Do Elites Discriminate Against Female Political Aspirants? Evidence from a Field Experiment

Online Appendix

KOSTANCA DHIMA*
Texas A&M University

*Ph.D. Candidate, Department of Political Science, Texas A&M University, 2935 Research Parkway, College Station, TX 77843. (dhima1@tamu.edu).

Online Appendix A: What about Quebec?

The email message from the hypothetical student in my audit experiment was written in English. I chose to send the emails in English to replicate as closely as possible the empirical strategy employed by Kalla, Rosenbluth and Teele (2018). This has implications for Quebec with its large French-speaking population. I now address some of these implications here.

One implication is that the overall response rate for the legislators in Quebec will be lower than the response rate for the legislators elsewhere in Canada. This is borne out in the data. There are 302 Quebec legislators in my sample, 89 (29%) of whom responded to my email. This response rate is significantly lower than the overall response rate in Canada as a whole, which was 49%. Notably, though, the response rate in Quebec was still higher than the response rate of 26% in the United States.

While the overall response rate in Quebec may well have been higher had I sent my emails in French, it's unclear whether this would affect any of the inferences that I draw with respect to the *directionality* of the responses. One way to evaluate this is to look to see if the pattern of responses from Quebec is different from that found elsewhere in Canada. In Table 1, I provide information about the response rates by treatment name and legislator gender *for Quebec only*. The results indicate that legislators in Quebec exhibit a pro-female bias of 7.1 percentage points. This is almost identical to the 7 percentage points pro-female bias observed in the full sample. As with the full sample, the female legislators in Quebec exhibit a higher pro-female bias (12 percentage points) than their male counterparts (5 percentage points). The magnitudes of the pro-female bias exhibited by the female and male legislators in Quebec are again almost identical to those observed in the full sample, where female legislators exhibited a pro-female bias of 11 percentage points and the male legislators exhibited a pro-female bias of 6 percentage points. Unlike with the full sample, the pro-female biases in Quebec do not reach conventional levels of statistical significance. However, this is likely due to the much smaller sample size used in the Quebec-only comparisons. The bottom line here is that the legislators in Quebec who did respond to my email message did so in an almost identical way to the legislators in the other provinces.

A second way to evaluate whether my inferences are affected by my decision to send the email message in English is to see what happens if I exclude the responses from Quebec. In Table 2, I provide information about the response rates by treatment name and legislator gender in the full sample *with Quebec excluded*. The results indicate that the legislators in the full sample with Quebec excluded exhibit a pro-

Table 1: Response Rates by Treatment Name and Legislator Gender: Only Quebec

	Male Name	Female Name	Difference	p-value of Difference
All Representatives	26% [19%, 33%] $N = 154$	33% [25%, 41%] $N = 148$	7.1 [−5, 19]	0.22
Female Legislators	24% [12%, 36%] $N = 54$	36% [21%, 50%] $N = 45$	12 [−6, 29]	0.19
Male Legislators	27% [18%, 36%] $N = 100$	32% [23%, 41%] $N = 103$	5 [−9, 19]	0.47

Note: The first two columns of Table 1 show the response rates to the email messages sent from male and female students for different sets of elected representatives. The third column indicates the percentage-point-difference in response rates, with positive differences indicating a pro-female bias. 95% confidence intervals are shown in square brackets. The fourth column indicates whether the differences are statistically significant. The p-values come from a linear probability model where *Email Response* is the dependent variable and the model includes strata fixed effects (gender of legislator) and robust standard errors clustered at the email account level.

