

Social Desirability, Hidden Biases, and Support for Hillary Clinton

Ryan L. Claassen, *Kent State University*

John Barry Ryan, *Stony Brook University*

ABSTRACT

An emerging consensus suggests that women are underrepresented in government because of biases in the recruitment process instead of biases at the ballot box. These results, however, are largely for legislative offices, and research suggests that “male” characteristics are generally associated with executive positions like the presidency. At the same time, some research demonstrates social desirability masks gender biases against women who seek the highest office in the land. We use the historic candidacy of Hillary Clinton to examine if she faces hidden biases in either the primaries or the general election. Two different methods for uncovering hidden biases embedded in national surveys demonstrate small hidden biases that are likely electorally inconsequential.

According to Gallup polls, the percentage of Americans that would vote for a well-qualified woman nominated by their party for the presidency stands at 92%—up 60 percentage points from when the question was first asked in 1937 (McCarthy 2015). Meanwhile, the 114th Congress set a new record with 104 women serving. And although political science research finds important biases in the candidate emergence process (Lawless and Fox 2010), many studies find little evidence of gender bias in voting (Dolan 2014; Hayes 2011; Pearson and McGhee 2013; but see Bauer 2016).¹

This election, however, is historic and unexplored territory as the evidence against gender bias in voting is centered on success in legislative elections. While the public generally accepts the idea of a woman representative, the stereotype of an executive often includes “male” characteristics such as aggressiveness and competence with the military (Huddy and Terklidsen 1993; Smith, Paul, and Paul 2007). It is also possible that some previous studies failed to uncover gender bias because individuals actively hide their biases. In this vein, Streb et al. (2008) used a list experiment to demonstrate there are hidden biases against women seeking the presidency as social desirability pressures lead individuals to hide their reluctance to support a woman seeking the White House.

Accordingly, we analyzed the results of two studies conducted on nationally representative samples designed to uncover both

implicit and explicit biases against the first woman to win a major party nomination for president, Hillary Clinton. The first study, conducted early in the campaign, looks for hidden biases during the primaries. The second study, conducted as the primaries concluded, looks for hidden biases in the general election. We find few hidden biases and, when it comes to explicit biases (such as the 8% that tell Gallup they would not vote for a woman), we find robust positive biases (e.g., those more likely to vote for woman) largely counteract negative biases.² The next woman who seeks the White House may face hidden gender biases if she is less well known, but most gender biases for and against Hillary Clinton appear to be fairly open.

STUDY 1: BIASES DURING THE PRIMARIES

There are several ways to search for social desirability bias and we use two different experimental methods in each of our two studies. The method in our first study has been used to detect social desirability with regards to race (Holbrook, Green, and Krosnick 2003) and partisan identification (Klar and Krupnikov 2016). Instead of asking respondents to provide their true responses, respondents are randomly assigned to either answer as if their goal was to make the *best* impression on others or to make the *worst* impression on others. If more people give a particular response in the *Best Impression* condition, then that response is judged to be socially desirable. If more people give a particular response in the *Worst Impression* condition, then the response is socially undesirable. If an equal percentage of people give a response in both conditions, then there is no social desirability bias associated with that response category.

In January 2016, a national sample of 810 respondents participated in this first study via Survey Sampling International (SSI)—a panel-based Internet survey company whose samples are commonly used in political science research (e.g., Kam 2012;

Ryan L. Claassen is professor of political science at Kent State University. He is author of *Godless Democrats and Pious Republicans? Party Activists, Party Capture, and the “God Gap”* published by Cambridge University Press. He may be reached at rclaasse@kent.edu.

John Barry Ryan is associate professor of political science at Stony Brook University. He is co-author of the book *Experts, Activists, and Democratic Politics: Are Electorates Self-Educating?* published by Cambridge University Press. He may be reached at john.ryan@stonybrook.edu.

Krupnikov, Piston, and Bauer 2015). Respondents were randomly assigned to either say who they would support in the Democratic primaries if they wanted to make the best impression or who they would support if they want to make the worst impression. Their choices were the three candidates seeking the nomination at the time: Hillary Clinton, Bernie Sanders, and Martin O'Malley.

The left side of figure 1 presents the results for all respondents, which gives a sense of the larger social desirability among the public. The right side displays how Democratic respondents in particular view the social desirability of voting for each candidate.

gender, race, and age. We found no differences in the level of social desirability for a Clinton candidacy between men and women or whites and minorities. However, among the youngest respondents—those under 25—53% said supporting Clinton gives the worst impression while only 27% said supporting Clinton gives the best impression ($p < .05$). This result goes against the conventional wisdom that respondents will hide negative gender attitudes by falsely stating they support a woman candidate. If there is social desirability bias in polls, this suggests polls would underestimate Clinton support among young people.

