

Why Public Employees Rebel: Guerrilla Government in the Public Sector

Gary E. Hollibaugh, Jr.¹
University of Notre Dame

Matthew R. Miles
Brigham Young University Idaho

Chad B. Newwander
Brigham Young University Idaho

Justin S. Vaughn
Boise State University

¹ Authors names are alphabetical. Each author had an equal contribution to this paper.

Why Public Employees Rebel: Guerrilla Government in the Public Sector

In late January of 2017, approximately one thousand foreign service officers wrote to “register [their] dissent to the State Department’s implementation” of President Donald’s Trump executive order that temporally banned immigrants from seven predominately Muslim countries.² Kathryn Sullivan, chief of the National Oceanic and Atmospheric Administration, registered her dissent to the House Committee on Science, Space, and Technology in December of 2015 after it subpoenaed internal research documents on climate change. She would not comply with the request.³ Kim Davis registered her dissent of the decision made by the Supreme Court to legalize same-sex marriage by refusing to sign marriage licenses for same-sex couples, temporarily ending up in jail as a result. Dan Quinn and Charles Martland, both Green Berets, registered their dissent to instructions from their commanding officers about to ignore child abuse from local allies in Afghanistan.⁴ Quinn was removed from his command and Martland faced disciplinary proceedings. Agent John Dodson of the Alcohol, Tobacco and Firearms registered his dissent by blowing the whistle in exposing an administrative practice that allowed guns, which were bought in the United States, to be taken into Mexico.⁵ A Department of Justice lawyer then attempted to tarnish Dodson’s reputation.

The registering of dissent and the ensuing reprisals is a behavioral reality within the confines of public service and public organizations. What Rosemary O’Leary (2014) refers to as guerrilla government, or “the actions taken by career public servants who work against the

² See the order at https://www.washingtonpost.com/r/2010-2019/WashingtonPost/2017/01/30/Editorial-Opinion/Graphics/Draft_Dissent_on_EOsec3.pdf.

³ Rein, Lisa. “She’s braved rough seas and space walks. Can she weather climate change skeptics?” *Washington Post*, December 15, 2015.

⁴ Goldstein, Joseph. “U.S. Soldiers Told to Ignore Sexual Abuse of Boys by Afghan Allies.” *The New York Times*, September 20, 2015.

⁵ Atkisson, Sharyl. “Agent: I was ordered to let U.S. guns into Mexico.” *CBS News*, March 3, 2011.

wishes—either implicitly or explicitly communicated—of their superiors,” is part and parcel of administration (xi). Even though such resistance should not come as a surprise, it does generate an expected level of concern. So-called deep state actors are accused of undermining democratic norms when they oppose and work against orders from their constitutional superiors, agency heads, and immediate supervisors.

Despite the gravity and/or grimness of following an ethic of dissent in public service, we do not fully understand the behavioral dynamics that push administrators away from compliance and toward dissent. In fact, there is “little empirical public administration research [that] has been devoted to unbureaucratic behavior,” such as guerrilla dissent aimed at resisting rules and orders from superiors (DeHart-Davis, 2007, p. 901). From a Behavioral Public Administration (BPA) perspective, this is partly due to the fact that the field has not given priority to Herbert Simon’s admonishment to study what administrators actually do. The descriptive and scientific account has often been displaced by the normative and philosophical (Dubnick, 1999). From a politics-administration dichotomy angle, this is partly due to the fact that an ethic of neutrality has not entirely dissipated (Thompson, 1985). Accepting a legitimate role of administrators to help govern the ship of state is challenged in a recurring manner in both theory and practice (Wamsley and Wolf, 1996). From the guerrilla government angle, this partly due to the fact that such activities are marked by irregularity and breadth (O’Leary, 2016). The variance is significant in how administrators undermine their superiors, ranging from being an actor of omission in allowing their superiors to fail to being an actor of commission in actively working against their superiors’ aims.

Due to these limitations, greater awareness is needed about the behavioral tendencies of administrators who carry out opposition against their superiors. This is important, because a core

aim of O’Leary’s insights into guerrilla government is to provide a deeper layer of knowledge about organizations and management, bureaucratic politics, and ethics. Furthermore, these insights into guerrilla government buffer and bolster another key aim of O’Leary’s research—providing knowledge of how to manage dissent. Without a more concrete understanding of the political and ethical motives that push one toward guerrilla government, it is difficult to ascertain how to manage dissent in a productive manner. Due to this information gap, we are not exactly aware if personal, ethical, or policy preferences push administrators to more likely engage in disruptive behaviors. And, we are not sure if the *likelihood of retribution* scares off administrators, or whether it has any effect at all. We are also not sure if the *type of retribution* hinders guerrilla activities, or even promotes them. The ability to begin parsing out the behavioral tendencies of guerrilla activity provides a contribution to not only understand the ethical conditions under which public servants choose to engage in guerrilla activities and the political tactics they prefer, but also deepens the knowledge base of how to manage dissent.

Thus, to better understand the determinants of guerrilla government, we present the results of a conjoint experiment embedded in a survey of federal bureaucrats. We find that the decisions to obey orders and/or engage in guerrilla government are dependent on many factors, including the personal views of the policy the respondent is asked to enact, the compatibility of the proposed policy with the respondent’s ethical framework, the probability that others will be harmed by the policy, and the status of the person making the request. Notably, we find that the probability and type of expected retribution for disobedience and/or guerrilla government has a substantively insignificant effect on respondents’ decisions. These results provide important context to previous studies of guerrilla government—which we discuss below—as well as studies of bureaucratic politics more generally.

Guerrilla Government

Engaging in guerrilla government activity is serious. The choice to give voice demystifies the neutrally competent civil servant and exposes an agential-directed civil servant who has an agenda (Wamsley, 1990). The consequence of subordinates who willingly use their discretion to disobey a superior generates legitimacy concerns—there is a general level of uneasiness of undemocratic actors pushing back against democratically-elected representatives (Overeem, 2005). But, there is also an uneasiness of not speaking truth to power (Hart 1984; Frederickson and Hart, 1985). In many ways, administrators are in an ambivalent position. They are expected to comply and to be agents. While compliance cannot be dismissed, guerrilla government operates on the assumption that public servants do use their discretion in a way to subvert. To understand the character of such activity, O’Leary frames the exercise of administrative agency, expressed through tactics of dissent, through three interdependent lenses: organizations and management, bureaucratic politics, and ethics.

