How Responsive are Political Elites? A Meta-Analysis of Experiments on Public Officials*

Mia Costa

Abstract

In the past decade, the body of research using experimental approaches to investigate the responsiveness of elected officials has grown exponentially. Given this explosion of work, a systematic assessment of these studies is needed not only to take stock of what we have learned so far about democratic responsiveness, but also to inform the design of future studies. In this article, I conduct the first meta-analysis of all experiments that examine elite responsiveness to constituent communication. I find that racial/ethnic minorities and messages sent to elected officials (as opposed to non-elected) are significantly less likely to receive a response. A qualitative review of the literature further suggests that some of these inequalities in responsiveness are driven by personal biases of public officials, rather than strategic, electoral considerations. The findings of this study provide important qualifications and context to prominent individual studies in the field.

Keywords: Meta-analysis, elite responsiveness, representation

INTRODUCTION

In the past decade, a burgeoning body of research has used experimental approaches to shed new light on how responsive public officials are to constituents. Specifically, over 50 audit experiments on elite responsiveness have been conducted since Butler and Broockman's (2011) initial study of constituent communication.[1] But the fact that any single study provides limited information has become especially poignant given recent publicized concerns about publication bias and

[1]Some of these experiments are excluded from this meta-analysis because they do not fit all the inclusion criteria. See Section 2.

minimal replication in scientific research (Collaboration, 2015; Singal, 2015). Given the explosion of work in this area and the novelty of methods being used, a systematic assessment of these studies is needed not only to take stock of what we have learned so far about democratic responsiveness, but also to inform the design of future studies taking on this question. Meta-analysis is one technique that can overcome the limitations of standard null hypothesis significance testing (Gill, 1999), as well as provide a comprehensive framework for understanding a heavily studied or burgeoning topic of research (Humphrey, 2011).

In this article, I conduct the first meta-analysis of all published and unpublished experiments that use political elites as subjects and where responsiveness to constituent communication is the outcome of interest. I estimate the overall rate at which government officials respond to constituent communication, as well as how this effect varies across different experimental designs. I am particularly concerned with uncovering whether the constituent's race, level of office being studied, and content of the message condition how political elites respond to constituent communication. These factors have been central to analyses of elite responsiveness.

Levels of responsiveness vary by study. Although Butler and Broockman (2011) found that officials were 5.1% more likely to respond to white constituents over black constituents, Einstein and Glick (2017) found they were 3.2% *more* likely to respond to blacks. When it comes to Latino/a constituents, we are even more unsure of the effect of racial or ethnic discrimination on responsiveness. Mendez (2014), for example, finds that state legislators respond to Latino constituents at a rate of 29.8%, while Mendez and Grose (2014) find that same response rate to be as high as 40.3%.

What explains the differences between findings? It is possible that methodological choices are responsible for the different levels of responsiveness. For example, Butler and Broockman (2011) email state legislators, while Einstein and Glick (2017) focus on lower-level public officials. Even compared to other studies with similar, lower-level officials, responsiveness levels vary (e.g. Butler and Crabtree, 2016; White et al., 2015). What is the effect of the level of office, or perhaps being elected versus non-elected, on elite responsiveness? Additionally, Butler and Broockman (2011) waited 30 days for responses, White et al. (2015) waited 62 days, and Mendez and Grose (2014) waited 14 days. Meanwhile, other studies waited as long as 10 months (Butler et al., 2012). Can differences in response rates be attributed to different response cut-off times?

The questions outlined above highlight why a meta-analytic review of the literature is important. Responsiveness is a multi-faceted conception that manifests in elite communication with constituents in various ways. Because much of the research in this area is relatively recent and pioneering in its own right, considerable diversity exists among the findings. I consider the methods used and design elements of each experiment to synthesize these findings and explain variations in elite responsiveness to constituent communication.

METHOD

For a study to be included in this analysis, it had to be a fully randomized controlled trial and use political elites as subjects and responsiveness to constituent communication as the outcome variable.[2] I searched for every published and unpublished study that might fall within this area. I used a wide variety of search tactics to be as comprehensive as possible, including searching library/journal databases, conference proceedings, pre-registration databases, and sending out calls to personal contacts and email listserves.[3]

This search resulted in a final dataset of 28 published and 13 unpublished experiments that could be located as of February 23, 2016, for a total of 41 experiments from 19 different papers and 1 book.[4] Table 1 includes the full list of papers from which the experiments were collected. The papers were published or written between 2011 and 2016 with the earliest experiment conducted in 2007 (Grose et al., 2015), highlighting the recent and rapid emergence of these kind of experiments in the discipline. Thirty-three studies focused on the U.S., 3 on China, 3 on Germany, 1 on South Africa, and 1 on the European Union.