Table 2: Response Rates by Treatment Name and Legislator Gender: Excluding Quebec

	Male Name	Female Name	Difference	p-value of Difference
All Legislators	49% [45%, 53%] $N = 734$	56% [53%, 60%] $N = 738$	7.2 [1, 13]	0.02
Female Legislators	47% [40%, 54%] $N = 205$	58% [51%, 64%] $N = 213$	11 [1, 19]	0.03
Male Legislators	50% [45%, 54%] $N = 529$	56% [51%, 60%] $N = 525$	6 [−1, 12]	0.07

Note: The first two columns of Table 2 show the response rates to the email messages sent from male and female students for different sets of elected representatives. The third column indicates the percentage-point-difference in response rates, with positive differences indicating a pro-female bias. 95% confidence intervals are shown in square brackets. The fourth column indicates whether the differences are statistically significant. The p-values come from a linear probability model where *Email Response* is the dependent variable and the model includes strata fixed effects (gender of legislator) and robust standard errors clustered at the email account level.

female bias of 7.2 percentage points. This is almost identical to the 7 percentage points pro-female bias observed in the full sample. Again, the female legislators exhibit a higher pro-female bias (11 percentage points) than their male counterparts (6 percentage points). The magnitudes of these pro-female biases exhibited by the female and male legislators are identical to those observed in the full sample, where female legislators exhibited a pro-female bias of 11 percentage points and the male legislators exhibited a pro-female bias of 6 percentage points. The pattern of statistical significance across the three groups – All Legislators, Female Legislators, and Male Legislators — is also identical to that found in the full sample.

We can also look at how the response rates differ across the different levels of government when we exclude the responses from Quebec. This information is shown in Table 3. The main thing to note here is that there is always a pro-female bias at each level of government. As in the full sample that includes Quebec, this pro-female bias is larger for female legislators than male legislators at the municipal and provincial levels. At the municipal level, the pro-women bias exhibited by female legislators (11 percentage points) is 2.8 times larger than that exhibited by male legislators (4 percentage points). At the provincial level, the pro-women bias exhibited by female legislators (14 percentage points) is 2.8 times larger than that exhibited by male legislators (5 percentage points). There is no substantive difference in the magnitude of the pro-women bias across female and male legislators at the federal level.

In summary, these additional analyses show that my inferences are robust to looking only at Quebec or excluding Quebec from the full sample. As a result, my decision to send my email message only in English is not problematic.

Table 3: Response Rates by the Treatment Name, Level of Office, and Legislator Gender: Excluding Quebec

	Male Name	Female Name	Difference	p-value of Difference
Municipal:				
All Legislators	58% [53%, 64%] $N = 339$	64% [59%, 69%] $N = 342$	6 [−0.2, 12]	0.06
Female Legislators	53% [43%, 63%] $N = 94$	64% [54%, 74%] $N = 89$	11 [−1, 22]	0.07
Male Legislators	60% [54%, 67%] $N = 245$	64% [58%, 70%] $N = 253$	4 [−2, 10]	0.20
Provincial:				
All Legislators	50% [44%, 57%] $N = 227$	58% [52%, 65%] $N = 240$	8 [−2, 18]	0.12
Female Legislators	49% [37%, 61%] $N = 69$	63% [52%, 74%] $N = 79$	14 [−4, 32]	0.11
Male Legislators	51% [43%, 59%] $N = 158$	56% [48%, 64%] $N = 161$	5 [−6, 16]	0.34
Federal:				
All Legislators	29% [22%, 35%] $N = 168$	35% [28%, 43%] $N = 156$	6 [−4, 17]	0.20
Female Legislators	31% [16%, 46%] $N = 42$	36% [21%, 50%] $N = 45$	5 [−14, 23]	0.61
Male Legislators	28% [20%, 36%] $N = 126$	35% [26%, 44%] $N = 111$	7 [−4, 18]	0.18

Note: The first two columns of Table 3 show the response rates to the email messages sent from male and female students for different sets of elected representatives separated by level of office. The third column indicates the percentage-point-difference in response rates, with positive differences indicating a pro-female bias. 95% confidence intervals are shown in square brackets. The fourth column indicates whether the differences are statistically significant. The p-values come from a linear probability model where *Email Response* is the dependent variable and the model includes strata fixed effects (gender of legislator) and robust standard errors clustered at the email account level.