A core feature of administrative behavior is how often organizations and management fail to secure employee acquiescence, and vice-a-versa, how administrators act against the preferred proclivities of their superiors. The recurring problem of employee non-compliance, disobedience, or even active resistance poses a fundamental challenge to organizational control and coordination. To curb the perceived problem of acting against organizational priorities, Chester Barnard refers to creating a zone of indifference where orders are “unquestionably acceptable” (1968, p. 169), Herbert Simon refines this area of indifference to one of active acceptance, and, Herbert Kaufman polishes this process even further, noting that organizations leverage employee buy-in by establishing “decisions [that] are preformed” and then expecting administrators to carry them out (2006, p. 91).

While Barnard, Simon, and Kaufman sought to understand how administrative behavior could be controlled to ensure predictable and rational results, there is also an awareness among them about the tendencies of individuals to push back. In other words, divergence is expected, especially within a public administration context (Gofen, 2014). The implementation of policy creates discrepancies, or differences in how administration is carried out (Lipsky, 2010). This is due to responding to different needs and conceptions of clients (Maynard-Moody and Musheno, 2003), unclear and numerous goals (Wilson, 1989), and divergent constitutional demands from constitutional superiors (Rohr, 1986). As a result of such conflicts, administrators become political actors who join in the pulling and hauling process of allocating resources and generating meaning (Allison and Zelikow, 1999; Meier and Bohte, 2007; Flyvbjerg, 1998).

While this discretionary power enables administrators to fulfill a political role, it also empowers them to be moral agents who decide how to use their agency in a proper way. For O'Leary, ethics involves more than thinking about what is right and wrong, it involves doing the right thing (p.16). At a general level, administrators may serve in this capacity by internalizing the ethos of the Constitution's checking function by pushing back against abusive and authoritarian actions (Spicer, 1995) and serving as balancing wheel to protect general regime values (Rohr, 1998). Their guerrilla acts can range from grand overtures that try to fundamentally change a governmental entity or policy (Newswander, 2012) or they can be small in scope that guide issues of public importance like finding justice for someone who was abused by public officials (Newswander, 2015). Leading by dissent enables administrators to steer issues of public importance. But, leading by dissent can also steer issues away from public importance and toward parochial and self-interested ends. Guerrilla activity can aim high or low.

At a specific level, administrators can employ a range of guerrilla tactics such as rule bending (DeHart-Davis, 2007; Borry, 2017) and whistle blowing (Near and Miceli, 1996) to do what they perceive is proper. With rule bending, Leisha DeHart-Davis finds that administrators with a commitment to public service are not compelled to act out. One explanation, according to DeHart-Davis, is that the “individuals who indicate commitment to public service have been indoctrinated to the ‘ethic’ of consistent rule application” (2007, p. 900). In contrast to this finding, the association between an ethical mandate and pushing back is challenged to an extent when examining whistle blowing activity. Yoon Jik Cho and Hyun Jin Song (2015) and James Caillier (2017) show that whistle blowing is more likely to occur if civil servants have a public service motivation.

Even with these inconsistent findings, there is a concern that retaliation may dampen guerrilla activities. Janet Near and Marcia Miceli (1996) advance findings that reporting wrong doing in the form of whistle blowing is not impacted by the threat of retaliation. Yet, there is conflicting evidence; they acknowledge that case studies have shown retribution as an indicator of whether someone is likely to blow the whistle. In parsing this out, employees in places of power are less likely to be punished while those in less secure positions are more likely to feel the wrath of their superiors (Miceli and Near, 1992; Near and Miceli, 2008). Cho and Song (2015) look at how federal employees are impacted by this dilemma. They find a negative association between personal costs and whistleblowing. Personal harm may impact the choice to give voice.

While retaliation may or may not impact resistance, there is an expectation that a concern for others or mitigating harms acts as a catalyst to act out to achieve the greater good or protect the vulnerable. In general, Erin Borry (2017) finds that an ethical climate is associated with rule-

bending instead of structural characteristics. Civil servants' failure to be straight arrows is associated with pushing back to protect broader organizational priorities. Furthermore, she also finds rule bending happens to help fellow colleagues, not just achieving organizational priorities. Cecilia Lavena (2016) confirms this general connection between ethics and resistance, but with an important difference—she notes that whistles are blown in order to achieve norms such as respect and fair treatment. Caillier (2017) adds more support to the importance of moral concerns, but with another important twist. The level of harm matters in deciding whether to raise an alarm. The seriousness of wrongdoing and whistle-blowing are positively associated with each other.

Yet, there is still a concern that structure, or organizations and management, explains dissent, not just ethics. Being responsive to one's superior or the rules administrators are expected to follow may contribute to understanding why dissent exists. Civil servants are expected to be compliant to one's political and bureaucratic bosses while also adhering to the prescribed rules and routines. As a result, resistance may be spurred on by the structure and order of power, not morality concerns. DeHart-Davis (2007) discovers that administrators are more likely to push back in the face of centralization while less likely to push back in the face of formalization. In contrast, Borry (2017) finds no impact between formalization and rule-bending and even a negative association between centralization and rule-bending.

These questions and findings regarding centralization and formalization get to the core of what is expected of civil servants. Despite an acknowledgment that the politics-administration dichotomy does not work in theory or practice, the notion of neutral competence has not withered away (West, 2005; Aberbach & Rockman, 1994). The expectation for administrators to be neutrally competent and avoid being partisan is still a potent ideal (Overeem, 2005).

Administrators should be guided by political leadership (Demir, 2009) and they should have a responsibility to give sound, independent advice in order to help their political superiors (Heclo, 1975). Yet, this ideal is often not translated in practice. At the local level, Tansu Demir and Ronald Nyhan find that “political guidance had a low and nonsignificant impact on the planning ability of public administrators” (2008, p. 91). Furthermore, they find that the neutral competence of administrators is not strong enough to explain why and what administrators choose to do. At the federal level, Terry Moe (1989) posits that administrators have a distinct and separate will from their political superiors, especially in relation to the president. Instead of being neutral, they engage in bureaucratic politics.

As a whole, these findings highlight that administrators are agential. The guerrilla government lenses of ethics, bureaucratic politics, and organizations and management point to civil servants who can and often do go their own way, even if that means going against one’s political superiors. They will be less likely to comply and more likely to engage in guerrilla government when instructions violate an individual’s personal code of ethics. When regular channels of communication are disrupted, or individuals feel like their perspective is not being valued or considered, they will be more likely to employ guerrilla government tactics. And, they have their own agendas and allies. They can leverage these positions and relationships to thwart political planning, exposing civil servants as active and adversarial players.