Each separate experiment, even if appearing in the same article, is included as its own case in the meta-analysis. For example, if requests sent to Congress members and state legislators are estimated separately in the same paper, I include both as separate observations. This is standard practice in meta-analyses, but to ensure the results are not conditioned by the articles they appear in, I also analyze the studies without "double-counting" experiments from the same article in the SI.

The results from each study are coded and transformed into a common metric so I can examine the consistency and magnitude of findings across all studies. To calculate the main effect of constituent communication on elite responsiveness, I record the proportion of constituent requests that received a response and, if available, meaningful response.[5] After the proportion of responses are coded or calculated, they are aggregated to produce a summary estimate of the overall effect of constituent communication on elite responsiveness. Since the studies vary widely

[2]Too much heterogeneity across these terms introduces uncertainty among the effects. Carefully defining each of these concepts (experiment, responsiveness, elites, communication) across a common metric is therefore necessary in order to ensure there is some level of homogeneity across the studies. See the Supporting Information (SI) for more detailed information on how the criteria were defined.
[3]See the SI for more information on this process.
[4]This is not including one additional study on Swedish officials that I exclude because, although relevant to the present analysis, Swedish officials are required by national law not only to respond to every public request, but are also explicitly required to respond equally to all requests and to offer email as a means of communication, resulting in a response rate of almost 100% (Adman and Jansson, 2017).
[5]When this information was not reported, I used online supplementary/replication material whenever available, or directly inquired with the authors, to manually calculate the number of responses received out of the total requests sent.

<div align="center">

Table 1

Papers/Books Included in Meta-Analysis, Listed Alphabetically by Last Name

</div>

Author(s)	Title	Journal	Year	Country
Benjamin Bishin, Thomas Hayes	Do elected officials service the poor? A field experiment on the U.S. congress	Unpublished	2016	US
Damien Bol, Thomas Gschwend, Thomas Zittel, Steffen Zittlau	The electoral sources of good political representation: A field experiment on German MPs	Unpublished	2015	Germany
David Broockman	Black politicians are more intrinsically motivated to advance blacks interests: A field experiment manipulating political incentives	AJPS	2013	US
Daniel Butler, David Broockman	Do politicians racially discriminate against constituents? A field experiment on State legislators	AJPS	2011	US
Daniel Butler	Representing the advantaged: How politicians reinforce inequality	Cambridge University Press	2014	US
Daniel Butler, Charles Crabtree	Moving beyond measurement: Adapting audit Studies to test bias-reducing interventions	Unpublished	2016	US
Daniel Butler, Chris Karpowitz, Jeremy Pope	A field experiment on legislators home styles: service versus policy	JOP	2012	US
Nicholas Carnes, John Holbein	Unequal responsiveness in constituent services? Evidence from casework request experiments in North Carolina	Unpublished	2015	US
Catherine De Vries, Elias Dinas, Hector Solaz	You have got mail! How intrinsic and extrinsic motivations shape legislator responsiveness in the european parliament	Unpublished	2015	EU
Greg Distelhorst, and Yue Hou	Ingroup bias in official behavior: A national field experiment in China	QJPS	2014	China
Kyle Dropp, Zachary Peskowitz	Electoral security and the provision of constituency service	JOP	2012	US
Katherine Levine Einstein, David Glick	Does race affect access to government services? An experiment exploring street level bureaucrats and access to public housing	AJPS	2017	US
Stephan Grohs, Christian Adam, Christoph Knill	Are some citizens more equal than others? Evidence from a field experiment	PAR	2015	Germany

<div align="center">

Table 1
(Continued)

</div>

Author(s)	Title	Journal	Year	Country
Christian Grose, Neil Malhotra, Robert Van Houweling	Explaining explanations: How legislators explain their policy positions and how citizens react	AJPS	2015	US
Andrew Janusz, Nazita Lajevardi	Differential responsiveness: Do legislators discriminate against hispanics?	Unpublished	2016	US
Gwyneth McClendon	Race and responsiveness: A field experiment with South Africa in politicians	JEPS	2016	South Africa
Matthew Mendez	Who represents the interests of undocumented latinos? A field experiment of state legislators	Unpublished	2014	US
Matthew Mendez, Christian Grose	Revealing discriminatory intent: Legislator preferences, voter identification, and responsiveness bias	Unpublished	2014	US
Tianguang Meng, Jennifer Pan, Ping Yang	Conditional receptivity to citizen participation: Evidence from a survey experiment in China	CPS	2014	China
Ariel White, Noah Nathan, Julie Faller	What do I need to vote? Bureaucratic discretion and discrimination by local election officials	APSR	2015	US