If there is social desirability bias in polls, this suggests polls would underestimate Clinton support among young people.

As the figure shows, there is no difference in the percentage of people who say Clinton is the answer that makes the best impression and the people who say Clinton is the answer that makes the worst impression. Instead, the figure suggests participants may feel more social desirability pressures when it comes to stating that one is a Sanders supporter—especially among Democrats.

Note that in both treatments a majority of Democrats respond, “Hillary Clinton.” As a result, there is no generalized social desirability bias among all Democrats because anyone who states they support Clinton only to avoid appearing sexist will be matched with someone who says they do not support Clinton for another reason related to socially desirability.

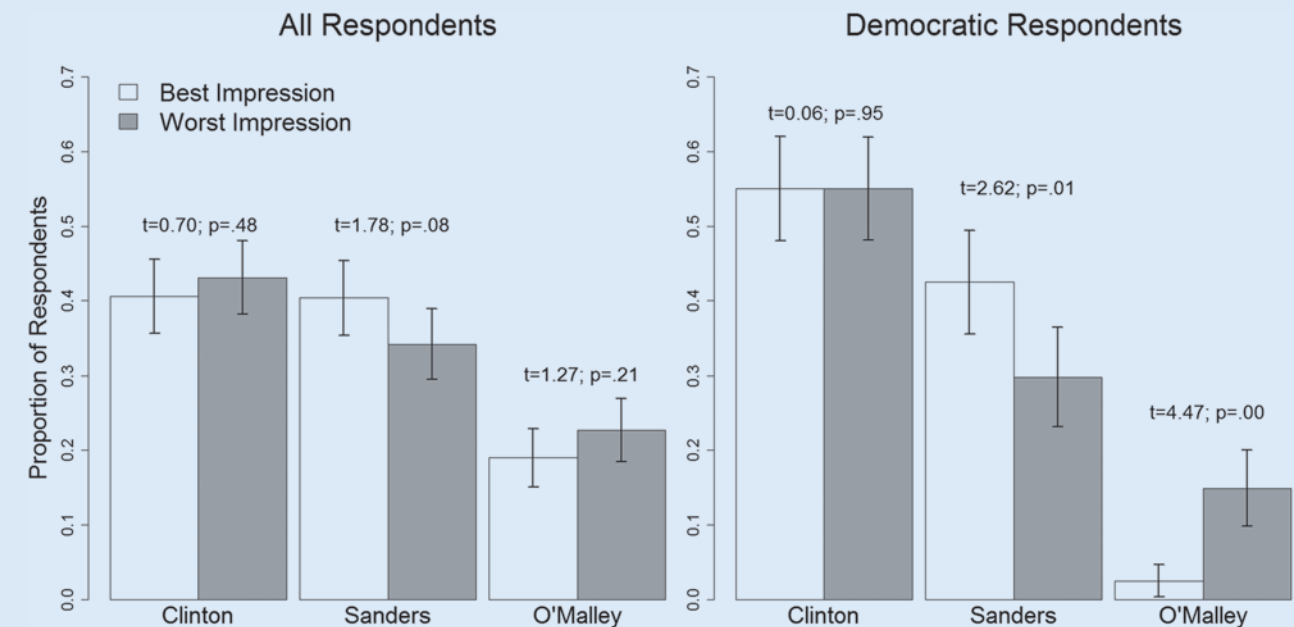
In addition to the simple analysis in figure 1, we attempted to see if the effect among Democrats was conditional on demographics that received a lot of attention during the primary:

STUDY 2: BIASES IN THE GENERAL ELECTION

Our next measure relies on a technique initially implemented by the Obama campaign in 2008 to uncover voters who tell a pollster they support Obama, but in the privacy of the voting booth select McCain (Issenberg 2012). After trying numerous alternatives, the campaign identified a question about whether one’s neighbors would be willing to vote for a well-qualified black candidate for president. This question worked well identifying people “who had high Obama-support scores but ended up backing McCain.” (Issenberg 2012, 296). Issenberg writes, “Something in that response—perhaps a feeling of being liberated to publicly share an unpopular opinion—convinced [the campaign] that the people who acknowledged their neighbors’ racism might really be confessing a view of their own” (Issenberg 2012, 296).