Data and Methods

Study Design

Participants. Our sampling frame consists of over 460,000 federal bureaucrats, whose information—including name and agency—were scraped from the federal salary database (<https://www.fedsdatacenter.com/federal-pay-rates/>) during April of 2017; work emails were

then constructed based on the name, agency, and email format(s) used in the relevant agency. Constructed email addresses were pinged in August 2017 to ensure accuracy, and the resulting sample consists of 460,369 employees from 95 federal departments and agencies.⁶ This innovative approach to sampling federal government employees has various costs and benefits. The primary benefit is that we are able to reach many more federal government employees than any prior external research has included in their sampling frame. The chief cost is that the federal government did not sanction the survey and many federal employees are instructed to immediately delete any emails from external sources. We are uncertain if we will be able to conduct a survey in this manner in the future, which resulted in a lengthy survey (15-20 minutes) that we requested employees complete without any compensation.

We began sending emails in December, 2017. Initially, because of concerns that email servers might mark the messages as SPAM, we gradually increased the number of email invitations we sent in each batch. We sent the first round of five thousand email invitations in batches of 400 email invitations per day. We only sent invitations on business days to increase the probability that federal employees would see the message as soon as it arrived. It took nearly two weeks for the first round of 5,000 email invitations to reach their intended recipients.

Of the 5,000 email invitations we sent in the first batch, 60% did not make it to their intended recipient because the IT servers marked them as SPAM and blocked them. Of the roughly 2,000 government employees who actually received the email, 18 began the survey and 17 completed the entire survey. The response rate on this first round is just under one percent with a 78% completion rate. By mid-December, 2017 we began sending the emails in batches of

⁶ Due to redactions and problems with pinging, some notable agencies—such as the Departments of Defense and Energy—are omitted from our data.

5,000 per day and we achieved a similar SPAM rate with a response rate of 2.1% and a completion rate of 51%.

In early January, 2018 we decided that sending emails in smaller batches did not change the probability that the federal government IT servers would mark the email invitations as SPAM. As such, we began sending invitations in batches of 65,000/week and scheduled them to arrive on Monday mornings at 8am Eastern Standard Time. The IT servers at some agencies were more likely to tag our emails as SPAM than at other agencies, which caused an imbalance in the number of people who responded by agency. We estimate that 60% of the 458,000 email invitations we sent did not make it to the inbox of the intended recipient. In the end, 6,848 federal employees began our survey (3.7% response rate) of which 44.5% (3,053) completed the entire survey.⁷

Procedure.

Ultimately, we are interested in what influences individual propensities for engaging in guerrilla government. Recall that O’Leary (2014) proposes that guerrilla government can be understood through three lenses: ethics, bureaucratic politics, and organizations and management. Rather than creating several separate experiments manipulating each of these in turn, we opted for a conjoint experimental design because it allows us to see how each works in conjunction with each other (Hainmueller, Hopkins, and Yamamoto 2013).

Table 1 describes each treatment to which a participant might be randomly exposed, and Figure 1 provides one example of what might have been seen by a participant in the survey.

⁷ Complete details about the survey, including tables comparing this survey to other surveys on key demographics and personality traits are in the online appendix.

Table 1: Possible Treatments in the Conjoint Experiment

PERSON TREATMENT	ETHICS TREATMENT	POLICY TREATMENT	SCOPE TREATMENT	RETRIBUTION PROBABILITY TREATMENT	RETRIBUTION TYPE TREATMENT
Person Making the Request is... [TREATMENT]	The proposed policy... [TREATMENT]...your personal code of ethics	You think that the policy is... [TREATMENT]...solution to a problem	If you follow the instructions, there is... [TREATMENT]...probability that many lives will be negatively affected	If you do not follow the instructions, there is... [TREATMENT]...probability that the person making the request will find out what you've done	Should you choose to disobey the instructions, you are likely to be... [TREATMENT]
Co-worker	Completely violates	The best possible	A small	A small	Immediately terminated
Immediate Supervisor	Mostly disagrees with	An appropriate	A decent	A decent	Written up
Agency Appointee	Is unrelated to	An unwise	A high	A high	Praised by outsiders as a whistleblower
U.S. President	Is in line with	The worst possible	An almost certain	An almost certain	Blacklisted
	Reinforces				

It is possible that you would never have to make a choice like the one we are about to present to you. Some of the scenarios may even seem impossible. However, we are interested in what you would do if you had to make a choice like this one. Remember that your responses to this are strictly anonymous.

Scenario 1

Imagine **the President of the United States** asked you to implement a policy that **mostly disagrees with** your personal code of ethics. In your opinion, the policy is **an unwise** solution to the problem. If you follow the instructions, there is a **high** probability that many lives will be negatively affected. If you do not follow the instructions, there is a **decent** possibility that the person making the request will find out what you have done. Should you choose to disobey the instructions, you are likely to be **written up**.

Scenario 2

Imagine **your immediate supervisor** asked you to implement a policy that **is unrelated to** your personal code of ethics. In your opinion, the policy is **an appropriate** solution to the problem. If you follow the instructions, there is a **almost certain** probability that many lives will be negatively affected. If you do not follow the instructions, there is a **decent** possibility that the person making the request will find out what you have done. Should you choose to disobey the instructions, you are likely to be **written up**.

Brief Summary of the two scenarios:

	<i>Scenario 1</i>	<i>Scenario 2</i>
<i>Requested by:</i>	the President of the United States	your immediate supervisor
<i>Ethics:</i>	mostly disagrees with	is unrelated to
<i>Your Personal View:</i>	an unwise	an appropriate
<i>Probability of Harm:</i>	high	almost certain
<i>Probability of Retribution:</i>	decent	decent
<i>Type of Retribution:</i>	written up	written up

Figure 1: Conjoint Treatment as Seen by the Respondents

After viewing the two scenarios, respondents were presented with a list of possible responses from which to choose. These include the following: Follow the instructions, delay

compliance for as long as possible, directly inform the person the decision is wrong, contact a member of the press and tell them what happened, leak the information to an anonymous public source (like Wikileaks). For each possible response, as Figure 2 indicates, the respondent was asked to indicate if they were more likely to have that response in scenario 1 or scenario 2, keeping in mind they might not choose to do any of these.

