UC Berkeley

UC Berkeley Previously Published Works

Title

Face Value? Experimental Evidence that Candidate Appearance Influences Electoral Choice

Permalink

<https://escholarship.org/uc/item/62b082sm>

Journal

Political Behavior, 39(1)

ISSN

0190-9320

Authors

Ahler, DJ
Citrin, J
Dougal, MC
[et al.](#)

Publication Date

2017-03-01

DOI

10.1007/s11109-016-9348-6

Peer reviewed

Face Value? Experimental Evidence that Candidate Appearance Influences Electoral Choice

Douglas J. Ahler¹ · Jack Citrin¹ · Michael C. Dougal¹ · Gabriel S. Lenz¹

Published online: 27 May 2016
© Springer Science+Business Media New York 2016

Abstract According to numerous studies, candidates' looks predict voters' choices—a finding that raises concerns about voter competence and about the quality of elected officials. This potentially worrisome finding, however, is observational and therefore vulnerable to alternative explanations. To better test the appearance effect, we conducted two experiments. Just before primary and general elections for various offices, we randomly assigned voters to receive ballots with and without candidate photos. Simply showing voters these pictures increased the vote for appearance-advantaged candidates. Experimental evidence therefore supports the view that candidates' looks could influence some voters. In general elections, we find that high-knowledge voters appear immune to this influence, while low-knowledge voters use appearance as a low-information heuristic. In primaries, however, candidate appearance influences even high-knowledge and strongly partisan voters.

Keywords Elections · Candidate appearance · Congressional elections · Primary elections · Heuristics

Introduction

On what basis do voters decide? The answer bears directly on debates over citizens' competence and the quality of the officials they elect. Confronted by evidence of widespread ignorance, ideological innocence, and the paucity of issue voting

Electronic supplementary material The online version of this article (doi:[10.1007/s11109-016-9348-6](https://doi.org/10.1007/s11109-016-9348-6)) contains supplementary material, which is available to authorized users.

✉ Gabriel S. Lenz
glenz@berkeley.edu

¹ Charles and Louise Travers Department of Political Science, University of California, Berkeley, 210 Barrows Hall #1950, Berkeley, CA 94720-1950, USA

(Converse 1964; Delli Carpini and Keeter 1996; Lenz 2012), scholars have looked to heuristics—simple rules of thumb—to haul voters onto the shores of rationality (Fiorina 1981; Key 1966; Lupia 1994; Popkin 1991). Voters undoubtedly do rely on informative heuristics, such as the state of the economy, party ties, and feelings about the incumbents. But how often do simple rules of thumb lead them astray from a more informed and appropriate choice (Kuklinski and Quirk 2000)?

One potentially worrisome heuristic is a candidate's appearance. Endorsing Mitt Romney for the 2012 Republican presidential nomination, Bob Dole declared, "So it looked to me like it would be either Romney or Newt [Gingrich] for the nomination, but... Romney *looks* like a president" (Laskas 2012, 88). Whether candidates look like presidents may not be entirely uninformative, but seems unlikely to provide much information about the candidates—a point we return to in the conclusion. Nevertheless, research implies that some voters evaluate candidates as Dole did in 2012. They vote for politicians whose appearance in photographs is judged more competent or attractive by naïve raters (those who neither know nor recognize the candidates) at higher rates in actual U.S. Senate and House elections, as well as in elections abroad (Atkinson et al. 2009; Ballew and Todorov 2007; Banducci et al. 2008; Berggren et al. 2010; Hall et al. 2009; King and Leigh 2009; Lawson et al. 2010; Mattes et al. 2010; Olivola and Todorov 2010; Rosar et al. 2008; Spezio et al. 2012; Todorov et al. 2005).

While these observational studies find that candidate appearance *correlates* with actual election results, these studies do not show that candidate appearance actually *causes* voters to change their minds in real-world elections. It is troubling for the quality of electoral choices if voters conflate mere physical appearance—stylish hair, straight teeth or a strong jaw—and actual competence, but the results of observational research cannot rule out several alternative explanations for the correlation between appearance and vote choice, leaving the causal mechanism ambiguous.

Do voters rely on the seemingly superficial heuristic of appearance when voting? Or is there an alternative explanation for the correlation between appearance and vote choice that reveals more competent voters? Foremost among the possible alternative explanations is the influence of candidate effort. Candidates who campaign harder or have more resources may look better because they also put more effort into their press materials (including photographs) or because their superior resources can pay for professional photographers, image consultants, \$600 haircuts, and the like. But if this is the case then the observationally estimated effects of candidate appearance could be entirely spurious, an artifact of the correlation between improved candidate appearance and other facets of campaign effort—voter mobilization, showing up at events, outreach and communication, etc.—that may be the real drivers of voters' choices.

Campaign effort—whether by the candidate or her party—could thus make candidates' looks appear to influence voters when they do not do so directly. Indeed, candidates who outspend their opponents do look better to naïve raters. Examining 2006 Senate races, Lenz and Lawson (2011, 584–5) find a 0.59 correlation between a candidate's spending advantage and appearance advantage and a 0.56 correlation between incumbency and appearance advantage. However, after controlling for

variables that might capture campaign effort, such as race competitiveness and candidate spending, Atkinson et al. (2009) find that candidate appearance still has a small effect on vote share.¹ Given these observational findings, how should we assess the possibility that the appearance-vote findings are spurious rather than causal?

Sorting out causation here is hard. Statistically controlling for variables such as competitiveness, spending, or incumbency is appropriate only if these variables are causes of candidate appearance, not consequences. If appearance is in fact causally prior to spending—e.g., better looking candidates can raise more money, win endorsements based on looking the part (like Governor Romney), or are likely to have won previous elections (and so become incumbents)—then researchers should not control for these variables because they could be consequences of candidate appearance (i.e., post-treatment).² Put differently, if candidates can raise more money or attract more volunteers because they are better looking, then controlling for such variables will bias estimates of the candidate appearance effect downward by incorrectly attributing part of the true effect of appearance to these variables.³ Given the possibility for complex causal relations among these variables, drawing firm inferences with observational data may be impossible.⁴

For these reasons, we test the influence of a candidate's appearance on voters with two experiments rather than with observational studies. We interviewed individuals just before elections in which they said they would likely participate and asked for their voting intention. Crucially, however, we randomly assigned participants to one of two conditions: (1) a control group received a ballot designed

¹ Several other observational results are inconsistent with the alternative explanation emphasizing the causal influence of campaign effort entirely explaining the observed effect of candidate appearance. Specifically, the effect of the candidate's appearance holds when professional photographers took the pictures in a standard format (Antonakis and Dalgas 2009; Klein and Rosar 2005), and when one statistically controls for differences in image quality and other aspects of the pictures, such as visible light (Lawson et al. 2010; Rosar et al. 2008). Additionally, appearance-advantaged candidates win in competitive races, where the candidates should be more comparable in quality and in resources (Antonakis and Dalgas 2009; Benjamin and Shapiro 2009). They also perform disproportionately well in systems where legislators compete against members of the same party (Berggren et al. 2010) and in non-partisan contests (Banducci et al. 2008; Martin 1978).

² In general, researchers should not control for variables that intervene between the treatment and the outcome, in this case, between candidate appearance and vote share. For a general discussion, see King (1991, 1049–50).

³ When estimating the effect of challenger appearance, Atkinson et al. (2009) carefully try to avoid post-treatment bias by measuring district competitiveness at least 1 year before the general election, when the challenger's identity is less clear (using the *Cook Political Report*). Nevertheless, these experts may already know the likely challengers and so may be influenced by their looks (making these ratings post-treatment).

⁴ Indeed, Atkinson et al. (2009, 236) are careful not to interpret their regression coefficient for incumbent appearance as a causal estimate. They suggest instead that appearance-advantaged incumbents (as challengers in a prior election) disproportionately select into competitive districts, which would bias their estimate of incumbent appearance downwards. This downward bias and, more generally, the causal complexity of observational studies on appearance provide reasons to turn to experimental studies such as ours.

to resemble the one actually used and (2) a treatment group received a ballot that also showed candidates' photos next to their names. To evaluate whether appearance directly influences votes, we simply compare the degree to which candidate appearance predicts vote intentions in the two conditions.