Note. AJPS=American Journal of Political Science, JOP=Journal of Politics, CPS=Comparative Political Studies, QJPS=Quarterly Journal of Political Science, APSR=American Political Science Review, JEPS=Journal of Experimental Political Science, PAR=Public Administration Review.

in sample size, the findings are weighted using the inverse variance of each study and between-studies τ^2, standard in random-effects models for meta-analysis.[6]

Finally, since the overall response rate is not generally the estimand of interest in this literature, the studies are tested for moderator effects, or variables that might influence the outcome of an experiment. I examine whether (1) the length of time researchers waited for a response in days, (2) the sender was a minority or not, (3) the request focused on constituency service or a policy issue, (4) the level of government (national or sub-national), (5) the elite is elected or not, and (6) the experiment was conducted in the U.S. or not, explains the variation in response rates across experiments. See Table 2 for summary statistics on these variables. The SI includes more details on how each of these variables was coded.

The types of moderators I can include are limited to those that are central to the design of the experiments. In other words, I am only able to incorporate factors that are consistently reported by the studies (i.e. the type of elite subject) or are

---

[6]See the SI for more detail on the method, as well as robustness tests of other weighting adjustments.

Table 2
Descriptive Statistics of Moderator Variables

	Freq. (#)	Percent (%)
Minority constituent	14	34.1
Service request	34	83.0
Sub-national	31	75.6
Elected official	34	83.0
In U.S.	33	80.5
	Mean (s.d.)	Median
Response cutoff (days)	77.9 (90.9)	35

part of the experimental designs themselves (i.e. the race of the constituent sender). Some additional design manipulations that exist in the literature, such as whether the constituent and public official share the same race or political party, have not yet been included in enough experiments to be meta-analyzed, but I do qualitatively take those into account when interpreting the results.

RESULTS

How Responsive are Political Elites Overall?

To answer how responsive political elites are to constituent requests, I fit a random-effects model using the *metafor* package in R (Viechtbauer, 2010) to compute a weighted mean of the effect sizes. For more details on random-effects models and additional robustness checks of the weighting technique used, see the SI.

Figure 1 is a forest plot that displays the results from each study. The columns on the left indicate whether the study is published and in what country the experiment took place. The last column indicates the observed response rate with 95% confidence intervals. In the middle, the observed effects are displayed by square boxes that are sized proportional to the precision of the estimates. At the bottom, the combined effect is represented by a diamond with the outer edges drawn to the confidence interval limits.

The average observed effect of the treatment is 0.529. That is, political elites respond to constituent communication 53% of the time. The main effect for only published studies is 0.542 and for unpublished studies is 0.50. Neither of these are statistically distinguishable than the combined effect for all studies, but see the SI for additional tests for publication bias. The response rates range from 0.19 to 0.79. The lowest response rate occurred in an experiment by Butler et al. (2012) conducted on U.S. state legislators. The two highest response rates were in Grohs et al. (2015) in Germany.