Below is a list of possible responses a person might have to the above scenarios. Keeping in mind that you might not choose to do any of these, please indicate if you are **more likely** to have that response in Scenario 1 or Scenario 2. Please make a selection for each possible response.

	Scenario 1	Scenario 2
Follow the instructions.	<input type="radio"/>	<input type="radio"/>
Delay compliance for as long as possible.	<input type="radio"/>	<input type="radio"/>
Directly inform the person that the decision is wrong.	<input type="radio"/>	<input type="radio"/>
Contact a member of the press and tell them about what happened.	<input type="radio"/>	<input type="radio"/>
Leak the information to an anonymous public source (like Wikileaks).	<input type="radio"/>	<input type="radio"/>

Figure 2: Conjoint Comparison Responses

Following this, they are asked to rate their absolute likelihood of each of the responses for each scenario; an example of this is presented in Figure 3.

On a scale from 0 to 10, where 0 indicates that you would absolutely not and 10 indicates that definitely would, how likely would you be to *leak information to an anonymous source* in each of the scenarios?

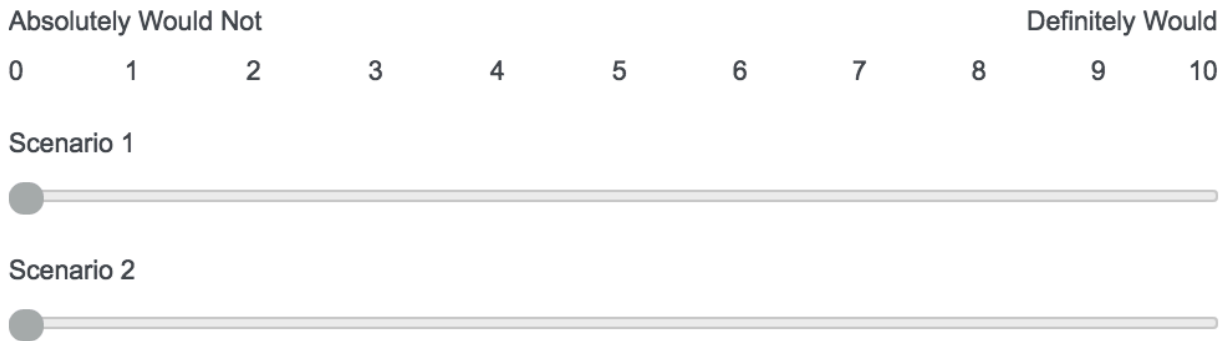


Figure 3: Conjoint Absolute Responses

This experimental design will allow us to isolate the effects of the different treatments while maintaining a balance of internal and external validity. Full randomization means that regression can be used to recover the treatment effects, and the experiment will have higher degrees of realism compared to other experiments that simply vary a single dimension (Bansak et al. 2017; Hainmueller, Hopkins, and Yamamoto 2013).

Results

Consistent with our expectations, the strongest predictors of employing guerrilla government tactics are the extent to which the policy violates an individual's code of ethics and whether the person thinks that the proposal is a good idea. As we are using a conjoint experiment in our analysis, we estimate average marginal component effects (AMCEs), per the recommendations of Hainmueller, Hopkins, and Yamamoto (2013). The AMCE is an estimate of the average extent to which a particular scenario component (e.g., the requester, the respondent's ethical conflicts, etc.) affects the dependent variable. Hainmueller, Hopkins, and Yamamoto

(2013) show that the AMCEs can be estimated by regressing the dependent variables on sets of indicator variables measuring the levels of each attribute; for example, *Ethics-Completely Violates*, *Ethics-Mostly Disagrees With*, *Ethics-Is Unrelated To*, and *Ethics-Reinforces* would be included as independent variables in such a regression to capture the effect of the *Ethics* treatment, with *Ethics-Is In Line With* as the baseline category. In this case, the coefficients recovered using OLS estimation would represent the AMCEs.

In the present context, we estimate models using five dependent variables—*Obey the Directive*, *Confront Requester*, *Delay Compliance*, *Contact Press*, and *Leak Anonymously*—that are each operationalized in two different ways: a *forced choice* operationalization and a *continuous rating* operationalization. For all dependent variables, the *forced choice* operationalization is a binary variable indicating whether the respondent indicated s/he would take the indicated action in the forced-choice question, and the *continuous rating* operationalizations indicate the likelihood of the respondent taking the action; as the latter is rescaled to a [0,1] interval, we interpret this as tantamount to the probability of taking the action in question.⁸

⁸ The AMCE can be thought of as analogous to a regression coefficient in an OLS model, in that it presents the average difference in predicted probability of choosing a particular action—for example, to *Obey*—given different component-level treatments. For example, consider the comparison in Figure 1 and the requested responses in Figure 2 and Figure 3. The respondent in question is implicitly rating two profiles—those listed in Scenario 1 and Scenario 2—and therefore appears in our data twice. If the respondent is more likely to obey in scenario 1, then the observation associated with scenario 1 receives a 1 for the *Obey* variable, and the *Obey* value for scenario 2 would be 0. Each scenario-respondent combination would then enter as individual observations in a regression model, and the estimated coefficients would represent the AMCEs. However, since this approach implies that each respondent appears in the data twice, and each of the other responses—*Confront Requester*, *Delay Compliance*, *Contact Press*, and *Leak Anonymously*—are correlated, we estimate AMCEs with standard errors clustered on respondent. We also present 90% (thick horizontal lines) and 95% (thin horizontal lines) confidence intervals about the estimated AMCEs.

Each panel in Figure 4 presents the estimated effects on the probability—relative to the indicated baseline for that category—that the respondent will obey (in the leftmost panels) or otherwise engage in the indicated form of guerrilla government.^{9,10} The top panel provides the estimates for the forced choice questions, and the bottom provides the estimates for the [0,1] indicators.

Focusing on the leftmost panels of Figure 4, we first examine the probability that someone will obey the direction given the various treatments. Regardless of how obedience is measured, the probability of retribution for disobeying has no consistent influence on whether or not someone will comply with directions, nor does the type of retribution. The status of the person making the request has a significant influence on compliance. If the instruction comes from the president, public employees are significantly more likely to obey than if the instruction comes from an agency appointee or their direct supervisor. However, if the request is made by a co-worker, public servants are significantly less likely to obey the instructions.