The “would your neighbor be willing to support this candidate” question that the Obama campaign used, was actually

Figure 1 Impressions of Vote Choices in the Democratic Primaries



Error bars are for 95% confidence intervals.

a question Gallup had fielded to measure implicit sexism in voting three years earlier. Only 1 in 10 survey respondents said they were unwilling to support a woman for president, while about 1 in 3 said their neighbors would be unwilling to do so (Feldmann 2007). In fact, although Streb et al. (2008) ultimately used a “list experiment,” they cite the Gallup version of the neighbor question as evidence of social desirability bias.


In the third week of April, 2016, as the nomination contests neared a conclusion, a new set of SSI respondents (N=964) was interviewed to see if the method the Obama campaign used to detect negative racial attitudes could uncover negative gender attitudes. We randomly assigned subjects to answer either the

standard question about willingness to support women candidates or the neighbor version. We also randomly assigned subjects to both versions of a question asking about plans to vote for Hillary Clinton, Donald Trump, or not vote.

By having measures that could capture both gender bias in general and regarding a specific candidate, we can estimate with some precision how many, if any, percentage points Clinton’s sex is likely to cost her. This also will enable us to explore whether the social desirability problems Streb et al. (2008) document regarding a hypothetical female candidate also affect voting plans towards the first actual female major party candidate. If there are hidden biases, we hope to uncover them using the “neighbor” question in an experiment.

Figure 2

Effects of “Neighbor” Question Treatment on Willingness to Support a Female Candidate and Presidential Vote Intention
 A. “Willingness” to Support Female Candidate
 B. Presidential Vote Intention



Error bars are for 95% confidence intervals.

The two versions of the “willingness to support a female candidate” question are below (neighbor version in brackets). Both questions had the same response options.

Regardless of the specific candidates who are running for president, would you [your neighbors] be more or less likely to support a candidate for president who is a woman, or wouldn't this matter to you [your neighbors]?

1. More likely to support a candidate who is a woman.
2. It would not matter.
3. Less likely to support a candidate who is a woman.
4. I don't know.

Figure 2 presents the basic results of these treatments. Figure 2A reveals that the different versions yield significantly different distributions with regards to the willingness to support a woman for president ($\chi^2=82.93$; 3 d.f.; $p<.001$). The difference is mostly movement from “doesn't matter” to the “don't know” category, but nevertheless the result is a distribution that is less congenial toward women candidates. So worries that social desirability bias may be causing underestimates of opposition to women candidates for president receive some validation from our experiment.

We hasten to add, however, that while the less congenial distribution could be problematic for Clinton, it may not necessarily translate into fewer votes. For example, if “don't know” responses are no more predictive of voting plans than “doesn't matter” responses, then the distribution differences wouldn't matter much when it comes to vote intentions. We will investigate this possibility later in this article.

Before that, we look directly at the matchup between Clinton and Trump. Figure 2B graphs the distributional differences

across the two versions—standard and neighbor—of the vote intention question. We can reject the null hypothesis that the distribution was statistically independent of question version ($\chi^2= 8.73$; 2 d.f.; $p<.05$). The candidates' marginals are not dramatically different, and Clinton holds a 20 percentage point lead in both treatments. Respondents, however, are about 7% less likely to say their neighbors “won't vote” compared to their reports about their own plans. In fact, the “won't vote” difference occurs almost entirely among Republicans. With the neighbor version, more Republicans report voting for Trump and Clinton. This suggests that some Republicans felt it was socially undesirable to admit defecting or *voting for their own candidate* in this election. We find no evidence in the vote intention questions, however, that social desirability will lead pollsters to overestimate support for the Clinton candidacy.

Our final analysis estimates the effect of willingness to support women candidates for president, in general, on 2016 vote intentions in order to assess the potential penalties and rewards Clinton faces for being a woman. Because of our 2x2 experimental design, we have roughly one quarter of our sample in each of the two question version combinations. By looking at rewards and penalties in each combination we compute a range of possibilities that varies depending on the question version and the estimated effects. Ultimately, this analysis allows us to see if the Clinton campaign should worry about the large number of subjects who report they “don't know” if their neighbors are willing to support a female candidate.

We examine rewards and penalties for Hillary Clinton using a logit model of vote intention among subjects who stated they would vote for either her (coded 1) or Donald Trump (coded 0).