Online Appendix B: Sample of Legislators

I obtained my sample of Canadian legislators from the Represent Civic Information API. The original sample consisted of $1,936$ legislators. After dropping legislators for whom an email address was not provided and 'duplicate' legislators who held multiple official roles, I obtained a sample of $1,779$ unique legislators. In this sample, there were 854 municipal legislators (28.6% women), 591 provincial legislators (31.6% women), and 334 federal legislators (26.7% women). Ultimately, five of these legislators had an invalid email address and so my results are based on a final sample of $1,774$ legislators who received an email.

Although my sample does not represent the full number of legislators across the three levels of government in Canada, the API states that it's "the most comprehensive source in Canada for elected officials and electoral districts." My sample of Canadian legislators is over nine times larger than that used in a recent audit study looking at the responsiveness of Canadian legislators to constituent requests for assistance (Loewen and MacKenzie, 2019). The API does not indicate why it has information about some elected representatives and not others. It appears that there is no information for legislators at the municipal and 'provincial' level in the three Canadian territories; the federal-level legislators are, however, included in the sample. I see no obvious reason why the legislators with contact information provided by the API would be qualitatively different when it comes to exhibiting a gender bias from the legislators for whom the API does not provide contact information.

My sample includes elected representatives from the following provincial legislatures: Legislative Assembly of Alberta, Legislative Assembly of British Colombia, Legislative Assembly of Manitoba, Nova Scotia House of Assembly, Legislative Assembly of Ontario, Legislative Assembly of Prince Edward Island, Assemblée nationale du Québec, and Legislative Assembly of Saskatchewan. It includes elected representatives from the following municipal councils: Ajax Town Council, Belleville City Council, Brampton City Council, Brantford City Council, Burlington City Council, Burnaby City Council, Caledon Town Council, Calgary City Council, Cambridge City Council, Cape Breton Regional Council, Charlottetown City Council, Conseil municipal de Beaconsfield, Conseil municipal de Dorval, Conseil municipal de Gatineau, Conseil municipal de Kirkland, Conseil municipal de Laval, Conseil municipal de Lévis, Conseil municipal de Montréal, Conseil municipal de Montréal-Est, Conseil municipal de Pointe-Claire, Conseil municipal de Saint-Jean-sur-Richelieu, Conseil municipal de Saint-Jérôme, Coquitlam City Council, County of Grande Prairie No. 1 Council, Edmonton City Council, Fredericton City Council, Georgina Town Council,

Grande Prairie City Council, Greater Sudbury City Council, Grimsby Town Council, Guelph City Council, Haldimand County Council, Halifax Regional Council, Hamilton City Council, Kelowna City Council, King Township Council, Kingston City Council, Kitchener City Council, Lambton County Council, Langley Township Council, Lethbridge City Council, Lincoln Town Council, London City Council, Mississauga City Council, Moncton City Council, Newmarket Town Council, Niagara Regional Council, North Dumfries Township Council, Oakville Town Council, Oshawa City Council, Ottawa City Council, Peel Regional Council, Peterborough City Council, Pickering City Council, Richmond City Council, Richmond Hill Town Council, Saanich District Council, Saint John City Council, Saskatoon City Council, St. Catharines City Council, Strathcona County Council, Surrey City Council, Thunder Bay City Council, Toronto City Council, Uxbridge Township Council. Vancouver City Council, Victoria City Council, Waterloo City Council, Waterloo Regional Council, Welland City Council, Whitby Town Council, Windsor City Council, and Winnipeg City Council.

Online Appendix C: Staff Responses

As I noted in the main text, there's no guarantee that the legislator is the person to receive and respond to the student's email message. This is true for all audit experiments of this type and not just the one discussed here. Thus, the unit of analysis is technically the email address of the legislator and not the legislator themself.