Some studies additionally measure the *quality* of the response received, i.e. whether it is "friendly" or "helpful." Specifically, 19 out of the 41 experiments
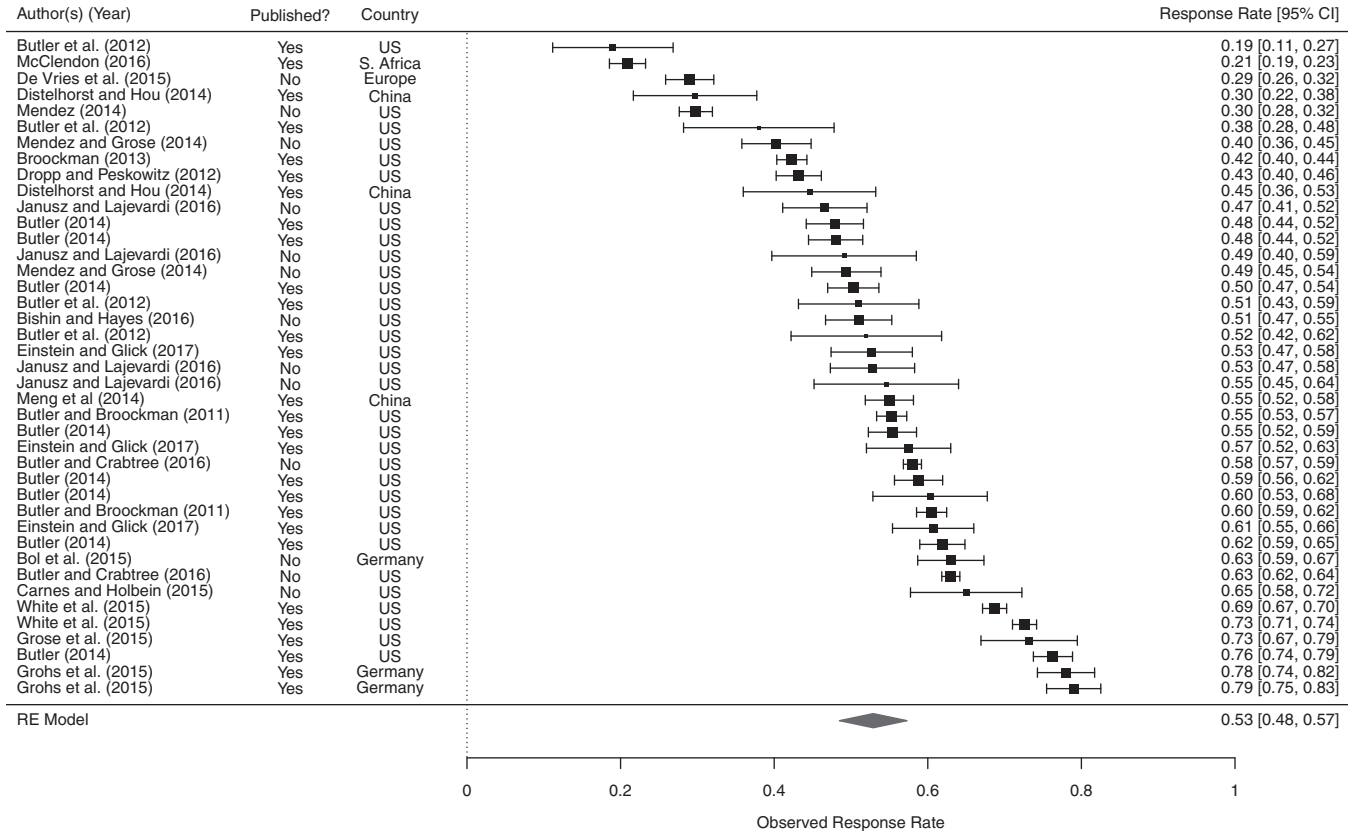
Author(s) (Year)	Published?	Country	Response Rate [95% CI]
Butler et al. (2012)	Yes	US	0.19 [0.11, 0.27]
McClendon (2016)	Yes	S. Africa	0.21 [0.19, 0.23]
De Vries et al. (2015)	No	Europe	0.29 [0.26, 0.32]
Distelhorst and Hou (2014)	Yes	China	0.30 [0.22, 0.38]
Mendez (2014)	No	US	0.30 [0.28, 0.32]
Butler et al. (2012)	Yes	US	0.38 [0.28, 0.48]
Mendez and Grose (2014)	No	US	0.40 [0.36, 0.45]
Broockman (2013)	Yes	US	0.42 [0.40, 0.44]
Dropp and Peskowitz (2012)	Yes	US	0.43 [0.40, 0.46]
Distelhorst and Hou (2014)	Yes	China	0.45 [0.36, 0.53]
Janusz and Lajevardi (2016)	No	US	0.47 [0.41, 0.52]
Butler (2014)	Yes	US	0.48 [0.44, 0.52]
Butler (2014)	Yes	US	0.48 [0.44, 0.52]
Janusz and Lajevardi (2016)	No	US	0.49 [0.40, 0.59]
Mendez and Grose (2014)	No	US	0.49 [0.45, 0.54]
Butler (2014)	Yes	US	0.50 [0.47, 0.54]
Butler et al. (2012)	Yes	US	0.51 [0.43, 0.59]
Bishin and Hayes (2016)	No	US	0.51 [0.47, 0.55]
Butler et al. (2012)	Yes	US	0.52 [0.42, 0.62]
Einstein and Glick (2017)	Yes	US	0.53 [0.47, 0.58]
Janusz and Lajevardi (2016)	No	US	0.53 [0.47, 0.58]
Janusz and Lajevardi (2016)	No	US	0.55 [0.45, 0.64]
Meng et al (2014)	Yes	China	0.55 [0.52, 0.58]
Butler and Broockman (2011)	Yes	US	0.55 [0.53, 0.57]
Butler (2014)	Yes	US	0.55 [0.52, 0.59]
Einstein and Glick (2017)	Yes	US	0.57 [0.52, 0.63]
Butler and Crabtree (2016)	No	US	0.58 [0.57, 0.59]
Butler (2014)	Yes	US	0.59 [0.56, 0.62]
Butler (2014)	Yes	US	0.60 [0.53, 0.68]
Butler and Broockman (2011)	Yes	US	0.60 [0.59, 0.62]
Einstein and Glick (2017)	Yes	US	0.61 [0.55, 0.66]
Butler (2014)	Yes	US	0.62 [0.59, 0.65]
Bol et al. (2015)	No	Germany	0.63 [0.59, 0.67]
Butler and Crabtree (2016)	No	US	0.63 [0.62, 0.64]
Carnes and Holbein (2015)	No	US	0.65 [0.58, 0.72]
White et al. (2015)	Yes	US	0.69 [0.67, 0.70]
White et al. (2015)	Yes	US	0.73 [0.71, 0.74]
Grose et al. (2015)	Yes	US	0.73 [0.67, 0.79]
Butler (2014)	Yes	US	0.76 [0.74, 0.79]
Grohs et al. (2015)	Yes	Germany	0.78 [0.74, 0.82]
Grohs et al. (2015)	Yes	Germany	0.79 [0.75, 0.83]
RE Model			0.53 [0.48, 0.57]