This research design sheds light on the appearance-vote relationship in a way that previous studies could not. It does so because random assignment rules out the alternative explanations. Since candidate effort—raising more money, shaking more hands, kissing more babies—cannot differentially influence voters in the photo condition, any effect we detect must be a result of viewing the photos of the candidates.⁵ The experimental design can thus provide internally valid evidence that candidate appearance influences vote choice, even when real world voters have other information about the candidates. However, random assignment only provides clean causal identification within the experiment on the mock ballots, which differ from actual ballot results and therefore do not necessarily provide generalizable (externally valid) estimates. Thus, this design may not tell us *how much* candidate appearance really matters to election outcomes, but that it *could* matter. Given concerns that the entire appearance-vote correlation could be from omitted variable bias, our studies make an important contribution by providing experimental evidence that at least a portion of the appearance-vote correlation is causal, even if they do not tell us precisely how much of that correlation is due to this causal effect.

We ran this experiment on 14 House races in the 2012 California congressional primary and 44 statewide races across 18 states in the 2012 general election. Using naïve raters' assessments of candidate appearance, we find that including candidates' photographs on the ballot does indeed lead participants to vote more often for appearance-advantaged candidates. Using actual candidates in the midst of an election as stimuli, we find that a substantively significant percentage of our participants (9 % in the general election races) voted differently than they otherwise would because they saw candidate photos. Although our emphasis is on the existence of an appearance effect, not its magnitude, we note that at face value this shift would be large enough to change the outcome in roughly 29 % of primary races and about 14 % of general election races.

The rest of the paper proceeds as follows: First, we describe our research design more fully and present candidate-level results for House primaries in California. Second, we replicate these results in higher-salience, statewide general election races. Third, we consider external validity and assess the robustness of our findings. Fourth, we conduct individual-level analyses for both studies and show that candidate appearance most heavily influences low-information voters in general elections and matters more in the earlier stages of campaigns. This important nuance in our results helps us assess the ramifications of these studies for voter competence and democratic accountability.

⁵ Of course, candidates' efforts to "improve" their appearance, as revealed through their photos, may contribute to any such causal effects.

Study 1: Appearance Advantage in the 2012 California House Primaries

Design and Procedures

Starting 10 days before the 2012 California primary, an Internet poll conducted by Survey Sampling International (SSI) interviewed 1268 registered voters from 14 of California's 53 House districts. The sample adequately represents registered voters on age, party registration, and political ideology. Fifty-three candidates ran in these 14 districts—11 females and 42 males.⁶ In terms of partisanship, there were 23 Republicans, 23 Democrats, and 10 no-party preference or other party candidates. Importantly, the survey's election results closely mirror the actual election results [see the online appendix (OA), section 1.1, which also presents the demographic characteristics of participants. The OA is available with the online version of this article at the journal website].

As noted, we randomly assigned participants to one of two conditions. Those assigned to the control condition received a ballot identical to the one they would see in the actual June 5 election (a top-two primary ballot with all candidates listed, regardless of party). In the treatment condition, we gave participants the same choice of candidates but also displayed black-and-white photographs of the candidates' faces next to their names. Figure 1 provides an example.

We measured the appeal of candidates' appearances in a separate survey by showing U.S. workers on Amazon's Mechanical Turk the photos and asking, "How good of a congressperson do you think this person would be?" (see OA section 1.2 for survey details). We use this general measure to sidestep the debate about which traits voters primarily respond to—e.g., competence (Todorov et al. 2005) or attractiveness (Banducci et al. 2008).⁷ Our measure, however, correlates more strongly with perceptions of facial competence (0.61) and attractiveness (0.60), than with perceptions of dominance (0.19, see OA section 2.3 for more detail). This general appearance measure predicts election outcomes as well as specific trait-based measures (Lawson et al. 2010). To obtain naïve ratings, we excluded ratings from California Mechanical Turk workers and from workers who reported recognizing the candidate (just 0.1 % of 9174 ratings). We gave participants in the experiments five options for their evaluations of "how good of a congressperson do you think this person would be," ranging from "much better than average" to "much worse than average." Mean candidate ratings ranged from a low of 2.4 on the five-point scale (Jim Reed, a Democratic candidate from district 1) to a high of 3.5 (Mary Bono Mack, an incumbent Republican from district 36), with a mean of 2.9 and a standard deviation of 0.27. For the analysis below, we recode this measure

⁶ We also ran the experiment in six California State Senate races. We do not pool these races with the House primaries in the analysis because photograph quality was noticeably lower. Instead, we present these results in OA section 1.2. Including them in the main analysis leaves our key findings unchanged.

⁷ No matter what measure of appearance we choose, that measure will also pick up other characteristics that correlate with it. One way to break these correlations would be to artificially alter candidate photos in order to experimentally vary these traits, but using altered photos would significantly reduce the external validity of our experiments. We thus use actual candidate photos and make no claim about the particular aspect of a candidate's appearance that influences voters.

Which primary candidate for The House of Representatives do you plan to vote for?

☐ Bruce Margolin Party Preference: Democratic Occupation: Marijuana Legalization Attorney	☐ David William Steinman Party Preference: Green Occupation: Consumer Rights Advocate
☐ Tim Pape Party Preference: Democratic Occupation: TV Editor	☐ Zeln E. Obagi Party Preference: Democratic Occupation: California Business Attorney
☐ Christopher David Party Preference: Republican Occupation: Entrepreneur	☐ Henry A. Waxman Party Preference: Democratic Occupation: United States Representative
☐ Steve Collett Party Preference: Libertarian Occupation: Certified Public Accountant	☐ Don't know yet
☐ Bill Bloomfield Party Preference: None Occupation: Small Business Owner	

Which primary candidate for The House of Representatives do you plan to vote for?

☐ Zeln E. Obagi Party Preference: Democratic Occupation: California Business Attorney	☐ David William Steinman Party Preference: Green Occupation: Consumer Rights Advocate
☐ Christopher David Party Preference: Republican Occupation: Entrepreneur	☐ Bill Bloomfield Party Preference: None Occupation: Small Business Owner
☐ Steve Collett Party Preference: Libertarian Occupation: Certified Public Accountant	☐ Bruce Margolin Party Preference: Democratic Occupation: Marijuana Legalization Attorney
☐ Tim Pape Party Preference: Democratic Occupation: TV Editor	☐ Don't know yet
☐ Henry A. Waxman Party Preference: Democratic Occupation: United States Representative	

Fig. 1 Example of control (*Top*) and treatment (*Bottom*) ballots in California primary experiment (Study 1). *Note:* The ballots showed the same information as the California 2012 primary ballot, except of course for the photos. This example is from district 33

so it captures *Appearance Advantage* within a district by subtracting from a candidate's (e.g., Mary Bono's) raw appearance rating the mean rating of all the candidates in her district. We rescale this variable to a 0–1 scale, but the results are robust to other coding procedures (see OA section 1.5).

Our dependent variable is a candidate's vote share in the treatment condition minus her vote share in the control condition (*Photo Condition minus No-Photo Condition Vote Share*). This measures the degree to which a candidate performs better or worse when voters see her face and her opponents' faces on the ballot. If appearance matters, candidates should receive an increasing vote share in the treatment condition as their appearance advantage grows. About 70 % of participants reported an intended vote choice—congressional primaries have low salience—so we lose 30 % of participants when we calculate the dependent variable. (The voting rate was similar across conditions.) We also exclude an additional 2 % because they said they would not vote in the actual election, leaving us with 851 participants. (The results are similar in the full sample.)