⁹ For example, the AMCE for an *Unwise* policy view in the *Obey* model is an estimate of how the probability of obeying changes if the respondent is told that a particular policy solution is “unwise” versus “appropriate.”

¹⁰ The underlying OLS models are in the Appendix.




Figure 4: Estimated Average Marginal Component Effects (AMCEs)

The scope of the policy, how well it conforms to an individual’s code of ethics, and how much someone agrees with the policy have the strongest influence on compliance. If the probability that many lives will be negatively affected is more than “small”, public employees are approximately twelve percentage points less likely to obey the instructions, regardless of

whether the *forced choice* or *continuous rating* operationalization of *Obey* is employed. In addition, when the directive completely violates their code of ethics, or if they perceive that the policy is an unwise course of action, public servants are about twenty percentage points less likely to obey, relative to the baseline of the policy being “in line with” their own code of ethics. For public managers, this suggests that people are more likely to comply with instructions they believe to be correct, regardless of the consequences. Rather than inventing new ways to punish people for failure to comply, it may be more effective to spend time persuading employees of the merits of a particular course of action. Open discussion on the ethical considerations of a policy may do more to invoke compliance than threats of termination.

In fact, neither the type of retribution nor the probability of retribution have much of an effect on individual propensities to engage in guerrilla tactics. In each of the models predicting the probability of engaging in guerrilla government, none of the retribution items achieve statistical significance. In terms of retributive behavior, this pattern of findings suggests that there is little public managers can do to prevent guerrilla government. However, another pattern of findings emerges that is equally instructive for public managers hoping to prevent guerrilla government.

First, the scope of the policy matters. When the probability of harm is small, public servants are significantly more likely to obey and much less likely to resist. However, as the scope climbs from small, to decent, the probability of guerrilla government increases. As such, one thing public managers can do is recognize that whenever a reasonable probability of public harm exists, the probability that public servants will engage in guerrilla tactics increases.

Second, public servants’ opinions are important. When public employees agree with the proposed policy change, they are more likely to obey and less likely to engage in guerrilla

government. A simple change from the employee thinking that the proposal is “appropriate” to an “unwise” course of action predicts a dramatic shift in the probability of compliance. Public employees who think that a policy proposal is unwise are about ten percentage points less likely to obey than if they think that it is an appropriate course of action. Public employees who think a course of action is unwise are ten percentage points more likely to confront their superiors, seven to ten percentage points more likely to delay compliance, five percentage points more likely to go to the press, and five percentage points more likely to leak information to another anonymous source than are those who think that a course of action is an appropriate solution. Effects are even stronger if we set the baseline level of approval to be thinking a particular course of action is “the best” solution to a problem.

Third, ethics make a difference. When a policy proposal does not conform to an individual’s code of ethics, they are significantly less likely to comply. Compared to a policy that is “in line with” a public employees code of ethics, when asked to implement a policy that “mostly disagrees” with their code of ethics a public employee is five to ten percentage points more likely to confront their supervisor, delay compliance, go to the press, or leak to another anonymous source. Although we cannot say what it is that causes a policy to violate an employee’s personal code of ethics, these results suggest that people are less likely to comply with directives that violate their ethics.

Finally, we do not find that individuals disproportionately favor one set of guerrilla tactics over others in response to the various treatments.¹¹ The effect sizes for the various forms

¹¹ While the baseline probabilities of each type of guerrilla government response are different—with the stated likelihoods in the (rescaled to the [0,1] scale) *Continuous Rating* responses ranging from 0.210 (for *Press*) to 0.701 (for *Confront*)—the likelihoods of engaging in the

of resistance are surprisingly similar. In short, once someone decides not to obey, there are many ways to rebel, and the relative likelihoods of each are comparable to one another.

Discussion and Conclusion

In order to better understand the causes of guerrilla government, this paper presents the results of a conjoint experiment embedded in a survey of federal bureaucrats. We find that decisions to engage in guerrilla government (or obey potentially questionable directives) are conditional on a multitude of factors—namely the bureaucrat’s personal views of the directive as a policy solution, the compatibility of the directive with the bureaucrat’s ethical framework, the status of the person issuing the directive, and the probability that the directive might cause harm to others. Notably, these decisions generally are *not* affected by the probability of retribution, or the expected type thereof.

These results provide important context to O’Leary’s (2014) guerrilla government thesis, as they reveal a consistent pattern that can be useful for public managers hoping to motivate compliance and deter guerrilla government. In particular, focusing on the specific guerrilla tactic may prove futile, since the marginal changes in the likelihoods of each type of tactic under analysis were approximately equal to one another. Instead, we recommend that when a policy proposal has the potential for harm, public managers focus their attention on persuading their employees of the merits of the new proposal. To the extent that public managers convince their subordinates that the policy is both “the best approach” and ethically moral, they should be better able to ensure compliance and deter these forms of guerrilla government.

different types of guerrilla government are comparably responsive to the component-level treatments.

Given the importance of ethics and the status of the requester to the decisions to obey or engage in guerrilla government, future research should focus microfounding and further contextualizing these factors. The roles of emotion, personality, and public service motivation should be examined in concert with ethical frameworks as potential interactive factors. For example, it seems plausible that those experiencing more negative emotions at the time a directive is issued—or those with more emotionally unstable or less agreeable personalities (for example)—might react more strongly to perceived violations of ethical norms. Moreover, considering the role of ethical frameworks might help further contextualize recent findings suggesting that individuals with greater public service motivation are more likely to engage in whistleblowing, especially in cases of particularly serious wrongdoing (Caillier 2016).

Our results also have implications for the broader study of bureaucratic politics in cognate disciplines such as economics, political science, sociology, and other social sciences. Standard models of bureaucratic politics often assume that the primary motivations for bureaucratic subversion and/or dissent are ideological in nature, and that bureaucratic discretion should be structured in a way to mitigate potential ideological divergence (e.g., Bendor and Meirowitz 2004; Epstein and O'Halloran 1994, 1999; Gailmard 2002; Hollibaugh 2015, 2017; but see Ting 2008). We show that, while ideological conflict (here, likely captured by the respondents' personal views of the policies as solutions to the problems they are intended to address) certainly affects respondents' propensity to engage in guerrilla government, ideological conflict is not the sole source of subversion or dissent. Ethics and perceived harm to others also matter, and the effects of these considerations on respondents' actions are comparable to those induced by policy-specific considerations. To this extent, our results should push more quantitatively- and formally-inclined scholars to consider nonideological and nonpolicy

incentives for bureaucratic subversion, with the goal of building more realistic formal and empirical models.