Observed Response Rate

Figure 1
Forest Plot of All Studies on Elite Responsiveness.

Figure Plots the Estimate and 95% Confidence Interval for Each Study. Estimates are Represented by the Black Boxes and Sized Proportional to Their Precision. Studies with Larger Boxes are Given More Weight in the Calculation of the Effect Size.

Table 3
Meta-Regression Analysis Estimating the Effect of Moderators on Elite Responsiveness

	(1)	(2)
Response cutoff	− 0.005	0.055
(log days)	(0.031)	(0.050)
Minority constituent	− 0.094*	− 0.046*
(baseline=Non-minority)	(0.047)	(0.006)
Service	0.003	0.250*
(baseline=Policy)	(0.071)	(0.044)
Sub-national	− 0.061	− 0.100*
(baseline=National)	(0.055)	(0.045)
Elected	− 0.181*	− 0.326*
	(0.055)	(0.130)
In U.S.	0.054	− 0.091
	(0.054)	(0.135)
Intercept	0.762*	0.509*
	(0.144)	(0.187)
Observations	31	31
Fixed effects?	No	Yes
R^2	0.489	0.999
τ^2	0.011 (0.003)	-

Note. *$p < 0.05$. Standard errors in parentheses. τ^2 represents the amount of heterogeneity among the true effects that is not already accounted for by the moderators. τ^2 estimator: Restricted maximum-likelihood estimation. Response cutoff is not reported for 10 of the studies. See the SI for a model without this variable to preserve the full *N*.

measured the quality of responsiveness. Using this alternative outcome variable in a random-effects model, the main effect of constituent communication on receiving a "good" response from a public official is 0.453, a difference of 8 points from receiving *any* response (0.53).[7]

Why Responsiveness Varies

Although the response rate from this set of studies was consistent across different robustness checks, it does not explain what causes some studies to find higher response rates than others. And of course, estimating the overall rate of response from public officials to constituent communication is not the main aim of these studies, but rather to examine if officials are more responsive to some constituents than others. Toward this end, I estimate a mixed-effects model, which is simply a random-effects model with covariates that may account for differences across studies. The first column in Table 3 presents the results from a model with these variables included.[8]

[7]See the SI for that full analysis.
[8]I also estimate models that include the year the experiment was conducted or the year the paper was written to control for the possibility that experimenters are getting better over time at soliciting responses. Neither of these variables are statistically significant when included, nor do they alter the

The first thing to note is that design-oriented variables that are typically thought to affect how likely researchers will observe a high response rate do not seem to have statistically significant effects. For example, studies with a later response cutoff do not necessarily record higher responsiveness. That is, waiting longer for responses does necessarily yield a higher number of responses. Additionally, although policy related communications are largely avoided by many researchers in lieu of service requests, possibly due to their dampening effect on responsiveness originally found in Butler et al. (2012), no such statistically significant effect between service and policy communication is discernible here. The assumption that service requests are much more likely to receive a response is often built into these experiments at the design stage. Butler et al. (2012) propose multiple theories for why elected officials should be less likely to respond to policy-oriented messages—offending the constituent if they disagree on issues, for example—and the benefits of service requests—such as defending local interests and cultivating an image of helpfulness. Yet, service communication is not necessarily prioritized by public officials over policy communication when estimating responsiveness across all studies. Not only is the coefficient not statistically significant for this variable, but the estimated difference is substantively very small.