External Validity

The experimental design provides internally valid evidence that candidate appearance can influence vote choice. In U.S. elections, of course, ballots do not show photos. So, to what extent do the findings help us understand the appearance effect in real-world elections? This question is difficult to answer. To make the findings as generalizable as possible, we took several steps: (1) we only interviewed registered voters, (2) we excluded respondents who said they would not vote, (3) we interviewed respondents close to Election Day, (4) and we showed respondents ballots with actual candidates from their own districts. Despite these efforts, however, there are reasons why we may be overestimating the real-world appearance effect, but there are also reasons why we may be underestimating it. We therefore devote a section of the paper, after presenting the results of both studies, to discussing external validity.

Candidate-Level Results

Figure 2 presents evidence that showing photos of the candidates to voters just before the election influences how they vote. The vertical axis presents the difference in candidate vote share between the photo (treatment) and no-photo (control) conditions, while the horizontal axis places candidates according to their within-district appearance advantage. The positive trend in Fig. 2 indicates that appearance-advantaged candidates—that is, those rated highly by the naïve Mechanical Turk raters—do receive more votes when voters see photos on the ballot.

For example, incumbent Democrat Henry Waxman (CA-33) suffers from a considerable appearance disadvantage, scoring poorly on the relative appearance measure (0.13). When participants saw the candidates' photos (which we show in the example ballot in Fig. 1), support for Waxman dropped by about 10 percentage points. In contrast, Waxman's most appearance advantaged opponent, an unknown

Fig. 2 Appearance-advantaged candidates in house primaries benefit in photo condition. *Note:* Each point represents a candidate in the 2012 U.S. House primaries in California. Observations weighted by the total number of respondents reporting a vote choice in the candidate's district

Democratic candidate named Tim Pape, scored near the top of the appearance measure (0.81) and received almost a 15-point boost in vote share when voters saw the photographs in the multicandidate primary. Waxman also lost out to “no party preference” (NPP) candidate Bill Bloomfield, who scored well on the appearance measure (0.62) and received a 10-point boost in vote share in the treatment group. Bloomfield went on to lose only narrowly to the powerful incumbent Waxman in the general election.

Critically, this effect cannot result from Pape or Bloomfield exerting more campaign effort or strategically choosing to enter this race. Since we randomly assigned voters to the photo and no-photo conditions, the candidates' campaigns—mostly absent in Pape's case anyway—could not have disproportionately influenced the treatment group.

This pattern holds on average across candidates, as the scatterplot in Fig. 2 shows. The slope of the best-fit line through the data points is 0.21, indicating that if the most appearance-advantaged and appearance-disadvantaged candidates (across all districts) received equal vote shares in the control condition, we would expect to see a 21-point difference in vote shares in the appearance-advantaged candidate's favor in the photo condition. More typically, a one standard deviation increase in appearance advantage, coupled with the treatment ballot, yields a 5.4-point boost, an effect that could alter outcomes in several of these primaries. The bivariate regression analysis in Table 1 presents the estimated slope of the best-fit line in Fig. 2 and shows that this 0.21 slope is unlikely to occur by chance alone [the 95 % confidence interval (95 % CI) for the estimate is 0.10 to 0.31]. This effect is robust

across important categories of candidates. As Columns 2 and 3 in Table 1 show, appearance-advantaged incumbents and challengers both benefit from having their photo shown. The next two columns show that appearance affects the fortunes of both viable and nonviable candidates, indicating that the photos did not simply serve to remind participants of real political information they had previously learned about prominent, photogenic candidates.⁸ Finally, Columns 6 and 7 show that appearance matters for both Democratic and Republican politicians: appearance-advantaged candidates from both parties fared significantly better in the treatment condition.

In addition, candidates' race and gender fail to explain the appearance effect. Column 8 of Table 1 presents the estimated effect of appearance-advantage, controlling for candidate race, gender, and incumbency status.⁹ Interestingly, the photos appear to change the impact of these characteristics on voting intentions: incumbent and male candidates perform significantly worse, and white candidates significantly better on the photo ballot. Even so, neither the substantive nor the statistical significance of the appearance advantage effect changes when we control for these covariates.

A potential concern is that voters' reactions to especially unprofessional looking candidates may drive this result. For some candidates, the only pictures available showed them wearing T-shirts or with disheveled hair. Rather than reacting to candidates' physical features, voters could be inferring candidates' competence based on their unprofessional presentation. However, when we exclude races with the most poorly rated candidates, those below the 25th percentile on overall photograph ratings, we find a larger treatment effect, not a smaller one (Table 1, Column 10, see OA section 1.6 for scatterplots and analyses with other cutoffs, which yield similar results). Thus, voters do not appear to be simply reacting negatively to candidates whose photos may signal incompetence.¹⁰

Study 2: Appearance Advantage in the 2012 General Elections

The results from the primary study imply that candidates' looks can directly affect voters. Our experimental design enables us to reject the argument that candidate effort or other confounding variables fully account for previous observational

⁸ The authors and a team of research assistants used endorsements, campaign finance data, previous office, number of competing co-partisans, and vote share in previous elections to classify candidates as viable or nonviable before the election. These assessments were largely holistic but are validated by the actual election results: A nonviable-classified candidate finished ahead of a viable-classified candidate in just 2 of the 13 races under study here, and neither of those candidates came close to advancing to the general election.

⁹ Ideally, we would address race and gender not with controls but by restricting the analysis to candidates matched on race and gender, but only three of the 14 races in Study 1 were so matched. We are, however, able to conduct this analysis in the second study.

¹⁰ Finally, to address separate concerns about candidate vote share not being independently distributed within district, OA section 4.6 shows that the photo-ballot respondents were 10 percentage points more likely to vote for the most appearance-advantaged candidate in their district compared to control ballot respondents ($p < 0.001$).

Table 1 Appearance-advantaged candidates in house primaries benefit in photo condition

Dependent variable: photo condition minus no-photo condition vote share	(1) All	(2) Challengers	(3) Inc.	(4) Non-viable	(5) Viable	(6) Dem.	(7) Rep.	(8) All with controls	(9) Races with min. face rating at/below 25th percentile	(10) Races with min. face rating above 25th percentile
Appearance advantage	0.21*** (0.05)	0.18*** (0.05)	0.30** (0.14)	0.14** (0.06)	0.19** (0.09)	0.25*** (0.08)	0.19** (0.08)	0.15*** (0.05)	0.11** (0.05)	0.28*** (0.08)
Incumbent								-0.09*** (0.03)		
White								0.05 (0.03)		
Male								-0.10*** (0.04)		
Constant	-0.09 (0.03)	-0.06 (0.03)	-0.17 (0.07)	-0.03 (0.04)	-0.10 (0.04)	-0.15 (0.04)	-0.04 (0.04)	0.01 (0.05)	-0.04 (0.02)	-0.12 (0.04)
Candidates (N)	53	39	14	24	29	23	23	53	16	37
R ²	0.23	0.24	0.28	0.18	0.14	0.30	0.23	0.41	0.25	0.26

This table shows candidate-level regressions (each column showing a separate model) with standard errors in parentheses. The Column 1 regression simply shows the 0.21 slope of the best-fit line in Fig. 2. The number of participants in these regressions is 851, all of whom are registered voters. We do not cluster the standard errors at the district level because, with only 14 clusters, clustering is unreliable (Angrist and Pischke 2009, ch. 8), but we do cluster by candidate in the individual-level analysis shown in OA section 4.5

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

findings of the appearance effect. However, because voters often pay little attention to primary campaigns and know little about the candidates, they may be more likely to rely on appearance as a low-information heuristic in this electoral setting. Would these results also hold in general election races? We examine this in Study 2.

Design and Procedures

In the 3 weeks prior to the general election on November 6, 2012, we recruited 2235 participants across 18 states through Amazon's Mechanical Turk. We asked these participants about their likely vote choices in anywhere between one and nine statewide races, which ranged from higher-salience races (i.e., gubernatorial or senatorial) to down ballot races (e.g., attorney general).¹¹ We selected the states before running the study because they possess sufficiently large populations of Mechanical Turk workers. The 44 contests included: 15 U.S. Senate races; three for governor; four each for attorney general and state treasurer; three each for lieutenant governor and secretary of state; two each for commissioner of insurance; state auditor, and superintendent of public insurance; and one for agricultural commissioner; labor commissioner; public land commissioner; railroad commissioner; presiding judge court of criminal appeals; and university board of regents (see OA section 2.1.1 for a list of races). The candidates included 26 females and 62 males. Of the 2235 participants in the study, we excluded four percent because they failed an attention test and an additional 10 percent because they reported not intending to vote, leaving us with 1933 participants (the results are similar in the full sample; see OA section 2.2).