References

- Aberbach, Joel D., & Rockman, Bert A. (1994). Civil servants and policymakers: Neutral or responsive competence?. *Governance*, 7(4): 461-469.
- Allison, Graham and Philip Zelikow. 1999. *Essense of Decision: Explaining the Cuban Missile Crisis*. Second Edition. Reading, MA: Longman.
- Bansak, Kirk, Jens Hainmueller, Daniel J Hopkins, and Teppei Yamamoto. 2017. Beyond the breaking point? Survey satisficing in conjoint experiments.
- Barnard, Chester I. 1968. *The Functions of the Executive*. Thirtieth Anniversary Edition. Cambridge, MA: Harvard University Press.
- Bendor, Jonathan and Adam Meirowitz. 2004. "Spatial Models of Delegation." *American Political Science Review* 98(2): 293-310.
- Borry, Erin L. 2017. "Ethical Climate and Rule Bending: How Organizational Norms Contribute to Unintended Rule Consequences." *Public Administration* 95(1): 78-96.
- Caillier, James Gerard. 2017. "Public Service Motivation and Decisions to Report Wrongdoing in U.S. Federal Agencies: Is This Relationship Mediated by the Seriousness of the Wrongdoing." *American Review of Public Administration* 47(7): 810-825.
- Cho, Yoon Jik and Hyun Jin Song. 2015. "Determinants of Whistleblowing Within Government Agencies." *Public Personnel Management* 44(4): 450-472.
- DeHart-Davis, Leisha. 2007. "The Unbureaucratic Personality." *Public Administration Review* 67(5): 892-903.

- Demir, Tansu. (2009). Politics and Administration: Three Schools, Three Approaches, and Three Suggestions. *Administrative Theory & Praxis*, 31(4): 503-532.
- Demir, Tansu., & Nyhan, Ronald. C. (2008). The Politics-Administration Dichotomy: An Empirical Search for Correspondence between Theory and Practice. *Public Administration Review*, 68(1): 81-96.
- Dubnick, Melvin J. 1999. "Demons, Spirits, and Elephants: Reflections on the Failure of Public Administration Theory." Paper presented at the Annual Meeting of the American Political Science Association, Atlanta, Georgia, September 2-5 1999.
- Epstein, David and Sharyn O'Halloran. 1994. "Administrative Procedures, Information, and Agency Discretion." *American Journal of Political Science* 38(3): 697-722.
- Epstein, David and Sharyn O'Halloran. 1999. *Delegating Powers: A Transaction Cost Politics Approach to Policy Making Under Separate Powers*. New York: Cambridge University Press.
- Flyvbjerg, Bent. 1998. *Rationality and Power: Democracy in Practice*. Chicago: University of Chicago Press.
- Frederickson, H. George and David K. Hart. 1985. "The Public Service and the Patriotism of Benevolence." *Public Administration Review* 45(5): 547-553.
- Gailmard, Sean. 2002. "Expertise, Subversion, and Bureaucratic Discretion." *Journal of Law, Economics, and Organization* 18(2): 536-555.
- Gofen, Anat. 2014. "Mind the Gap: Dimensions and Influence of Street-Level Divergence." *Journal of Public Administration Research and Theory* 24(2): 473-493.

- Hainmueller, Jens, Daniel J. Hopkins, and Teppei Yamamoto. 2013. "Causal Inference in Conjoint Analysis: Understanding Multidimensional Choices via Stated Preference Experiments." *Political Analysis* 22(1): 1-30.
- Hart, David K. 1984. "The Virtuous Citizen, the Honorable Bureaucrat, and "Public" Administration." *Public Administration Review* 44(SI): 111-120.
- Heclo, Hugh. (1975). OMB and the Presidency-the Problem of Neutral Competence. *The Public Interest*, (38): 80.
- Hollibaugh, Gary E., Jr. 2015. "Naïve Croynism and Neutral Competence: Patronage, Performance, and Policy Agreement in Executive Appointments." *Journal of Public Administration Research and Theory* 25(2): 341-372.
- Hollibaugh, Gary E., Jr. 2017. "The Incompetence Trap: The (Conditional) Irrelevance of Agency Expertise." *Journal of Public Administration Research and Theory* 27(2): 217-235.
- Kaufman, Herbert. 2006. *The Forest Ranger: A Study in Administrative Behavior*. Special Reprint Edition. Washington, DC: RFF Press.
- Lavena, Cecilia Florencia. 2016. "Whistle-Blowing: Individual and Organizational Determinants of the Decision to Report Wrongdoing in the Federal Government." *American Review of Public Administration* 46(1): 113-136.
- Lipsky, Michael. 2010. *Street-Level Bureaucracy: Dilemmas of the Individual in Public Services*. Thirtieth Anniversary Edition. New York: Russell Sage Foundation.

- Maynard-Moody, Steven and Michael Musheno. 2003. *Cops, Teachers, Counselors: Stories from the Front Lines of Public Service*. Ann Arbor, MI: University of Michigan Press.
- Meier, Kenneth J. and John Bohte. 2007. *Politics and the Bureaucracy: Policymaking in the Fourth Branch of Government*. Fifth Edition. Belmont, CA: Wadsworth Cengage Learning.
- Miceli, Marcia. P. and Janet. P. Near. 1992. *Blowing the Whistle: The Organizational and Legal Implications for Companies and Employees*. New York: Lexington.
- Moe, Terry M. (1985). The Politicized Presidency. In J. Chubb & P. E. Peterson (Eds.), *The New Direction in American Politics* (pp. 235-271). Washington, D.C.: Brookings Institution.
- Near, Janet P. and Marcia P. Miceli. 1996. "Whistle-Blowing: Myth and Reality." *Journal of Management* 22(3): 507-526.
- Near, Janet P. and Marcia P. Miceli. 2008. "Wrongdoing, Whistle-Blowing, and Retaliation in the U.S. Government: What Have Researchers Learned From the Merit Systems Protection Board (MSPB) Survey Results?" *Review of Public Personnel Administration* 28(3): 263-281.
- Newswander, Chad B. 2012. "Moral Leadership and Administrative Statesmanship: Safeguards of Democracy in a Constitutional Republic." *Public Administration Review* 72(6): 886-874.
- Newswander, Chad B. 2015. "Guerrilla Statesmanship: Constitutionalizing an Ethic of Dissent" *Public Administration Review* 75(1): 126-134.

