal of Politics, 71*(2), 429–441.
Gilens, M. (2001). Political ignorance and collective policy preferences. *American Political Science Review, 95*(2), 379–396.
Grunwald, M. (2013). Immigrant son: Marco Rubio wants to sell the GOP on a path to citizenship for undocumented Americans. So why is his mom calling? *Time*. Retrieved from https://time.com.
Hambleton, R. K., Swaminathan, H., & Rogers, H. J. (1991). *Fundamentals of item response theory*. Newbury Park: Sage.
Hoefer, M., Rytina, N., & Baker, B. C. (2012). *Estimates of the unauthorized immigrant population residing in the United States: January 2011*. U.S. Office of Immigration Statistics: Policy Directorate.
Hutchings, V. L. (2001). Political context, issue salience, and selective attentiveness: constituent knowledge of the Clarence Thomas confirmation Vote. *Journal of Politics, 63*(3), 846–868.
Iyengar, S. (1986). Whither political information. ANES Pilot Study Report, No. nes002253.
Iyengar, S. (1990). Shortcuts to political knowledge: the role of selective attention and accessibility. In J. A. Ferejohn & J. H. Kuklinski (Eds.), *Information and democratic processes*. Urbana: University of Illinois Press.
Kinder, D. R., & Winter, N. (2001). Exploring the racial divide: Blacks, Whites, and opinion on national policy. *American Journal of Political Science, 45*(2), 439–453.
Kuklinski, J. H., Quirk, P. J., Jerit, J., Schwieder, D., & Rich, R. F. (2000). Misinformation and the currency of democratic citizenship. *The Journal of Politics, 62*(3), 790–816.
Lupia, A. (2006). How elitism undermines the study of voter competence. *Critical Review, 18*, 217–233.
Luskin, R. C. (1987). Measuring political sophistication. *American Journal of Political Science, 31*(4), 856–899.
Luskin, R. C., & Bullock, J. G. (2011). "Don't know" means "don't know": DK responses and the public's level of political knowledge. *Journal of Politics, 73*(2), 547–557.
McGraw, K., & Pinney, N. (1990). The effects of general and domain-specific expertise on political memory and judgment. *Social Cognition, 8*(1), 9–30.
Metzler, R. (2013). Favorite son? Marco Rubio poised to challenge Dems in 2016. *U.S. News and World Report*.
Mondak, J. J. (1999). Reconsidering the measurement of political knowledge. *Political Analysis, 8*, 57–82.
Mondak, J. J. (2001). Developing valid knowledge scales. *American Journal of Political Science, 45*(1), 224–238.
Mondak, J. J., & Anderson, M. R. (2004). The knowledge gap: a reexamination of gender-based differences in political knowledge. *The Journal of Politics, 66*, 492–512.
Mondak, J. J., & Davis, B. C. (2001). Asked and answered: knowledge levels when we will not take "don't know" for an answer. *Political Behavior, 23*(3), 199–224.
Pantoja, A. D., Ramirez, R., & Segura, G. M. (2001). Citizens by choice, voters by necessity: patterns in political mobilization by naturalized Latinos. *Political Research Quarterly, 54*(4), 729–750.
Pantoja, A. D., & Segura, G. M. (2003). Fear and loathing in California: contextual threat and political sophistication among Latino voters. *Political Behavior, 25*(3), 265–286.

Peffley, M., & Hurwitz, J. (2010). *Justice in America: the separate realities of Blacks and Whites*. New York: Cambridge University Press.
Price, V., & Zaller, J. (1993). Who gets the news? Alternative measures of news reception and their implications for research. *Public Opinion Quarterly, 57*, 133–164.
Prior, M. (2013). Visual political knowledge: a different road to competence? *The Journal of Politics., 72*(3), 747–766.
Prior, M., & Lupia, A. (2008). Money, time, and political knowledge: distinguishing quick recall and political learning skills. *American Journal of Political Science, 52*(1), 169–183.
Sanbonmatsu, K. (2003). Gender-related political knowledge and the descriptive representation of women. *Political Behavior, 25*(4), 367–388.
Sijtsma, K. (2009). On the use, the misuse, and the very limited usefulness of Cronbach's alpha. *Psychometrika, 74*(1), 107–120.
Spector, P. E. (1992). *Summated rating scale construction*. Thousand Oaks: Sage.
Stegmueller, D. (2011). Apples and oranges? The problem of equivalence in comparative research. *Political Analysis, 19*, 471–487.
Sturgis, P., Allum, N., & Smith, P. (2008). An experiment on the measurement of political knowledge in surveys. *Public Opinion Quarterly, 72*(1), 90–102.
Tichenor, D. J. (2002). *Dividing lines: the politics of immigration control in America*. Princeton: Princeton University Press.
Verba, S., Schlozman, K. L., & Brady, H. E. (1995). *Voice and equality: civic voluntarism in American politics*. Cambridge: Harvard University Press.
Zaller, J. (1986). Analysis of information items in the 1985 ANES pilot study. ANES Pilot Report, No. nes002261.
Zaller, J. (1990). Political awareness, elite opinion leadership, and the mass survey response. *Social Cognition, 8*(1), 125–153.
Zaller, J. R. (1992). *The nature and origins of mass opinion*. New York: Cambridge University Press.
Zeller, R. A., & Carmines, E. G. (1980). *Measurement in the social sciences: the link between theory and data*. New York: Cambridge University Press.