As in Study 1, we randomly assigned participants to one of two conditions: a standard ballot or a ballot with candidate photographs. Participants assigned to the control condition received a mock ballot that included the statewide races they would see on the actual ballot on Election Day. In the treatment condition, we gave participants the same ballot but added black-and-white candidate photographs. In both conditions, the ballot showed the candidate's party affiliation.

We then measured candidate appearance in a separate survey on Amazon's Mechanical Turk, asking participants "How good of an elected official (e.g., Sen. or Gov.) do you think this person would be?" (see OA section 2.3 for survey details). To obtain naïve assessments, we excluded ratings when participants reported recognizing the candidate. Participants rated the candidates on the same 1–5 scale as used in the previous study, and their ratings have a mean of 3.1 and a standard deviation of 0.32. We measure *Appearance Advantage* in the direction of the Republican candidate (Republican Appearance minus Democratic Appearance), rescaled to a 0–1 range (see Fig. 3 for the ratings). Since all these contests featured just two candidates—one Democrat and one Republican—we can conduct the analysis at the race-level as opposed to the candidate-level. Thus, our dependent variable is the Republican vote share in the photo condition minus the Republican

¹¹ Survey dates: October 17–November 2. Election Day was November 6. We also asked participants about a handful of multicandidate races and single-candidate judicial retention elections. Since analyzing races with only one or more than two candidates introduces complications, we relegate analysis of these races to the OA (see OA section 5.1). The results are consistent with the overall findings in the paper.

vote share in the no-photo condition (*Photo Condition minus No-Photo Condition Vote Share*).

Race-Level Results

Even in general election races, voting intentions differed (in the aggregate) when the ballot included photos of the candidates' faces. Figure 3 plots the relationship between Republican appearance advantage (horizontal axis) and the difference between the Republican candidate's treatment vote share and control voter share (vertical axis). The slope of the regression line is 0.20, implying that the Republican candidate with the largest appearance advantage—Bob Corker in the Tennessee Senate race—is predicted to benefit by 20 percentage points over the most appearance-disadvantaged Republican—Steve Royal in the North Carolina State Treasurer race—when voters see candidates' faces on the ballot. Put in terms of a difference we are more likely to observe, a one standard deviation improvement in Republican appearance advantage, coupled with showing photographs to voters, would yield an expected 4-point boost in vote share for the Republican candidate. The estimated slope is significant [95 % CI (0.05, 0.35)] in a bivariate regression (see Column 1, Table 2) and robust to a number of alternative specifications.

As we noted earlier, a possible concern is that candidates' race or gender is the basis of the appearance advantage. To rule out this possibility, we estimate the bivariate model only for races in which the two candidates shared the same race and gender (Column 2, Table 2). If anything, the estimated effect of appearance advantage is larger in this subset of elections. A statistically and substantively significant effect also remains after controlling for candidate gender, race, and

Fig. 3 Appearance-advantaged candidates in statewide general elections benefit in photo condition. *Note:* Each point represents a state-level general election from 2012. Observations weighted by the total number of respondents reporting a vote choice in the race

Table 2 Appearance-advantaged candidates in statewide general elections benefit in photo condition

Dependent variable: photo condition minus no-photo condition vote share	(1) All	(2) Matched on race and gender	(3) All with controls	(4) Senate and governor	(5) Other	(6) Races with min. face rating at/below 25th percentile	(7) Races with min. face rating above 25th percentile
Appearance advantage (for Republican)	0.20*** (0.07)	0.23* (0.13)	0.15* (0.08)	0.24* (0.14)	0.19** (0.08)	0.16 (0.10)	0.23* (0.13)
Incumbent			0.01 (0.02)				
Female Republican			0.05 (0.04)				
Female democrat			0.00 (0.03)				
White Republican			−0.02 (0.10)				
White democrat			0.02 (0.09)				
Constant	−0.09** (0.04)	−0.11 (0.07)	−0.07 (0.13)	−0.13 (0.08)	−0.07 (0.04)	−0.09* (0.05)	−0.09 (0.07)
Candidates (N)	44	26	44	18	26	19	25
R ²	0.15	0.12	0.21	0.16	0.18	0.14	0.12

This table shows candidate-level regressions (each column showing a separate model) with standard errors in parentheses. The number of participants in this analysis is 1933. Dependent variable: photo condition minus no-photo condition vote share (coded so that higher values indicate greater Republican vote share). We do not cluster the standard errors at the state level because, with only 18 clusters, clustering is unreliable (Angrist and Pischke 2009, ch. 8), but we do cluster by participant and contest in the individual-level analysis shown in Tables 3 and 4

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

incumbency in our full sample (Column 3, Table 2).¹² The finding also holds in U.S. Senate and gubernatorial races (Column 4, Table 2), which tend to be more prominent, as well as lower-ticket races (Column 5, Table 2). In sum, the effect of the photos is robust across a variety of specifications, and holds in lower-salience primary races and in higher-salience general election races.

While generally more professional than primary candidates, particularly unprofessional looking candidates could still drive these findings. As in the primary, however, we observe a similar effect when we exclude races with the

¹² Research has found differential effects of candidate appearance where one or both candidates are female (Chiao et al. 2008; Poutvaara et al. 2009). Unfortunately, we lack a sufficient number of races to shed further light on this topic (half of the races are male-male and the other half are mostly female-male races).

lowest-rated candidates (Table 2, Column 7). These results therefore imply that candidates who look the part drive the main findings, not candidates who play the clown.

External Validity

Across a variety of elections, appearance-advantaged candidates tend to benefit when their photographs appear on the ballot, while appearance-disadvantaged candidates tend to lose support. Since the studies randomly assigned participants to the photo or no-photo condition, they demonstrate that candidate appearance can exert a causal effect on voters' choices, and are therefore inconsistent with concerns about endogeneity and the characterization of previous observational findings of appearance effects as largely or entirely spurious. Do these findings generalize to the real world? We now assess external validity.

Do Ballot Photographs Trigger Memories?

One alternative interpretation that would undermine external validity is that the photos do not directly influence votes, but indirectly do so by triggering memories of the candidates. If candidates who exert greater effort in their campaigns also look better, then the photos of those candidates could trigger positive memories about appearance-advantaged candidates and so could produce the experimental findings reported above even if voters did not actually judge candidates on their appearance. Attractive candidates tend to receive more media attention (Waismel-Manor and Tsfaty 2011), which could exacerbate this potential bias. To assess this possibility, we asked participants in Study 2 factual questions specific to their states' races after the treatment. Viewing the photos failed to increase recall of relevant facts: participants assigned to the photo ballot performed 0.7 % worse on these questions than participants assigned to the control ballot ($p = 0.33$)—the opposite of what we would expect if the photographs cued memories of other information. As we noted above, we also find the appearance effect among nonviable candidates and downballot candidates. Since these candidates usually lack conspicuous campaigns, these findings help further rule out this alternative.¹³

Overestimating or Underestimating the Appearance Effect?

Of course, ballots in the U.S. do not show photos of the candidates (though they do in some countries, such as Brazil). These experiments, therefore, do not necessarily tell us that appearance matters in real American elections, only that it might matter. Are there reasons to think we are substantially overestimating or underestimating the “true” appearance effect?

¹³ Another interpretation of the finding is that candidate age—as discerned from the pictures—influences treated participants to change their votes. Previous studies, however, have found that controlling for age, using various functional forms, leaves the appearance-vote relationship unchanged (Lawson et al. 2010, 581; Todorov et al. 2005).