- O'Leary, Rosemary. 2014. *The Ethics Of Dissent: Managing Guerrilla Government*. Second Edition. Washington, DC: CQ Press.
- O'Leary, Rosemary. 2017. "The New Guerrilla Government: Are Big Data, Hyper Social Media and Contracting Out Changing the Ethics of Dissent?" *PS: Political Science & Politics* 50(1): 12-22.
- Overeem, Patrick. 2005. "The Value of The Dichotomy: Politics, Administration, and The Political Neutrality of Administrators." *Administrative Theory & Praxis* 27(2): 311-329.
- Rohr, John A. 1986. *To Run a Constitution: The Legitimacy of the Administrative State*. Lawrence, KS: University Press of Kansas.
- Rohr, John A. 1998. *Public Service, Ethics & Constitutional Practice*. Lawrence, KS: University Press of Kansas.
- Spicer, Michael W. 1995. *The Founders, the Constitution, and Public Administration: A Conflict in Worldviews*. Washington, DC: Georgetown University Press.
- Thompson, Dennis F. 1985. "The Possibility of Administrative Ethics." *Public Administration Review* 45(5): 555-561.
- Ting, Michael M. 2008. "Whistleblowing." *American Political Science Review* 102(2): 249-267.
- Wamsley, Gary L. (1990). "The Agency Perspective: Public Administration as Agential Leaders." In *Refounding Public Administration*, eds. Gary L. Wamsley et al., 19-29. Newbury Park, CA: SAGE Publications.
- Wamsley, Gary L. and James F. Wolf. 1996. *Refounding Democratic Public Administration: Modern Paradoxes, Postmodern Challenges*. Thousand Oaks, CA: SAGE Publications.

West, William. F. (2005). Neutral competence and political responsiveness: an uneasy relationship. *Policy Studies Journal*, 33(2): 147-160.

Wilson, James Q. 1989. *Bureaucracy: What Government Agencies Do and Why They Do It*. New York: Basic Books.

Appendix

The American Government Employees Survey (AGES)

Our sampling frame consists of over 460,000 federal bureaucrats, whose information—including name and agency—were scraped from the federal salary database (<https://www.fedsdatacenter.com/federal-pay-rates/>) during April of 2017; work emails were then constructed based on the name, agency, and email format(s) used in the relevant agency. Constructed email addresses were pinged in August 2017 to ensure accuracy, and the resulting sample consists of 460,369 employees from 95 federal departments and agencies. This innovative approach to sampling federal government employees has various costs and benefits. The primary benefit is that we are able to reach many more federal government employees than any prior external research has included in their sampling frame. The chief cost is that the federal government did not sanction the survey and many federal employees are instructed to immediately delete any emails from external sources. We are uncertain if we will be able to conduct a survey in this manner in the future, which resulted in a lengthy survey (15-20 minutes) that we requested employees complete without any compensation.

We began sending emails in December, 2017. Initially, because of concerns that email servers might mark the messages as SPAM, we gradually increased the number of email invitations we sent in each batch. We sent the first round of five thousand email invitations in batches of 400 email invitations per day. We only sent invitations on business days to increase the probability that federal employees would see the message as soon as it arrived. It took nearly two weeks for the first round of 5,000 email invitations to reach their intended recipients.

Of the 5,000 email invitations we sent in the first batch, 60% did not make it to their intended recipient because the IT servers marked them as SPAM and blocked them. Of the

roughly 2,000 government employees who actually received the email, 18 began the survey and 17 completed the entire survey. The response rate on this first round is just under one percent with a 78% completion rate. By mid-December, 2017 we began sending the emails in batches of 5,000 per day and we achieved a similar SPAM rate with a response rate of 2.1% and a completion rate of 51%.

In early January, 2018 we decided that sending emails in smaller batches did not change the probability that the federal government IT servers would mark the email invitations as SPAM. As such, we began sending invitations in batches of 65,000/week and scheduled them to arrive on Monday mornings at 8am Eastern Standard Time. The IT servers at some agencies were more likely to tag our emails as SPAM than at other agencies, which caused an imbalance in the number of people who responded by agency. We estimate that 60% of the 458,000 email invitations we sent did not make it to the inbox of the intended recipient. In the end, 6,848 federal employees began our survey (3.7% response rate) of which 44.5% (3,053) completed the entire survey.

The demographic characteristics of these panels closely resemble that of the United States population on several important traits. The primary difference is that our sample is more educated than the general population. Table A.1 displays the demographics of this sample compared to American Community Survey (2014), Amazon's Mechanical Turk (adapted from Berinsky, Huber and Lenz (2012)), and two large-scale more nationally representative survey samples, the Annenberg National Election Study (Johnston et al. 2008) and the Cooperative Congressional Election Study (Ansolabehere and Shaffner 2018). Amazon's Mechanical Turk is an online marketplace where people hire laborers for a variety of tasks. Since the mid-2000's researchers have been offering people money to participate in online survey experiments through

Amazon’s Mechanical Turk. Recently, scholars have spent considerable effort trying to determine the quality of the samples that are usually obtained through this service (Mullinix et al. 2015). The following table shows that this sample is much more representative of the US population on key variables than samples obtained through Amazon’s Mechanical Turk and largely identical to the nationally representative sample collected in the Annenberg National Election Study.