As for variables that do have a statistically significant effect, minority constituents are almost 10 percentage points less likely to receive a response than non-minority constituents ($p < 0.05$). This is consistent with many individual studies that have shown requests from racial and ethnic minorities are given less attention overall, and particularly when the recipient official does not share their race (Broockman, 2013; Butler and Broockman, 2011; Distelhorst and Hou, 2014). It is possible, however, that these results would vary based on the specific racial or ethnic group of the sender. Latinos, for example, do not turn out to vote at as high of a rate as African Americans in the United States, so there may be less electoral incentive for political elites to respond to communication from Latinos. Although I use the minority/non-minority dichotomy because of the limited number of studies that focus on racial/ethnic biases in response rates, when I do include separate indicators for Blacks and Latinos in the U.S., Latinos are 14.2 percentage points less likely to receive a response than white constituents, compared to a 7.3 percentage point deficit for Blacks.

Moreover, elected political elites significantly decrease the observed response rate by 18% ($p < 0.05$) as opposed to non-elected elites, such as bureaucratic officials. It is possible that elected officials are more bombarded with constituent emails on a daily basis and are therefore unable to answer every request. Some scholars have also found that elected officials are more likely to respond to constituents in their

---

coefficients of the other model variables. To preserve the degrees of freedom I exclude them from the final model presented here. Also note that 10 observations are lost because the response cutoff is not available in all studies. However, when I exclude that variable and run the model with the full set of observations, there are no statistically or substantively significant changes in the effects that are reported here. See the SI for that full analysis.

own representational jurisdiction (e.g. Broockman, 2013), and that they tailor their responses to suit the constituent's policy preferences (Grose et al., 2015). Public officials who are not directly elected theoretically do not have to make the same considerations.

Of course, given the relatively small number of studies in this analysis, it may be difficult to estimate these moderator effects with high levels of precision. Thus, just because a moderator is not statistically significant in the model does not mean that it does not matter for affecting response rates. The second column in Table 3 reports the same analysis but with fixed effects included for each study, since some of these variables—such as response cutoff, service/policy, elected/not elected, national/sub-national— are typically not randomized within the same study; that is, the variation occurs *across* studies. Generally speaking, the results are mostly consistent in this model. The coefficients for service requests and sub-national public officials are now statistically significant, although in the same direction as in the first model. Service requests are more likely to receive a response, whereas requests to sub-national officials yield lower response rates. The coefficient for minority constituent is about half the size as it is in the first model, but it remains statistically significant. Finally, the fixed effects model produces a coefficient for the elected official variable that is larger than in the first model. Overall, although the magnitude of some of the effects estimated in the first model differ when using study-level fixed effects, these differences generally do not substantially alter the overall patterns uncovered in model 1.[9]

Strategic Considerations or Personal Bias?

The analysis above uncovered that elite responsiveness is not equal across all conditions. Yet the mechanisms behind these inequalities are not as clear. For example, contextual factors, such as district competitiveness, could affect how often legislators respond to constituent communication. Since it is rare for any two studies to measure these contextual variables in the same way (if they are accounted for at all), I am unable to meta-analyze their effects on response rates. However, a qualitative overview of these theoretically-relevant, contextual variables can shed light on the meaning of the average effects in the analysis above.

Although it is difficult to causally identify the impact of contextual factors on responsiveness since they are not randomly assigned, some scholars have examined whether treatment effects vary across contexts. In studies that examine the effect of race and ethnicity on responsiveness, the main theoretical question of interest is whether strategic, electoral considerations cause officials to be less responsive to

---

[9]One limitation of this analysis is that some moderators might not translate across countries. Although understanding the effects of the experiments in each country would be most ideal, there are not enough studies in each country to conduct a sub-analysis for each or include country fixed effects. I therefore test the effect of moderators only on experiments conducted in the United States. See the SI for results for models that only focus on cases in the United States.

minorities, rather than personal, intrinsic bias. For example, several studies have examined whether legislators have a heightened electoral incentive to respond to minorities when they comprise a larger share of the district's population. Although this is measured in different ways across the four studies that account for this factor (proportion of African Americans or Latinos in the district, "high"-minority populations versus "low"-minority populations, etc.), three of the four studies did not find a statistically significant relationship between levels of minorities in the population and responsiveness (Einstein and Glick, 2017; Mendez and Grose, 2014; Janusz and Lajevardi, 2016). White et al. (2015) did uncover a mild difference but with one important caveat: Although local election officials were more biased against Latinos in low-Latino localities, this was highly correlated with Voting Rights Act (VRA) coverage. Since high-Latino localities are also covered by the VRA, it is unclear whether the lack of VRA coverage or lower electoral incentive to respond to Latinos is driving the effect.

Among the two studies that empirically considered the effects of VRA coverage (Butler, 2014; White et al., 2015), the results are mixed. White et al. (2015) find that bias against Latinos is 7.5% points lower in places covered by the VRA. Yet Butler (2014) find that legislators from states covered under the VRA exhibited the same level of bias as legislators in states not covered under the VRA. It is therefore unclear whether anti-discrimination laws are effective at reducing bias in responsiveness to minorities.