There are arguments for both positions. One obvious reason we may be overestimating the effect is priming. When the ballot shows voters the pictures, they may more often cast their vote for the appearance-advantaged candidate simply because this makes the candidate's looks more perceptually salient (i.e., at the top of voters' minds) than it would be the case when voters normally go to the polls.¹⁴ A second reason for overestimating is that respondents misjudge their likelihood of voting, especially in the primary, and the photo ballots may disproportionately influence nonvoters.

On the other hand, we could be underestimating the effect because of noncompliance. By showing the photos to the treated group, we intended all treated participants to know how the candidates looked and all control participants not to know, which would constitute perfect compliance. In practice, however, some (even many) voters in the control group doubtless knew what the candidates looked like because the campaigns had already "treated" them. This is an instance of one-way noncompliance. Just as in a medical trial where some in the placebo condition take the real drug instead of the placebo, such noncompliance would lead us to underestimate the magnitude of the treatment effect. The presence of this kind of noncompliance should lead the experimentally-induced appearance effect to diminish as Election Day approaches because, as a result of campaign activity and media coverage, the awareness and hence the influence of the candidates' looks would have grown among the control group. Moreover, since general election campaigns are much more pervasive than primary campaigns, this compliance problem seems more problematic in Study 2.

Is the upward bias from priming and nonvoters greater than the downward bias from noncompliance? Because we cannot measure these directly, we cannot definitively say. Nevertheless, we can indirectly observe the effects (or lack thereof). Take noncompliance. As the campaign heats up, we would expect increasing numbers of control participants to know what candidates look like because the campaign is "treating" them. Consequently, candidate appearance should increasingly predict vote choice in the control group. Since we conducted Study 2 over 17 days, we can test this prediction. When we estimate the effect of appearance separately in the first half of the control group to take the survey and in the second half, we find support for this predicted pattern. As we show in OA section 3.1, the effect of appearance in the control group increases significantly over time. Noncompliance in the control group therefore likely biases our estimate of an appearance-effect downwards. To some extent, we can correct for this downward bias by estimating the photo-ballot's effect earlier in the campaign, when noncompliance is less of a problem (i.e., the control group is not already voting based on appearance). When we re-estimate the effect shown in Column 1 of Table 2 in the first half of the study, the estimate rises considerably, from 0.20 to 0.32 (see OA section 3.1 and 3.2, which discuss statistical significance).

¹⁴ Several classroom and lab studies have conducted experiments on appearance effects (e.g., Johns and Shephard 2007; Rosenberg and McCafferty 1987; Spezio et al. 2008). Our experiment builds on these by examining whether candidate appearance can influence real-world voters' decisions in actual elections.

We may also be able to indirectly assess the upward bias from priming appearance. If we are finding an appearance effect because people cannot resist voting for attractive candidates when they can see their photos, appearance should predict votes consistently in the treated group over time. In fact, however, whereas the appearance-vote relationship increases over time in the control group, it appears to decrease by about half in the treatment group, though this drop is not always statistically significant (see OA section 3.2). This decrease may occur because the campaigns inform voters about other aspects of the candidates, and so voters rely less on appearance as a low-information heuristic as Election Day approaches. This suggests that we are finding more than just an irresistible response to the photographs; thus, such a response may not be leading us to overestimate substantially the appearance effect in the real world. In light of these shifts in the magnitude of the treatment effect over time, we examined whether the importance of other variables, such as party or incumbency, also changed over time, but they do not.¹⁵

We can also indirectly assess the extent of upward bias from including nonvoters. Although we cannot measure which respondents truly did intend to vote, we know that the most politically knowledgeable voters vote at higher rates (Delli Carpini and Keeter 1996, 226–27). As we show in the next section, the photo ballot induced just as much appearance voting among high-knowledge individuals as it did among low-knowledge individuals in the primary study. In the primary study, therefore, we are not overestimating the effect by counting too many low-knowledge respondents as voters. In the general election study, the pattern is different: we primarily observe the effect among low-knowledge respondents. But the implications are less clear because low-knowledge individuals are more likely to vote in general elections.

Assessing whether the experimental effects we observe are externally valid—that is whether we are overestimating or underestimating the real-world effect of appearance on voters—is difficult. The findings in this section, however, should assuage concerns that these estimates represent a large overstatement. Noncompliance (participants already knowing what the candidates look like) probably suppresses the appearance effect estimates and “facial priming” does not seem to entirely overwhelm participants.

As we noted at the beginning of the paper, our research design cannot tell us precisely how much candidate appearance *really* matters to election outcomes, but that it *could* matter. Because we conducted these experiments with actual voters and actual candidates just prior to an election, we can be more confident that the results from our mock ballots reflect voters’ actual use of candidate appearance in real world choices. In other words, these experiments provide evidence for the existence of the effect in real world elections, but the magnitude of that effect in actual elections is less clear.

¹⁵ We also tested for over time patterns in Study 1. Since we conducted Study 1 over fewer days and since primary campaigns usually pale in comparison to general election campaigns, we might not expect to see such patterns, which is what we find.

Substantive Significance: Can Appearance Change Who Wins?

Although our focus is on the existence of the effect, not the magnitude, what would be the consequence for outcomes if the estimated effect manifested itself in actual U.S. elections? To assess this, we consider how election outcomes would have changed in the absence of the appearance effect. We do this by removing the predicted effect of appearance from candidates' actual vote share.¹⁶ In the primary elections (Study 1), the rankings of candidates change in four of the 14 districts (29 %). In the 11 races with more than two candidates, the top two finishers change in two elections (districts 6 and 33 have different second-place finishers). In the general elections (Study 2), the winner changes in six of 44 races (14 %). While these simple extrapolations from the experimental estimates may overstate the real world impact of candidate appearance on election outcomes, even if we conservatively assume that we are overestimating the appearance effect by twofold in the general election, we find that the winner would change in four general election races. Thus, to the extent that advertising in close races can even approach the salience of photographs on a ballot, we would expect a candidate's appearance to influence electoral outcomes.

These experimental effects are similar in size to Atkinson et al.'s observational effects in 1992–2006 Senate races (and these results preclude alternative explanations). They conclude, however, that the appearance effect is too small to change outcomes of any Senate races during that period, which seems inconsistent with our finding that some outcomes would change in the elections we examine. A key difference between our analysis and Atkinson et al. is that we simulate different counterfactuals. They ask what would have happened if the challenger were an average looking candidate, a counterfactual that ignores the impact of incumbent appearance. In our analysis, we simulate what would happen if both candidates had the same appearance.

To investigate the disparity between our results and Atkinson et al., we reanalyzed their data using a specification similar to that used for our experiments (see OA section 5.4 for details). In contrast with their result, we find that the appearance effect (estimated in the following subsets of Senate data) is large enough to change the winner in 9 % of the Senate races between 1992 and 2006. The winner changes in 21 % of races that the *Cook Report* deemed competitive (tossups) at least 1 year before the election. In races with an incumbent, the sample they focus on, 6 % of races would have a different outcome, while the winner would change hands in 11 % of open seat races.

¹⁶ We estimate this by multiplying candidate appearance advantage by the appearance effect reported in Column 1 of Table 2, adding the constant to the outcome, and then subtracting that result from the candidate's actual vote share in the 2012 election.

Candidate Appearance as a Low-Information Heuristic

The candidate-level analyses demonstrate that appearance-advantaged candidates experience greater success when we attach their photographs to the ballot. A natural follow-up question is, “Why?” Research in psychology finds that people rely on appearance most heavily when evaluating others they know little about; that is, they use appearance as a low-information heuristic (Bar et al. 2006; Hassin and Trope 2000; Zebrowitz 1997). Research in political science has only begun to investigate whether the same holds for voting based on appearance, but at least one study suggests that it does. Lenz and Lawson (2011) find that voters rely most on appearance when they watch a relatively high amount of television but know little about politics (see also Brusattin 2012; Riggio and Riggio 2010). To see if this pattern holds in the present studies, we conduct individual-level analyses on the experimental data. These analyses also provide important robustness checks for the candidate-level findings reported above.