The key independent variables are: (1) a dummy variable capturing whether the subject received the neighbor vote choice question or the standard version; (2) a dummy variable capturing whether the subject received the neighbor “willingness” question or the standard version; (3) a series of dummy variables measuring how the subject answered the “willingness” question with *Doesn't Matter* as the reference category; (4) interactions between the “willingness” treatment and the “willingness” answers. The model also controls for party, age, education, race, and respondent gender.

The coefficient estimates of the model are shown in table 1. The statistically significant “less likely” and “more likely” dummy variable coefficients reveal the existence of *both* positive and negative biases (relative to those that say “doesn't matter”). The null interaction coefficients suggest

Table 1

Logit Model of Support for Clinton by Question Treatment

	Coef.	Std. Err.	p-value
Vote Intention Neighbor Question	0.102	0.196	0.603
Willingness Neighbor Question	-0.240	0.233	0.302
More Likely	1.536	0.542	0.005
Willingness Neighbor*More Likely	-0.564	0.733	0.442
Less Likely	-1.888	0.595	0.002
Willingness Neighbor*Less Likely	-0.061	0.799	0.939
Don't Know	-1.272	0.792	0.108
Willingness Neighbor*Don't Know	1.201	0.850	0.158
Age	-0.057	0.063	0.367
Party	0.543	0.049	0.000
Woman Respondent	0.128	0.198	0.519
Education	0.219	0.124	0.077
White	-0.899	0.389	0.021
Black	0.143	0.497	0.773
Hispanic	1.025	0.288	0.000
Constant	0.893	0.472	0.058
N		731	
AIC		697.24	

Dependent variable coded 1 if subject supports Clinton and 0 if subject supports Trump. Subjects who state they will not vote are dropped from this analysis.

similar effects regardless of question version—though we note the “don’t know” interaction is the most robust and, because it is nearly equal and opposite the “don’t know” dummy variable coefficient, it indicates the large number of “don’t know” respondents in the neighbor version are indistinguishable from “doesn’t matter” respondents and unlikely to penalize Clinton for being a woman.

Turning to the substance of these effects, first we use the standard questions as a baseline to see if the various neighbor question combinations uncover any hidden biases. In all parts of table 2, the four columns show the values for each of the four question treatment combinations. Part A of table 2 shows the predicted probability of declaring an intent to vote for Clinton by the subject’s answer to the “willingness” question based on the model reported in table 1. Part B displays the distribution of willingness responses within each of the four question treatment combinations. Part C of the table multiplies part A by part B of the table to obtain the proportion of people who are both in that “willingness” category and answer the vote intention question with “Hillary Clinton.” The sums of the columns in the final part of the table are the estimated vote for Clinton in each question treatment combination.

The first column, which displays the result if the respondent is asked the standard version of both questions, is a baseline of support for Clinton that one would see in most polls. This baseline support includes the penalty she suffers because some percentage of people does not want a woman president. The other columns look for the hidden biases. Here, we see that Clinton can lose up to 5 percentage points off her lead, but she still maintains a large lead in all treatment combinations. In the case where the subjects receive only the neighbor version of the vote intent question, she

suffers no penalty at all. These results suggest there could be some hidden gender bias against a woman seeking the presidency, but, in this particular matchup, it is electorally inconsequential.

In figure 3, we also examine how each category of “willingness” affects Clinton’s vote compared to the baseline of “doesn’t matter” within each treatment. This allows us to see the penalties and rewards Clinton suffers and gains based on the distribution of subjects who give an answer other than “doesn’t matter.” To get the values in figure 3, we subtract the probability of voting for Clinton in each category from the probability in “doesn’t matter” (see table 2 part A) then multiply the difference by the proportion of respondents in that category (see table 2 part B).

Not surprisingly, the largest differences in probabilities from the “doesn’t matter” category are in the “less likely” category. Yet, as figure 3 shows, the penalties Clinton suffers are largely canceled out by the rewards she receives from subjects in the “more likely” category. Interestingly, while the substantive effect is smaller among “more likely” voters—because most subjects for whom it does not matter support Clinton—there are more people in the “more likely” category than the “less likely” category, which bolsters the reward. Again, because rewards and penalties tend to balance one another in figure 3, it indicates polls will not overestimate Clinton’s support dramatically and that biases are unlikely to undermine polling leads (assuming they are maintained).

CONCLUSION

In our first study, we found no evidence of social desirability in Clinton support except among the youngest voters. Even among that group, however, the results suggest the bias works in the opposite direction from what conventional wisdom regarding social desirability suggests, with respondents saying that supporting Clinton gives a bad impression because it is socially desirable for young people to say they support Bernie Sanders.