It is also possible that legislators are more likely to respond to all constituents, minorities or not, if they are up for re-election or the district is particularly competitive. Yet none of the studies testing for this possibility found that this was the case (Butler and Broockman, 2011; Butler, 2014; Janusz and Lajevardi, 2016; Mendez and Grose, 2014).

Finally, a few studies examined whether responsiveness correlate with total population size (Janusz and Lajevardi, 2016; Mendez and Grose, 2014; White et al., 2015). None of these studies found a statistically significant relationship between the total population in a legislator's district and the extent to which that legislator discriminated against non-whites.

The one factor that did seem to have a consistent effect on mediating racial biases is not a contextual characteristic about the localities, but a personal characteristic of the public official. Specifically, whether an official was the same race as the constituent influenced their propensity to respond in multiple studies (Broockman, 2013; Butler, 2014; Butler and Broockman, 2011; Mendez, 2014). This minority ingroup/outgroup effect goes above and beyond electoral incentives. Although all legislators are more likely to respond to fellow partisans, bias against racial outgroups remains among both white and minority legislators (Butler, 2014). Broockman (2013), for example, found that African American legislators were much more likely than other legislators to respond to African Americans even when they purportedly lived outside their district. This finding suggests that

African American legislators are more "intrinsically motivated" to respond to constituents of the same race, regardless of the electoral incentive to do so. Increasing the electoral incentive for white legislators to respond to minorities also does not close the gap in their response rates to whites versus minorities. As Butler (2014) concludes, "while there is evidence that strategic considerations regarding voters' perceived partisanship might partially motivate the observed patterns of discrimination, there remain significant levels of discrimination that cannot be explained by strategic responses alone" (108).

DISCUSSION

This analysis provides a number of qualifications to many prominent individual findings in the literature. For example, the difference between service requests and policy requests was not statistically significant and the magnitude of the coefficient was quite small. This is notable considering the vast majority of studies use service, instead of policy, as the focus of the communication, reflecting the widespread assumption that service requests are more likely to receive a response (a result originally found in Butler et al. (2012)).

Moreover, although individual studies provide a range of different estimates for the effect of race on response rates, this meta-analysis finds a substantively large and statistically significant 10 percentage point decrease in responsiveness to racial and ethnic minority constituents. When I analyzed Latino and African American constituents separately, Latinos were significantly less likely than whites to receive a response from government officials compared to African Americans. This is a more precise estimate that offers clarification to the relatively large body of research on this topic.

To be sure, there are a multitude of other factors that could affect responsiveness that have not yet been studied via this experimental approach. For example, elected officials might be much more likely to respond when there are electoral incentives to do so. A qualitative review of this literature suggested that while contextual variables such as anti-discrimination laws and strategic considerations may help to close the deficit in responsiveness to racial and ethnic minorities, there remain personal, intrinsic biases that remain unaccounted for. These findings not only provide an overall sense of how responsive officials are to constituent requests, but also help to organize, frame, and understand the literature on elite responsiveness.

SUPPLEMENTARY MATERIALS

To view supplementary material for this article, please visit https://doi.org/10.1017/XPS.2017.14.

REFERENCES

Adman, Per and Hanna Jansson. 2017. "A Field Experiment on Ethnic Discrimination Among Local Swedish Public Officials." *Local Government Studies* 44(1): 44–63.

Bishin, Benjamin and Thomas Hayes. 2016. "Do Elected Officials Service the Poor? A Field Experiment on the U.S. Congress." Paper presented at the annual Southern Political Science Associate conference, San Juan, Puerto Rico (January 9, 2016). https://www.dropbox.com/s/ahhctobf0gb1vbc/Do%20Elected%20Officials%20Service%20the%20Poor%3F%20%20A%20Field%20Experiment%20on%20Congress..pdf?dl=0

Bol, Damien, Thomas Gschwend, Thomas Zittel and Steffen Zittlau. 2015. "The Electoral Sources of Good Political Representation A Field Experiment on German MPs." Paper presented at the Annual Meeting of the European Political Science Association, Vienna (June 25–27, 2015). https://www.dropbox.com/s/vw7682r0qwto85b/The%20Electoral%20Sources%20o%f%20Good%20Political%20Representation-A%20Field%20Experiment%20on%20German%20M%Ps.pdf?dl=0.

Broockman, David E. 2013. "Black Politicians are More Intrinsically Motivated to Advance Blacks' Interests: A Field Experiment Manipulating Political Incentives." *American Journal of Political Science* 57(3): 521–536.

Butler, Daniel M. 2014. *Representing the Advantaged: How Politicians Reinforce Inequality*. New York, NY: Cambridge University Press.