It's often possible to identify when a staff member has sent the email response rather than the legislator. This is because the response either comes from a different email address or the person responding introduces themselves as the person responsible for communicating on behalf of the official. On this basis, it appears that about 38% of the email responses come from staff members. In general, we would not expect staff members to express their own views or opinions when engaging in official business. Instead, we'd expect them to express views that are consistent with those of the legislator for whom they work. Of course, if this is true, then we should see similar results to those reported in the main text when we exclude the email responses that have been identified as coming from staff members. This is exactly what I find. In Table 4, I provide information about the response rates by treatment name and legislator gender *when the staff responses are excluded*. The key thing to note is that there's still always a statistically significant pro-female bias. The overall pro-female bias is 6 percentage points in this smaller sample and 7 percentage points in the 'full' sample. Email accounts associated with both female and male legislators continue to exhibit a pro-female bias in this smaller sample. In contrast to the full sample, the pro-female bias exhibited by email accounts associated with female legislators is not larger than that exhibited by email accounts associated with male legislators.

As one might expect, the extent to which staff members send the email responses increases as we move from the local to the national level. While 17% of the responses at the municipal level came from staff, 79% of the responses at the federal level did so. This may help to explain why my results in the full sample, particularly with respect to the gender of the legislator, are weakest at the federal level — there's evidence of a pro-female bias for all legislators at the federal level, female and male legislators but it's never statistically significant. However, I can't rule out the possibility that the lack of statistical significance at the federal level is simply a result of the significantly smaller sample size at this level of government.

Table 4: Response Rates by Treatment Name and Legislator Gender: Excluding Staff Responses

	Male Name	Female Name	Difference	p-value of Difference
All Representatives	34% [30%, 37%] $N = 735$	40% [37%, 44%] $N = 708$	6 [2, 12]	0.01
Female Legislators	33% [26%, 39%] $N = 221$	40% [33%, 46%] $N = 197$	7 [−1, 15]	0.08
Male Legislators	34% [30%, 38%] $N = 514$	41% [36%, 45%] $N = 511$	7 [0.2, 13]	0.04

Note: The first two columns of Table 4 show the response rates to the email messages sent from male and female students for different sets of elected representatives. The third column indicates the percentage-point-difference in response rates, with positive differences indicating a pro-female bias. 95% confidence intervals are shown in square brackets. The fourth column indicates whether the differences are statistically significant. The p-values come from a linear probability model where *Email Response* is the dependent variable and the model includes strata fixed effects (gender of legislator) and robust standard errors clustered at the email account level.

Online Appendix D: Including Responses After the 2 Week Cutoff

In the main text, I focus on only those email responses that came during the first two weeks after my initial email was sent. In this appendix, I show that my inferences are robust to the inclusion of the email responses that came after the two week cutoff. In total, 50 responses came after the two week cutoff: 11 were from legislators at the municipal level, 24 from legislators at the provincial level, and 15 from legislators at the federal level.

In Table 5, I provide information about the response rates by treatment name and legislator gender in the sample *with later responses included*. The results continue to show a pro-women bias of 7 percentage points, which is identical to the 7 percentage points pro-women bias observed in the original sample. Female legislators continue to exhibit a higher pro-women bias (10 percentage points) than their male counterparts (5 percentage points). The magnitudes of these pro-women biases are almost identical to those observed in the original sample, where female legislators exhibited a pro-women bias of 11 percentage points and the male legislators exhibited a pro-women bias of 6 percentage points.

Table 5: Response Rates by Treatment Name and Legislator Gender: Including Later Responses

	Male Name	Female Name	Difference	p-value of Difference
All Legislators	48% [45%, 52%] $N = 888$	55% [51%, 58%] $N = 886$	7 [1, 12]	0.02
Female Legislators	47% [41%, 53%] $N = 259$	57% [51%, 63%] $N = 258$	10 [2, 18]	0.02
Male Legislators	49% [45%, 53%] $N = 629$	54% [50%, 58%] $N = 628$	5 [−1, 11]	0.10

Note: The first two columns of Table 5 show the response rates to the email messages sent from male and female students for different sets of legislators. The third column indicates the percentage-point-difference in response rates, with positive differences indicating a pro-female bias. 95% confidence intervals are shown in square brackets. The fourth column indicates whether the differences are statistically significant. The p-values come from a linear probability model where *Email Response* is the dependent variable and the model includes strata fixed effects (gender of legislator) and robust standard errors clustered at the email account level.