Individual-Level Analysis for State-Level General Election Races

We begin with the state-level general election contests (Study 2). If voters use appearance as a low-information heuristic, we might expect the appearance effect to vanish among those with better cues. In particular, politically knowledgeable voters may know enough about the candidates to not fall back on appearance. We might also expect it to diminish among strong partisans (individuals identifying as “strong” Democrats or Republicans). Since the ballot showed party labels and the contests pitted one Democrat against one Republican, strong partisans can rely on party rather than appearance as a cue. Of course, decades of research consistently show that voters rely heavily on partisanship in their voting decisions (Campbell et al. 1960; Schaffner and Streb 2002).

To test these predictions, we estimate these models at the individual level to increase the precision of the estimates. Our dependent variable is whether participants *Vote Republican* in a given race (coded Republican 1, Democrat 0). As in the race-level analysis, we measure *Appearance Advantage* in terms of the Republican candidate (Republican Appearance minus Democratic Appearance) rescaled to a 0–1 range. We use a linear probability model (ordinary least squares) but the results are the same with probit estimation (see OA section 4.1).¹⁷ We cluster the standard errors at the election race and participant level (see OA section 4.1 for alternative specifications that yield similar findings). We measure general political *Knowledge* with a four-item scale and classify as highly knowledgeable participants who answered three or more items correctly (see OA section 4.2 for wording).

Table 3 shows the results. In Column 1, we regress Republican vote choice on an indicator for the photo ballot (*Treatment*), appearance advantage (for Republican), and the interaction between the treatment and appearance advantage

¹⁷ We use linear probability models because they are consistent under weak assumptions and the estimates are simpler to interpret, especially with interaction terms (Ai and Norton 2003).

[*Treatment × Appearance Advantage (for Republican)*]. This interaction is the coefficient of interest—it tests whether candidate appearance predicts vote choices better in the treatment group than in the control group. Consistent with the candidate-level findings above, the interaction is positive and statistically significant. Its size, 0.20, implies that participants in the photo condition are 20-percentage points more likely to vote for the most appearance-advantaged candidates compared to the least advantaged one [95 % CI (0.06, 0.34)].

The next four columns of Table 3 test the low-information heuristic predictions. As expected, candidate appearance has a significantly higher effect for low-knowledge voters (Column 3) than high-knowledge voters, for whom the observed coefficient is near zero (Column 4). Similarly, candidate appearance affects weak partisans and independents (Column 5) more than strong partisans (Column 6), though this difference (0.22 vs 0.18) is not statistically significant.

Because high-knowledge voters may also be strong partisans, Table 4 examines the impact of candidate appearance on vote choice for subsets of participants based on both variables. It shows that knowledge, not partisanship, appears most protective. Among low-knowledge voters, we find that candidate appearance has a large and significant effect on weak partisans and independents (Column 1) and strong partisans (Column 2). Among high-knowledge voters, we find no appearance effect among either group (Columns 3 and 4). Taken together, these findings support the low-information heuristic interpretation of appearance effects. By implication, if voters were better informed about politics, they would not rely on candidate appearance.¹⁸

Individual-Level Analysis for U.S. House Primaries in California

In comparing general and primary elections, there are reasons to expect that partisan and high-knowledge voters would be more likely to rely on candidates' appearance in the primary context. While party labels clearly provide important cues in most general election races, they are significantly less valuable in primary contests, which usually feature multiple candidates from the same party. (Ten of the 14 races included in Study 1 did so.) Furthermore, voters know much less about congressional candidates than gubernatorial and senatorial candidates like those included in our general election study (Krasno 1997). In fact, knowledge is so low in congressional primaries that even politically knowledgeable voters appear largely ignorant of candidates' policy positions, except for what they can glean from the candidate's party (Ahler et al. 2016). Consequently, all voters in primaries, including politically knowledgeable strong partisans, may be more likely to rely on candidate appearance in their voting decisions because they know so little about the candidates and cannot rely on party labels as a guide.

Consistent with this reasoning, voters in Study 1 appear about equally susceptible to candidate appearance, regardless of their partisanship and knowledge: both high-

¹⁸ In OA sections 4.3–4.5, we find evidence that strong partisanship can diminish the appearance effect, especially among high-knowledge individuals. We find this in downballot races (no senatorial and gubernatorial races) and when we substitute local for general knowledge.

Table 3 Voters favor appearance-advantaged candidates at higher rates on the photo ballot

Dependent variable: vote Republican indicator variable	(1) All participants	(2) Matched on race and gender	(3) Low knowledge	(4) High knowledge	(5) Weak/indep.	(6) Strong partisan
Treatment	-0.09* (0.05)	-0.11 (0.07)	-0.20*** (0.07)	0.01 (0.06)	-0.14** (0.06)	-0.05 (0.07)
Appearance advantage (for Republican)	-0.03 (0.06)	0.03 (0.14)	-0.06 (0.08)	-0.01 (0.09)	-0.05 (0.09)	-0.07 (0.08)
Treatment × appearance advantage (for Republican)	0.20*** (0.07)	0.24** (0.11)	0.39*** (0.11)	0.01 (0.08)	0.22*** (0.08)	0.18* (0.10)
Votes (N)	4816	2918	2324	2492	2826	1626
R ²	0.00	0.00	0.01	0.00	0.00	0.00

This table shows individual-level regressions (each column showing a separate model). The dependent variable is coded Republican vote 1 and Democratic vote 0. Constant not shown. Standard errors clustered at the individual and race-level in parentheses

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

Table 4 Low-information voters are most susceptible to candidate appearance

Dependent variable: vote Republican indicator variable	(1) Low knowledge and non-strong partisan	(2) Low knowledge and non-partisan	(3) High knowledge and strong partisan	(4) High knowledge and strong partisan
Treatment	-0.18** (0.07)	-0.24** (0.10)	-0.10 (0.09)	0.13 (0.10)
Appearance advantage (for Republican)	-0.10 (0.11)	-0.05 (0.13)	0.00 (0.13)	-0.08 (0.08)
Treatment × appearance advantage (for Republican)	0.34*** (0.11)	0.47*** (0.18)	0.09 (0.13)	-0.11 (0.09)
Votes (N)	1475	694	1351	932
R ²	0.01	0.03	0.00	0.01

This table shows individual-level regressions (each column showing a separate model). The dependent variable is coded Republican vote 1 and Democratic vote 0. Constant not shown. Standard errors clustered at the individual and race-level in parentheses

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

knowledge and low-knowledge voters who saw the ballot with photos voted for appearance-advantaged candidates more frequently. Similarly, we find that strong partisans (those who place themselves at 1 or 7 on the 7-point party ID scale) vote for appearance-advantaged candidates at a similar rate to independents and weak partisans when they see photos of the candidates. We present these findings in the SI (see section 4.4).

Across both studies, therefore, appearance matters for voters' choices. But voters also seem to rely on appearance less when they know more about politics. The apparent moderating effects of knowledge and (less consistently) partisanship in the higher-salience general election contests, combined with a more wide-ranging effect of candidate appearance in House primaries, supports the view that candidate appearance acts as a low-information heuristic that voters discard when they have are more informed or can use more reliable cues. If voters knew more about the candidates, we might not find that candidate appearance—a piece of information that may carry little signal—influences vote choices. As it is, however, we find that looks sometimes may matter enough to affect electoral outcomes.

Conclusion

This study took earlier research finding a sizeable correlation between candidates' appearances and their electoral fortunes as a starting-point. Given that this result understandably evokes familiar normative concerns about citizen competence and pessimism about democratic accountability, scholars rightly raised questions about the authenticity of the appearance-choice connection, arguing that omitted variables may have generated a relationship that in reality is spurious. Given that candidate appearance correlates with so many other variables—candidate spending, incumbency, and incumbent vulnerability—it seemed plausible that candidate appearance was not a direct cause of voters' behavior.