In study 2, the neighbor version of the “willingness to vote for a woman” question creates a less congenial distribution, but most of the difference is in additional “don’t know” responses, which prove benign when it comes to voting intentions. And although the neighbor version does detect some biases standard polling questions overlook, the hidden biases are small and unlikely to reverse a strong Clinton lead. While we caution that other methods might find biases our “neighbor” experiment did not detect, the fact we find biases (hidden and overt, positive and negative) suggests the method is effective. We also note that small overall biases mask the fact that strong biases against women candidates are

Estimates based on table 1’s model and computed using the “margins” command in STATA.

Table 2

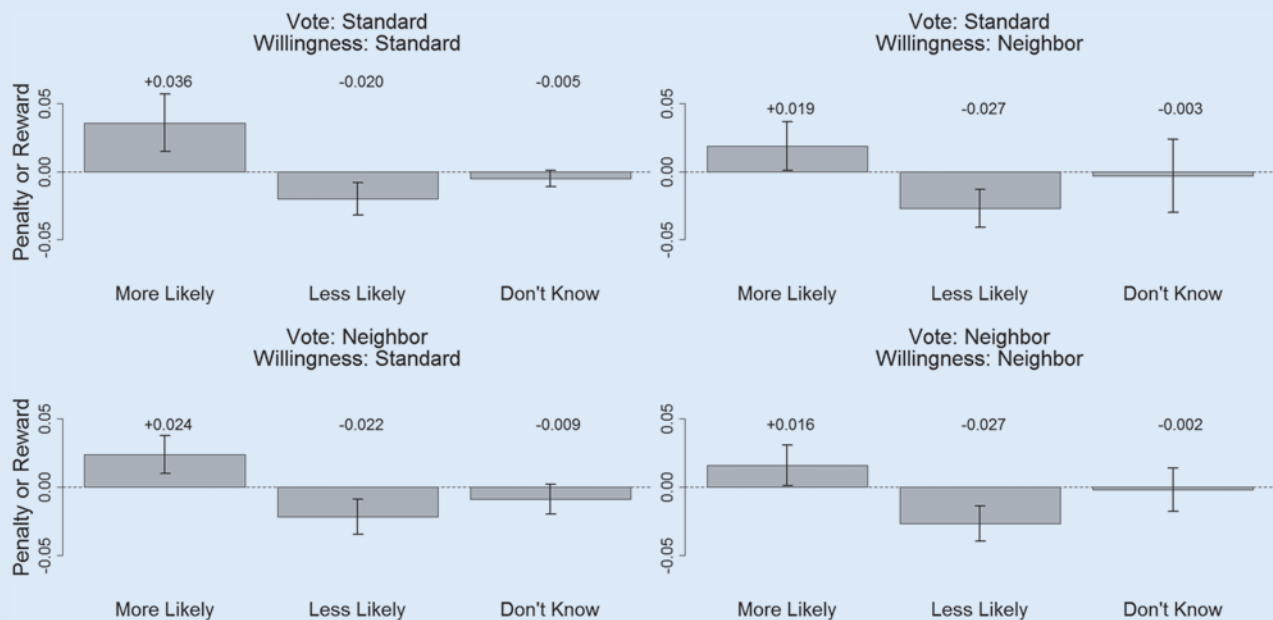
Calculating Support for Clinton in Each Treatment Combination

Treatments				
Willingness Question	Standard	Neighbor	Standard	Neighbor
Vote Question	Standard	Standard	Neighbor	Neighbor
A. Predicted Probability of Clinton Vote Given “Willingness” Answer				
More Likely	0.842	0.746	0.853	0.760
Doesn’t Matter	0.640	0.602	0.655	0.618
Less Likely	0.330	0.283	0.346	0.299
Don’t Know	0.431	0.591	0.448	0.607
B. Proportion Who Gave Each “Willingness” Answer (Columns Add to 1)				
More Likely	0.176	0.137	0.116	0.111
Doesn’t Matter	0.729	0.486	0.762	0.631
Less Likely	0.069	0.087	0.077	0.086
Don’t Know	0.027	0.290	0.044	0.172
C. Estimated Clinton Vote (Part A*Part B)				
More Likely	0.148	0.102	0.099	0.084
Doesn’t Matter	0.466	0.293	0.500	0.390
Less Likely	0.023	0.025	0.027	0.026
Don’t Know	0.011	0.171	0.020	0.104
Vote Total	0.648	0.590	0.645	0.604
Difference from Standard	—	-0.058	-0.003	-0.044

Estimates based on table 1’s model and computed using the “margins” command in STATA.