In Table 6, I show how the response rates differ across the different levels of government when I include responses that were received after two week cutoff. Again, we see a pro-women bias at each level of government. As before, this pro-women bias is larger for female legislators than male legislators at the municipal and provincial levels.

In Table 7, I look at how the response rates differ across the different types of political parties. As in the original sample, legislators from the left-leaning parties respond at higher rates (61%) to female political aspirants than legislators from the right-leaning parties (52%). The pro-women bias exhibited by the legislators in the left-leaning parties is again much larger — over three time as large — as that exhibited by the legislators in the right-leaning parties.

In sum, my results are robust to the inclusion of the email responses that came after the two week cutoff used in the main text.

Table 6: Response Rates by the Treatment Name, Level of Office, and Legislator Gender: Including Later Responses

	Male Name	Female Name	Difference	p-value of Difference
Municipal:				
All Legislators	55% [51%, 60%] $N = 421$	62% [57%, 66%] $N = 429$	7 [2, 11]	0.01
Female Legislators	50% [41%, 59%] $N = 123$	61% [52%, 70%] $N = 119$	12 [0.4, 23]	0.04
Male Legislators	58% [52%, 63%] $N = 298$	62% [57%, 67%] $N = 310$	4 [−1, 10]	0.12
Provincial:				
All Legislators	47% [41%, 53%] $N = 295$	54% [48%, 60%] $N = 295$	7 [−3, 17]	0.19
Female Legislators	48% [38%, 59%] $N = 93$	60% [50%, 70%] $N = 93$	12 [−4, 28]	0.15
Male Legislators	47% [40%, 53%] $N = 202$	51% [44%, 58%] $N = 202$	4 [−7, 16]	0.44
Federal:				
All Legislators	33% [26%, 40%] $N = 172$	38% [30%, 45%] $N = 162$	4 [−7, 15]	0.42
Female Legislators	37% [22%, 52%] $N = 43$	39% [24%, 54%] $N = 46$	2 [−20, 24]	0.86
Male Legislators	32% [24%, 40%] $N = 129$	37% [28%, 46%] $N = 116$	5 [−5, 16]	0.32

Note: The first two columns of Table 6 show the response rates to the email messages sent from male and female students for different sets of legislators separated by level of office. The third column indicates the percentage-point-difference in response rates, with positive differences indicating a pro-female bias. 95% confidence intervals are shown in square brackets. The fourth column indicates whether the differences are statistically significant. The p-values come from a linear probability model where *Email Response* is the dependent variable and the model includes strata fixed effects (gender of legislator) and robust standard errors clustered at the email account level.

Table 7: Response Rates by the Treatment Name, Party Ideology, and Legislator Gender: Including Later Responses

Received Email	Male Sender	Female Sender	Difference	p-value of Difference
Left Party:				
All Legislators	51% [41%, 62%] $N = 90$	61% [51%, 70%] $N = 99$	10 [−4, 23]	0.17
Female Legislators	54% [37%, 70%] $N = 39$	59% [45%, 73%] $N = 49$	5 [−16, 27]	0.62
Male Legislators	49% [35%, 63%] $N = 51$	62% [48%, 76%] $N = 50$	13 [−8, 34]	0.21
Center Party:				
All Legislators	40% [32%, 48%] $N = 141$	47% [39%, 55%] $N = 144$	7 [−5, 20]	0.24
Female Legislators	38% [22%, 53%] $N = 40$	56% [40%, 71%] $N = 45$	18 [2, 34]	0.03
Male Legislators	41% [31%, 50%] $N = 101$	43% [33%, 53%] $N = 99$	3 [−12, 18]	0.69
Right Party:				
All Legislators	49% [41%, 58%] $N = 138$	52% [44%, 61%] $N = 124$	3 [−11, 17]	0.69
Female Legislators	65% [44%, 86%] $N = 23$	63% [42%, 83%] $N = 24$	−2 [−26, 20]	0.81
Male Legislators	46% [37%, 55%] $N = 115$	50% [40%, 60%] $N = 100$	4 [−11, 19]	0.60