To overcome the indeterminacy about the meaning of these observational studies, we designed two experimental studies. These randomized experiments supported the notion that candidate appearance can affect voters' stated choices. When we exposed voters (in a treatment group) to photos of the candidates not long before Election Day, they reported intending to vote for appearance-advantaged candidates at higher rates and appearance-disadvantaged candidates at lower rates more often than those (in a control group) shown a ballot without photographs of the candidates. Since candidate effort (or other omitted variables) could not differentially influence voters in the photo conditions compared to voters in the control conditions, we no longer have to worry about the plausible alternative interpretations. By introducing exogeneity into a morass of endogenous relationships, we determine that candidate appearance does seem to have a direct, causal influence on voters.

The effect of candidate appearance on voting in these experiments is robust. It holds up in primary and general election races, among incumbents and challengers, among viable and nonviable candidates, among Democratic and Republican candidates, among up-ballot races (senator and governor) and down-

ballot races (e.g., attorney general), and in contests between candidates matched on race and gender. It also holds up in candidate-level analyses and in individual-level analyses. These effects show up even though voters are casting ballots for real-world candidates in their districts not long before Election Day and even though the ballots provide other information, such as candidate party, incumbency status, and occupation. While our experiments cannot precisely estimate the real world effect, they demonstrate that candidate appearance can influence voters. Even under the conservative assumption that the true effect is half the experimental estimate, candidate appearance changes election outcomes. We also presented evidence that the effect is unlikely to be due to the photo condition triggering memories of candidates or from an irresistible tendency of participants to vote based on candidate looks when the ballot shows photos (priming). Finally, we show the findings with a Mechanical Turk sample and a demographically representative sample of registered voters in California (via SSI).

Some scholars argue that voters generally possess the cues necessary to reach “as-though informed” decisions (Lupia 1994; Lupia and McCubbins 1998; Popkin 1991). Our results present evidence that this is not always the case: candidate appearance, an arguably uninformative cue, can influence voters’ choices and electoral outcomes. Since merely showing photographs of candidates’ faces produces changes in voting intentions, our findings raise questions about the quality of some voter heuristics and, more broadly, about citizens’ ability to hold politicians democratically accountable.

This, of course, all assumes that appearance is not an informative heuristic, that is, that appearance is uncorrelated with traits voters intentionally desire in the candidates. So we should ask, do looks provide valid information about candidates’ abilities? Although some studies find small correlations between attractiveness and IQ scores (e.g., Zebrowitz et al. 2002), most researchers conclude that the inferences about competence people draw from faces fail to correspond with reality (Alley 1988; Cohen 1973; Hassin and Trope 2000; Kalick et al. 1998), though none investigate candidates for political office.¹⁹

Researchers in other fields have found that people routinely make costly decisions based on facial inferences even when those inferences are shown to be uninformative, such as when lending money online (Ravina 2012), eliminating competitors on a television game show (Belot et al. 2012), playing incentivized trust games (Wilson and Eckel 2006), and incentivized public goods games (Andreoni and Petrie 2008).²⁰ Research on the beauty wage premium in the labor market reaches similar conclusions (e.g., Fletcher 2009). Summarizing years of his own and others’ research on the beauty premium, Hamermesh (2011) concludes that it results from taste-based discrimination, which he calls lookism, writing, “We have met the enemy and he is us” (Hamermesh 2011, 122). So, voters may be doing the same

¹⁹ Unpublished work by one of the authors finds that competent looking incumbents are no more effective in Congress, nor are they evaluated as being more effective by peers in the North Carolina legislature (see OA section 6).

²⁰ Other examples include students evaluating their professors’ teaching (Hamermesh and Parker 2005) and economists electing officers to the American Economic Association (Hamermesh 2006).

here: voting on candidate appearance even though candidate appearance is not informative about competence or other relevant traits. Consistent with this interpretation, voters say they place little weight on appearance in their voting decisions (see OA section 5.1). Of course, research on candidate appearance is still in its early stages and our design cannot demonstrate why candidate appearance matters, only that it does.

The notion that candidates' looks should not matter pervades popular conversation about politics. From Bill Clinton's 2012 stump speech charging that Mitt Romney predicated his campaign on "looking like a president" (Nelson 2012) to the 2013 kerfuffle over President Obama's comment that California has the "best looking attorney general," opinion leaders contend that we should not judge our politicians based on how they look. According to our findings, however, this normative ideal fails to describe voters' behavior, especially among those with little other information to go by.

More optimistically, our findings also point to a remedy. In the general election, political knowledge—and, to a lesser degree, partisan attachment—protected voters from this superficial tendency to use looks as a cue (though in primaries, where information is generally scarce, they did not). Opinion leaders who truly believe that appearance should not matter can potentially alleviate this tendency through campaigns that inform voters about candidate characteristics more relevant to effective governance. Appearance will likely cease to matter only when most voters possess more substantive guides to their choices.

Acknowledgments We thank Luke Edwards, Aaron Kaufman, Aidan McCarthy, Tony Valeriano, and Kelsey White for research assistance and students in the fall 2012 Presidential Elections and Democratic Accountability class for help collecting candidate photos and candidate knowledge questions. We are also grateful for comments from David Doherty, Laura Stoker, conference participants at the 2013 Midwest Political Science Association Annual Meeting and the 2013 West Coast Experiments Conference at Stanford University, and research workshop participants at both the University of California, Berkeley and the Wharton School of the University of Pennsylvania. Finally, we thank the Institute of Governmental Studies at the University of California, Berkeley for funding. Replication code and data is available from <https://dataverse.harvard.edu/dataverse/facevalue>.

Compliance with Ethical Standards

Human Rights and Informed Consent The studies were approved by the Committee for the Protection of Human Subjects (CPHS) at the University of California, Berkeley. All procedures performed in studies involving human participants were in accordance with the ethical standards of CPHS and with the 1964 Helsinki declaration and its later amendments or comparable ethical standards.

References

- Ahler, D. J., Citrin, J., & Lenz, G. S. (2016). Do open primaries improve representation? An experimental test of california's 2012 top-two primary. *Legislative Studies Quarterly*, 41(2), 237–268.
- Ai, C., & Norton, E. C. (2003). Interaction terms in logit and probit models. *Economics Letters*, 80(1), 123–129.
- Alley, T. R. (1988). *Social and applied aspects of perceiving faces, resources for ecological psychology*. Hillsdale, N.J.: Lawrence Erlbaum Associates.
- Andreoni, J., & Petrie, R. (2008). Beauty, gender and stereotypes: Evidence from laboratory experiments. *Journal of Economic Psychology*, 29(1), 73–93.