Butler, Daniel M. and Charles Crabtree. 2016. "Moving Beyond Measurement: Adapting Audit Studies to Test Bias-Reducing Interventions." Working Paper. https://www.dropbox.com/s/p8brs5dr5xmgwdv/Moving%20Beyong%20Measurement.pdf?dl=0.

Butler, Daniel M., Christopher F. Karpowitz and Jeremy C. Pope. 2012. "A Field Experiment on Legislators' Home Styles: Service versus Policy." *The Journal of Politics* 74(02): 474–486.

Butler, Daniel M. and David E. Broockman. 2011. "Do Politicians Racially Discriminate Against Constituents? A Field Experiment on State Legislators." *American Journal of Political Science* 55(3): 463–477.

Carnes, Nicholas and John Holbein. 2015. "Unequal Responsiveness in Constituent Services? Evidence from Casework Request Experiments in North Carolina." Working Paper. http://people.duke.edu/~nwc8/carnes_and_holbein.pdf.

Collaboration, Open Science. 2015. "Estimating the Reproducibility of Psychological Science." *Science* 349(6251): aac4716.

Costa, Mia. 2017. "Replication Data for: How Responsive are Political Elites? A Meta-Analysis of Experiments on Public Officials." Harvard Dataverse, doi: 10.7910/DVN/0HDTYM.

De Vries, Catherine, Elias Dinas, and Hector Solaz. 2015. "You Have Got Mail! How Intrinsic and Extrinsic Motivations Shape Legislator Responsiveness in the European Parliament." Paper presented at the annual Southern Political Science Association conference, New Orleans, LA (January 15–17, 2015). http://catherinedevries.eu/EPResposiveness_Feb2015.pdf.

Distelhorst, Greg and Yue Hou. 2014. "Ingroup Bias in Official Behavior: A National Field Experiment in China." *Quarterly Journal of Political Science* 9(2): 203–230.

Dropp, Kyle and Zachary Peskowitz. 2012. "Electoral Security and the Provision of Constituency Service." *The Journal of Politics* 74(01): 220–234.

Einstein, Katherine Levine and David M. Glick. 2017. "Does Race Affect Access to Government Services?: An Experiment Exploring Street Level Bureaucrats and Access to Public Housing." *American Journal of Political Science* 61(1): 100–116.

Gill, Jeff. 1999. "The Insignificance of Null Hypothesis Significance Testing." *Political Research Quarterly* 52(3): 647–674.

Grohs, Stephan, Christian Adam, and Christoph Knill. 2015. "Are Some Citizens More Equal than Others? Evidence from a Field Experiment." *Public Administration Review* 76(1): 155–164.

Grose, Christian R., Neil Malhotra and Robert P. Van Houweling. 2015. "Explaining Explanations: How Legislators Explain Their Policy Positions and How Citizens React." *American Journal of Political Science* 59(3): 724–743.

Humphrey, Stephen E. 2011. "What Does a Great Meta-Analysis Look Like." *Organizational Psychology Review* 1(2): 99–103.

Janusz, Andrew and Nazita Lajevardi. 2016. "Differential Responsiveness: Do Legislators Discriminate Against Hispanics?" Paper presented at the annual Midwest Political Science conference, Chicago, Il (April 5, 2014). updated version: http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2799043.

McClendon, Gwyneth H. 2016. "Race and Responsiveness: An Experiment with South African Politicians." *Journal of Experimental Political Science* 3(1): 60–74.

Mendez, Matthew. 2014. "Who Represents the Interests of Undocumented Latinos? A Field Experiment of State Legislators." Working Paper, University of Southern California. http://ssrn.com/abstract=2592754.

Mendez, Matthew and Christian Grose. 2014. "Revealing Discriminatory Intent: Legislator Preferences, Voter Identification, and Responsiveness Bias." USC CLASS Research Paper No. 14-17. http://ssrn.com/abstract=2422596.

Meng, Tianguang, Jennifer Pan, and Ping Yang. 2014. "Conditional Receptivity to Citizen Participation Evidence From a Survey Experiment in China." *Comparative Political Studies* 0010414014556212.

Singal, Jesse. 2015. "The Case of the Amazing Gay-Marriage Data: How a Graduate Student Reluctantly Uncovered a Huge Scientific Fraud." *NYMAG.com*.

Viechtbauer, Wolfgang. 2010. "Metafor: Meta-Analysis Package for R." *R package version* 2010: 1–0.

White, Ariel R., Noah L. Nathan, and Julie K. Faller. 2015. "What Do I Need to Vote? Bureaucratic Discretion and Discrimination by Local Election Officials." *American Political Science Review* 109(01): 129–142.