Note: The first two columns of Table 7 show the response rates to the email messages sent from male and female students for different sets of legislators separated by party ideology. The third column indicates the percentage-point-difference in response rates, with positive differences indicating a pro-female bias. 95% confidence intervals are shown in square brackets. The fourth column indicates whether the differences are statistically significant. The p-values come from a linear probability model where *Email Response* is the dependent variable and the model includes strata fixed effects (gender of legislator) and robust standard errors clustered at the email account level.

Online Appendix E: Analyzing the Content of the Email Responses

In this appendix, I further analyze the content of the email responses from the legislators. As I note in the main text, there are different ways to analyze the content of the email responses. We can think of these approaches as the *'micro'-mentorship approach*, the *quality of response approach*, and the *length of response approach*. Below, I describe each of the approaches in more detail and then present the results. It's important to note, though, that to avoid possible post-treatment bias, the analyses presented in this appendix are *not* conditional on having received an email response. As Coppock (2019, 1) explains, an email "response is a post-treatment outcome" and "conditioning on post-treatment outcomes 'de-randomizes' an experiment in the sense that the resulting treatment and control groups no longer have potential outcomes that are in expectation equivalent." To avoid conditioning on having received a response, I redefine the outcome variables I'm about to create to include non-responses in the 0 category. What this means in practice is that an outcome variable that, say, indicates whether helpful advice was provided is coded as 0 if either (i) no response was received or (ii) a response was received but did not provide helpful advice; the outcome variable is coded 1 only if helpful advice was provided.

'Micro'-mentorship Approach

The *'micro'-mentorship approach* to evaluating the content of the email responses comes from Kalla, Rosenbluth and Teele (2018). They created several indicators of micro-mentorship by qualitatively coding whether email responses (1) were meaningful, (2) provided praise, (3) offered help, (4) provided a warning, or (5) provided advice.

1. *Meaningful Response*: The email response was coded 1 it was determined to "contain real content".
2. *Praise*: The email response was coded 1 if it either "praises student for an interest in a political career" or provides vague praise such as "good luck with everything" or "hope this helps."
3. *Offer to help*: The email response was coded 1 if it indicated a willingness on the part of the legislator to meet, talk on the phone, or email further, or a general offer to follow up such as "If you have any other specific questions, please let me know."
4. *Warning*: The email response was coded 1 if it contained an explicit statement not to run, an encouragement to consider other career paths, or a warning of time commitment, work-life balance challenges, the difficultly of finding time for family, the challenges of fundraising, or the loss of privacy.
5. *Substantive Advice*: The email response was coded 1 if it contained either practical advice (e.g., motivational advice, get a business job, go to law school, get a different type of job, become involved in

local community groups, attend local party or political meetings, volunteer, get a mentor, fundraising advice, run for student government, learn about the issues, get a good education, always put your values first, stay loyal to your political party) or personality/image advice (e.g., always have a professional appearance, have thick skin, learn to be extroverted, learn to deal with conflict).