- Angrist, J. D., & Pischke, J.-S. (2009). *Mostly harmless econometrics: An empiricist's companion*. Princeton: Princeton University Press.
- Antonakis, J., & Dalgas, O. (2009). Predicting elections: Child's play! *Science*, 323(5918), 1183.
- Atkinson, M. D., Enos, R. D., & Hill, S. J. (2009). Candidate faces and election outcomes. *Quarterly Journal of Political Science*, 4(3), 229–249.
- Ballew, C. C. I. I., & Todorov, A. (2007). Predicting Political Elections from rapid and unreflective face judgments. *Proceedings of the National Academy of Sciences of the United States of America*, 104(46), 17948–17953.
- Banducci, S. A., Karp, J. A., Thrasher, M., & Rallings, C. (2008). Ballot photographs as cues in low-information elections. *Political Psychology*, 29(6), 903–917.
- Bar, M., Neta, M., & Linz, H. (2006). Very first impressions. *Emotion*, 6(2), 269–278.
- Belot, M., Bhaskar, V., & Van De Ven, J. (2012). Beauty and the sources of discrimination. *Journal of Human Resources*, 47(3), 851–872.
- Benjamin, D. J., & Shapiro, J. M. (2009). Thin-slice forecasts of gubernatorial elections. *The Review of Economics and Statistics*, 91(3), 523–536.
- Berggren, N., Jordahl, H., & Poutvaara, P. (2010). The looks of a winner: Beauty and electoral success. *Journal of Public Economics*, 94(1–2), 8–15.
- Brusattin, L. (2012). Candidate visual appearance as a shortcut for both sophisticated and unsophisticated voters: Evidence from a Spanish online study. *International Journal of Public Opinion Research*, 24(1), 1–20.
- Campbell, A., Converse, P. E., Miller, W. E., & Stokes, D. E. (1960). *The American voter*. Chicago: University Of Chicago Press.
- Chiao, J. Y., Bowman, N. E., & Gill, H. (2008). The political gender gap: Gender bias in facial inferences that predict voting behavior. *PLoS ONE*, 3(10), e3666.
- Cohen, R. (1973). *Patterns of personality judgment*. New York, NY: Academic Press.
- Converse, P. E. (1964). The nature of belief systems in mass publics. In D. E. Apter (Ed.), *Ideology and discontent*. New York: Free Press.
- Delli Carpini, M. X., & Keeter, S. (1996). *What Americans know about politics and why it matters*. New Haven: Yale University Press.
- Fiorina, M. P. (1981). *Retrospective voting in American national elections*. New Haven: Yale University Press.
- Fletcher, J. M. (2009). Beauty vs. brains: Early labor market outcomes of high school graduates. *Economics Letters*, 105(3), 321–325.
- Hall, C. C., Goren, A., Chaiken, S., & Todorov, A. (2009). Shallow cues with deep effects: Trait judgments from faces and voting decisions. In E. Borgida, C. M. Federico, & J. L. Sullivan (Eds.), *Political psychology of democratic citizenship* (pp. 73–99). Oxford: Oxford University Press.
- Hamermesh, D. S. (2006). Changing looks and changing “discrimination”: The beauty of economists. *Economics Letters*, 93(3), 405–412.
- Hamermesh, D. S. (2011). *Beauty pays: Why attractive people are more successful*. Princeton, N.J.: Princeton University Press.
- Hamermesh, D. S., & Parker, A. (2005). Beauty in the classroom: Instructors' pulchritude and putative pedagogical productivity. *Economics of Education Review*, 24(4), 369–376.
- Hassin, R., & Trope, Y. (2000). Facing faces: Studies on the cognitive aspects of physiognomy. *Journal of Personality and Social Psychology*, 78(5), 837–852.
- Johns, R., & Shephard, M. (2007). Gender, candidate image and electoral preference. *British Journal of Politics and International Relations*, 9(3), 434–460.
- Kalick, S. M., Zebrowitz, L. A., Langlois, J. H., & Johnson, R. M. (1998). Does human facial attractiveness honestly advertise health? Longitudinal data on an evolutionary question. *Psychological Science*, 9(1), 8–13.
- Key, V. O. (1966). *The responsible electorate: Rationality in presidential voting, 1936–1960*. Cambridge, MA: Harvard University Press.
- King, G. (1991). “Truth” is stranger than prediction, more questionable than causal inference. *American Journal of Political Science*, 35(4), 1047–1053.
- King, A., & Leigh, A. (2009). Beautiful politicians. *Kyklos*, 62(4), 579–593.
- Klein, M., & Rosar, U. (2005). Physische Attraktivität Und Wahlerfolg. Eine Empirische Analyse Am Beispiel Der Wahlkreiskandidaten Bei Der Bundestagswahl 2002. *Politische Vierteljahresschrift*, 46(2), 263–287.

- Krasno, J. S. (1997). *Challengers, competition, and reelection: Comparing senate and house elections*. New Haven, CT: Yale University Press.
- Kuklinski, J. H., & Quirk, P. J. (2000). Reconsidering the rational public. In A. Lupia & M. D. McCubbins (Eds.), *Elements of reason: Cognition, choice, and the bounds of rationality* (pp. 153–182). New York: Cambridge University Press.
- Laskas, J. M. (2012). Bob dole: Great American. *GQ*(July), 54, 88–90.
- Lawson, C., Lenz, G. S., Myers, M., & Baker, A. (2010). Looking like a winner: Candidate appearance and electoral success in new democracies. *World Politics*, 62(4), 561–593.
- Lenz, G. S. (2012). *Follow the leader? How voters respond to politicians' policies and performance*. Chicago: University of Chicago Press.
- Lenz, G. S., & Lawson, C. (2011). Looking the part: Television leads less informed citizens to vote based on candidates' appearance. *American Journal of Political Science*, 55(3), 574–589.
- Lupia, A. (1994). Shortcuts versus encyclopedias: information and voting behavior in California insurance reform elections. *American Political Science Review*, 88(1), 63–76.
- Lupia, A., & McCubbins, M. D. (1998). *The democratic dilemma: Can citizens learn what they really need to know?*. New York: Cambridge University Press.
- Martin, D. S. (1978). Person perception and real-life electoral behaviour. *Australian Journal of Psychology*, 30(3), 255–262.
- Mattes, K., Spezio, M., Kim, H., Todorov, A., Adolphs, R., & Alvarez, R. M. (2010). Predicting election outcomes from positive and negative trait assessments of candidate images. *Political Psychology*, 31(1), 41–58.
- Nelson, C. M. (2012). In denver, clinton says romney plan 'doesn't add up. *Wall Street Journal Blog*. Accessed April 2013. <http://blogs.wsj.com/washwire/2012/10/31/in-denver-clinton-says-romney-plan-doesnt-add-up>.
- Olivola, C. Y., & Todorov, A. (2010). Elected in 100 Milliseconds: Appearance-based trait inferences and voting. *Journal of Nonverbal Behavior*, 34(2), 83–110.
- Popkin, S. L. (1991). *The reasoning voter: Communication and persuasion in presidential campaigns*. Chicago: University of Chicago Press.
- Poutvaara, P., Jordahl, H., & Berggren, N. (2009). Faces of politicians: Babyfacedness predicts inferred competence but not electoral success. *Journal of Experimental Social Psychology*, 45(5), 1132–1135.
- Ravina, E. (2012). Love & loans: The effect of beauty and personal characteristics in credit markets. Available at SSRN 1101647.
- Riggio, H. R., & Riggio, R. E. (2010). Appearance-based trait inferences and voting: Evolutionary roots and implications for leadership. *Journal of Nonverbal Behavior*, 34(2), 119–125.
- Rosar, U., Klein, M., & Beckers, T. (2008). The frog pond beauty contest: Physical attractiveness and electoral success of the constituency candidates at the North Rhine-Westphalia state election 2005. *European Journal of Political Research*, 47(1), 64–79.
- Rosenberg, S. W., & McCafferty, P. (1987). The image and the vote: Manipulating voters' preferences. *Public Opinion Quarterly*, 51(1), 31–47.
- Schaffner, B. F., & Streb, M. J. (2002). The partisan heuristic in low-information elections. *Public Opinion Quarterly*, 66(4), 559–581.
- Spezio, M. L., Loesch, L., Gosselin, F., Mattes, K., & Alvarez, R. M. (2012). Thin-slice decisions do not need faces to be predictive of election outcomes. *Political Psychology*, 33(3), 331–341.
- Spezio, M. L., Rangel, A., Alvarez, R. M., O'Doherty, J. P., Mattes, K., Todorov, A., et al. (2008). A neural basis for the effect of candidate appearance on election outcomes. *Social Cognitive and Affective Neuroscience*, 3(4), 344–352.
- Todorov, A., Mandisodza, A. N., Goren, A., & Hall, C. C. (2005). Inferences of competence from faces predict election outcomes. *Science*, 308(5728), 1623–1626.
- Waismel-Manor, I., & Tsafati, Y. (2011). Why do better-looking members of congress receive more television coverage? *Political Communication*, 28(4), 440–463.
- Wilson, R. K., & Eckel, C. C. (2006). Judging a book by its cover: Beauty and expectations in the trust game. *Political Research Quarterly*, 59(2), 189–202.
- Zebrowitz, L. A. (1997). *Reading faces*. Boulder, CO: Westview Press.
- Zebrowitz, L. A., Hall, J. A., Murphy, N. A., & Rhodes, G. (2002). Looking smart and looking good. *Personality and Social Psychology Bulletin*, 28(2), 238–249.