Figure 3

Rewards and Penalties for Clinton by willingness to Support Answer within Each Treatment Condition



Estimates based on Table 1's model. Values greater than 0 indicate Clinton is rewarded for her gender; values less than 0 indicate Clinton is punished for her gender.

largely neutralized by strong biases for them. It is important to recognize biases exist even if they tend to cancel out.

These results can change depending on the course of the campaign and may not hold for other candidates. Campaign events and issue focus can make a candidate's gender more or less salient (Bauer 2016). Further, Hillary Clinton's 25-year presence in national politics has given her a profile unlike any woman before her in American history. Looking forward, the results of this and future elections will also enable researchers to disentangle factors that are unique to Clinton as a candidate from more general questions of gender bias in voting and, win or lose, this election will likely influence those biases in interesting ways. Having looked for open and hidden gender bias, we conclude it does not pose a major threat to the Clinton campaign at this point in time.

ACKNOWLEDGMENTS

The authors thank Yanna Krupnikov, Jonathan D. Caverley, and Adam Seth Levine for the help with data collection and their comments on this project. ■

NOTES

1. At the same time, concerns about possible gender bias might decrease women's willingness to run for office and change how they conduct their campaigns (Dittmar 2015).
2. Of course, this does not preclude the possibility that voters associate male candidates with different traits, for example, than female candidates. While interesting, our study focuses on the overall effects of gender bias on voting intentions leaving the perceptual foundations of those biases to future work.

REFERENCES

Bauer, Nichole M. 2016. "Emotional, Sensitive, and Unfit for Office? Gender Stereotype Activation and Support for Female Candidates." *Political Psychology* 36 (6): 691–708.

- Dittmar, Kelly. 2015. *Navigating Gendered Terrain: Stereotypes and Strategy in Political Campaigns*. Philadelphia: Temple University Press.
- Dolan, Kathleen. 2014. "Gender Stereotypes, Candidate Evaluations, and Voting for Women Candidates: What Really Matters?" *Political Research Quarterly* 67 (1): 96–107.
- Feldmann, Linda. 2007. "Hillary Clinton Targets Women's Vote." *The Christian Science Monitor* February 1, 2007.
- Hayes, Danny. 2011. "When Gender and Party Collide: Stereotyping in Candidate Trait Attribution." *Politics & Gender* 7: 113–65.
- Holbrook, Allyson L., Melanie C. Green, and Jon A. Krosnick. 2003. "Telephone Versus Face-to-Face Interviewing of National Probability Samples with Long Questionnaires: Comparisons of Respondent Satisficing and Social Desirability Response Bias." *Public Opinion Quarterly* 67 (1): 79–125.
- Huddy, Leonie and Nayda Terklidsen. 1993. "The Consequences of Gender Stereotypes for Women Candidates at Different Levels and Types of Office." *Political Research Quarterly* 46 (3): 503–25.
- Issenberg, Sasha. 2012. *The Victory Lab: The Secret Science of Winning Campaigns*. New York: Crown.
- Kam, Cindy D. 2012. "Risk Attitudes and Political Participation." *American Journal of Political Science* 56 (4): 817–36.
- Klar, Samara and Yanna Krupnikov. 2016. *Independent Politics: How American Disdain for Parties Leads to Political Inaction*. New York: Cambridge University Press.
- Krupnikov, Yanna, Spencer Piston, and Nichole Bauer. 2015. "Saving Face: Identifying Voter Responses to Black Candidates and Female Candidates." *Political Psychology* 37 (2): 253–73.
- Lawless, Jennifer L. and Richard Fox. 2010. *It Still Takes a Candidate: Why Women Don't Run for Office*. New York: Cambridge University Press.
- McCarthy, Justin. 2015. "In US, Socialist Presidential Candidates Least Appealing." *Gallup Report* June 22, 2015.
- Pearson, Kathryn and Eric McGhee. 2013. "What It Takes to Win: Questioning 'Gender Neutral' Outcomes in US House Elections." *Politics & Gender* 9 (4): 439–62.
- Smith, Jessi L., David Paul, and Rachel Paul. 2007. "No Place for a Woman: Evidence for Gender Bias in Evaluations of Presidential Candidates." *Basic and Applied Social Psychology* 29 (3): 225–33.
- Streb, Matthew J., Barbara Burrell, Brian Frederick, and Michael A. Genovese. 2008. "Social Desirability Effects and Support for a Female American President." *Public Opinion Quarterly* 72 (1): 76–89.