Quality of Response Approach

The *quality of response approach*, which has recently been developed by Costa (2020), classifies "quality" and "satisfying" responses as those that are "not automated, answers the question, arrives promptly, is at least 400 characters in length, and includes a named greeting, invitation to follow up, link to a website, and a sign-off" (Costa, 2020, 15). Responses that have all of these criteria are coded as 1. If a response is missing one of these criteria, then a specified amount is subtracted or 'discounted' from 1. The specific discount formula used by Costa (2020) is

$$\begin{aligned} \textit{Discount} = \textit{Automated} \times 0.209 + \textit{No named greeting} \times 0.048 + \textit{No invite follow up} \times 0.0613 \\ + \textit{Did not answer question} \times 0.116 + \textit{No website link} \times 0.0533 + \textit{No sign-off} \times 0.0333 \\ + \textit{Characters} < 400 \times 0.0005 + \textit{Days until response} \times 0.0033. \end{aligned}$$

The components of the discount formula are defined as:

- *Automated*: Since I classify automated responses as non-responses, any automated response in my sample actually receives a discount factor of 1. In other words, the overall quality of response for automated responses is always 0.
- *No named greeting*: This is coded 1 if the response is not personalized and 0 otherwise.
- *No invite follow up*: This is coded 1 if the response did not invite the student to follow up with further questions and 0 otherwise.
- *Did not answer question*= This is coded 1 if the response did not answer the student's question and 0 otherwise.
- *No website link*: This is coded 1 if the response did not include a website link and 0 otherwise.
- *No sign-off*: This is coded 1 if the response did not include a sign-off (eg. Sincerely, best, regards, etc) and 0 otherwise.
- *Characters*<400: This is coded 1 if the response was fewer than 400 characters and 0 otherwise.
- *Days until response*: This is a count of the number of days until a response came; this is capped at 30 days.

To calculate the overall *quality* of an email response, I simply subtract the discount scores from 1 for all responses that were received. To avoid post-treatment bias, I code "no replies" as having a discount factor of 1 and, thus, an overall quality score of 0.

Length of Response Approach

The *length of response approach* assumes that longer responses are more substantively meaningful than shorter ones. For this metric, I evaluate (i) the word count, (ii) the log word count, and (ii) the number of characters in a response. To avoid post-treatment bias, I code "no replies" as having 0 words/characters and as $log(1)$ for the log word count.

Results

Having described the three different approaches, we can now turn to the results. In Table 8, I provide information about the content of the email responses by treatment name. Overall, I find a pro-women bias across all nine of the different metrics – the numbers in the *Difference* column are always positive. This pro-women bias is statistically significant in eight of these nine cases. The only metric on which the pro-women bias is not statistically significant is when we look at whether the email responses provides praise.

In Table 9, I provide information about the content of the email responses by treatment name and legislator gender. In regards to the *'micro'-mentorship approach*, I find a pro-women bias exists among both female and male legislators across all metrics except for *Praise*. This pro-women bias is larger and only statistically significant among male legislators. When it comes to the *quality of response approach*, I find a statistically significant pro-women bias among both female and male legislators. On the *length of response approach*, I also find a pro-women bias, with email responses to female students being significantly longer than those to male students. This pro-women bias is only statistically significant for male legislators.

Table 8: Quality of Response by Treatment Name

	Male Name	Female Name	Difference	p-value of Difference
	$N = 888$	$N = 886$		
"Micro"-mentorship Approach				
Meaningful Response	12% [9%, 14%]	16% [14%, 18%]	4 [1, 7]	0.01
Praise	8% [7%, 10%]	10% [8%, 12%]	2 [−1, 4]	0.35
Offer to help	23% [20%, 26%]	29% [26%, 32%]	6 [0.5, 11]	0.03
Warning	1% [0.4%, 2%]	3% [2%, 4%]	2 [0.3, 4]	0.03
Substantive advice	11% [9%, 13%]	14% [12%, 17%]	3 [1, 6]	0.01
Quality of Response Approach				
Quality Response	0.31 [0.29, 0.34]	0.37 [0.34, 0.39]	0.06 [0.02, 0.09]	0.00
Length of Response Approach				
Word count	36 [30, 41]	46 [39, 52]	10 [3, 17]	0.01
Log word count	0.76 [0.71, 0.82]	0.90 [0.84, 0.96]	0.14 [0.05, 0.22]	0.00
Character count	161 [137, 184]	206 [176, 235]	45 [14, 76]	0.01

Note: The first two columns of Table 8 evaluate the content of the email responses depending on whether the original email message was sent by a male or female student. The third column indicates the difference in the content of the email response, with positive differences indicating a pro-women bias. 95% confidence intervals are shown in square brackets. The fourth column indicates whether the differences are statistically significant.

Table 9: Quality of Response by Treatment Name and Legislator Gender

	Male Legislators				Female Legislators			
	Male Sender	Female Sender	Difference	p-value of Difference	Male Sender	Female Sender	Difference	p-value of Difference
	$N = 629$	$N = 628$			$N = 259$	$N = 258$		
"Micro"-mentorship Approach								
Meaningful Response	11% [9%, 14%]	16% [13%, 19%]	5 [1, 9]	0.01	13% [9%, 17%]	15% [10%, 19%]	2 [−4, 8]	0.49
Praise	7% [5%, 10%]	9% [7%, 11%]	2 [−1, 5]	0.28	11% [7%, 15%]	11% [7%, 15%]	0 [−6, 6]	0.99
Offer to help	25% [21%, 28%]	30% [27%, 34%]	5 [0.3, 11]	0.04	19% [14%, 24%]	26% [21%, 31%]	7 [−2, 16]	0.12
Warning	1% [0.2%, 2%]	4% [2%, 5%]	3 [0.5, 5]	0.02	1% [−0.2%, 2%]	2% [0.03%, 3%]	1 [−2, 3]	0.71
Substantive advice	10% [8%, 13%]	15% [12%, 18%]	5 [0.4, 8]	0.03	12% [8%, 16%]	13% [9%, 17%]	1 [−5, 6]	0.76
Quality of Response Approach								
Quality Response	0.32 [0.29, 0.34]	0.36 [0.34, 0.39]	0.04 [0.004, 0.09]	0.03	0.30 [0.26, 0.35]	0.38 [0.34, 0.42]	0.08 [0.02, 0.14]	0.01
Length of Response Approach								
Word count	33 [28, 38]	43 [36, 50]	10 [2, 18]	0.02	42 [31, 53]	53 [39, 67]	11 [−5, 27]	0.19
Log word count	0.77 [0.70, 0.84]	0.89 [0.82, 0.96]	0.12 [0.01, 0.22]	0.03	0.75 [0.63, 0.86]	0.94 [0.83, 1.05]	0.19 [0.04, 0.35]	0.02
Character count	148 [123, 173]	194 [162, 226]	46 [8, 84]	0.02	193 [140, 246]	235 [170, 300]	42 [−32, 117]	0.25

Note: The first two columns of Table 9 evaluate the content of the email responses from male legislators. The third column indicates the difference in the content of the responses, with positive differences indicating a pro-women bias. 95% confidence intervals are shown in square brackets. The fourth column indicates whether the differences are statistically significant. The fifth and sixth columns evaluate the content of the email responses from female legislators. The seventh column indicates the difference in the content of the responses, with positive differences indicating a pro-women bias. The 95% confidence intervals are shown in square brackets. The eighth column indicates whether the differences are statistically significant.

References

Coppock, Alexander. 2019. "Avoiding post-treatment bias in audit experiments." *Journal of Experimental Political Science* 6(1):1–4.

Costa, Mia. 2020. "Citizen evaluations of legislator-constituent communication." *British Journal of Political Science.* Published online January 30. https://doi.org/10.1017/S0007123419000553.

Kalla, Joshua L., Frances Rosenbluth and Dawn Teele. 2018. "Are you my mentor? A field experiment on gender, ethnicity, and political self-starters." *Journal of Politics* 80(1):337–41.

Loewen, Peter John and Michael Kenneth MacKenzie. 2019. "Service representation in a federal system. A field experiment." *Journal of Experimental Political Science* 6(2):93–107.