Table A-1: Summary of General Population Survey Demographics

Demographics	AGES 2018	ACS 2014 Estimates	MTurk	NAES 2008	CCES 2016
Female	52.2%	50.8%	60.1%	56.62%	51.7%
Age (mean years)	50-54	37.4 (median)	20.3	50.05	48
Education (% completing some college)	96%	-	-	62.86%	58%
White	67.5%	73.8%	83.5%	79.12%	72.4%
Black	12.5%	12.6%	4.4%	9.67%	12.2%
Asian	3.2%	5.0%	-	2.53%	3.5%
Latinx	6.4%	16.9%	-	6.3%	7.3%
Multi-Racial	4.7%	2.9%	-	2.37%	2.3%
Party Identification					
Democrat	30.5%	-	40.8%	36.67%	36.2%
Independent	30.1%	-	34.1%	20.82%	31.7%
Republican	24.6%	-	16.9%	30.61%	28.4%
Other Party/DK	14.6%				
<i>N</i>	3,053	-	484-551	19,234	52,900

Table A-2: OLS Results for *Forced Choice* Questions

	Obey	Confront	Delay	Press	Leak
Type of Retribution (Baseline = Praised as Whistleblower)					
Blacklisted	-0.032* (0.018)	0.013 (0.018)	0.018 (0.018)	-0.018 (0.018)	-0.007 (0.018)
Immediate Termination	-0.022 (0.018)	0.033* (0.018)	0.012 (0.018)	0.002 (0.018)	0.006 (0.018)
Written Up	-0.003 (0.017)	0.042** (0.017)	0.018 (0.017)	-0.011 (0.017)	-0.013 (0.017)
Probability of Retribution (Baseline = Small)					
R: Almost Certain	-0.014 (0.017)	0.004 (0.018)	0.018 (0.018)	0.022 (0.018)	0.026 (0.018)
R: Decent	0.001 (0.017)	0.002 (0.018)	-0.003 (0.018)	-0.001 (0.018)	-0.008 (0.018)
R: High	-0.013 (0.018)	0.007 (0.018)	0.008 (0.018)	0.013 (0.018)	-0.001 (0.018)
Probability of Harm (Baseline = Small)					
H: Almost Certain	-0.102*** (0.017)	0.068*** (0.017)	0.083*** (0.017)	0.082*** (0.017)	0.064*** (0.017)
H: Decent	-0.075*** (0.017)	0.063*** (0.017)	0.040** (0.017)	0.035** (0.017)	0.050*** (0.017)
H: High	-0.108*** (0.017)	0.104*** (0.017)	0.096*** (0.017)	0.081*** (0.017)	0.081*** (0.017)
Personal View of Policy (Baseline = Appropriate)					
The Best	0.014 (0.017)	-0.030* (0.017)	-0.026 (0.017)	-0.019 (0.017)	-0.030* (0.017)
The Worst	-0.134*** (0.017)	0.128*** (0.017)	0.098*** (0.017)	0.091*** (0.017)	0.074*** (0.017)
Unwise	-0.111*** (0.017)	0.101*** (0.017)	0.067** (0.017)	0.056*** (0.018)	0.048*** (0.018)
Agreement with Ethics (Baseline = In Line With)					
Completely Violates	-0.176*** (0.019)	0.157*** (0.019)	0.129*** (0.020)	0.108*** (0.020)	0.109*** (0.020)
Mostly Disagrees	-0.103*** (0.019)	0.087*** (0.019)	0.079*** (0.020)	0.058*** (0.020)	0.075*** (0.020)
Reinforces	-0.057*** (0.019)	0.027 (0.019)	0.000 (0.019)	0.003 (0.020)	0.022 (0.020)
Unrelated	-0.054*** (0.020)	0.027 (0.020)	0.042** (0.020)	0.031 (0.020)	0.035* (0.020)
Person Making Request (Baseline = Immediate Supervisor)					
A Coworker	-0.099*** (0.018)	-0.011 (0.018)	0.009 (0.018)	0.009 (0.018)	-0.002 (0.018)
An Agency Appointee	-0.046*** (0.018)	-0.058*** (0.017)	0.013 (0.018)	0.030* (0.018)	0.014 (0.018)
The President	0.038** (0.018)	-0.145*** (0.018)	-0.020 (0.018)	-0.008 (0.018)	-0.012 (0.018)
Constant	0.754*** (0.027)	0.360*** (0.027)	0.342*** (0.027)	0.369*** (0.027)	0.380*** (0.027)
Number of Respondent-Profiles	6426	6426	6426	6426	6426

R ²	0.050	0.051	0.026	0.020	0.017
Adjusted R ²	0.047	0.048	0.024	0.017	0.015
F-Test	17.576***	17.939***	9.172***	6.903***	5.996***

Note: Coefficients are OLS estimates recovered from regressing the binary forced choice responses on the indicators for the references component treatments. Per Hainmueller, Hopkins, and Yamamoto (2013), the coefficients are equivalent to the AMCEs. Each respondent contributed two observations to the data, since each respondent was asked to choose between two different profiles. As such, we cluster our standard errors (in parentheses) on the respondent. Two-tailed tests: *** p < 0.01, ** p < 0.05, * p < 0.1

Table A-3: OLS Results for *Continuous Rating Questions*

	Obey	Confront	Delay	Press	Leak
Type of Retribution (Baseline = Praised as Whistleblower)					
Blacklisted	-0.013 (0.014)	0.018 (0.014)	0.028* (0.015)	0.010 (0.012)	-0.001 (0.013)

Immediate Termination	0.033** (0.014)	-0.016 (0.014)	0.010 (0.015)	0.010 (0.012)	-0.004 (0.012)
Written Up	0.026* (0.014)	0.008 (0.014)	-0.002 (0.016)	-0.005 (0.013)	-0.014 (0.013)
Probability of Retribution (Baseline = Small)					
R: Almost Certain	-0.017 (0.014)	0.030** (0.014)	0.017 (0.015)	0.021 (0.013)	0.010 (0.013)
R: Decent	-0.003 (0.014)	0.039*** (0.014)	0.028* (0.015)	0.024* (0.012)	0.028** (0.013)
R: High	-0.026* (0.014)	0.024* (0.014)	0.041*** (0.015)	0.019 (0.012)	0.012 (0.013)
Probability of Harm (Baseline = Small)					
H: Almost Certain	-0.089*** (0.014)	0.079*** (0.014)	0.073*** (0.016)	0.043*** (0.013)	0.036*** (0.013)
H: Decent	-0.055*** (0.014)	0.050*** (0.015)	0.052*** (0.015)	0.026** (0.013)	0.016 (0.013)
H: High	-0.081*** (0.014)	0.070*** (0.014)	0.061*** (0.015)	0.028** (0.012)	0.025** (0.012)
Personal View of Policy (Baseline = Appropriate)					
The Best	0.034** (0.014)	-0.036** (0.015)	-0.014 (0.015)	0.008 (0.012)	0.001 (0.012)
The Worst	-0.148*** (0.014)	0.094*** (0.014)	0.132*** (0.015)	0.067*** (0.013)	0.061*** (0.013)
Unwise	-0.102*** (0.014)	0.096*** (0.014)	0.101*** (0.015)	0.048*** (0.013)	0.048*** (0.013)
Agreement with Ethics (Baseline = In Line With)					
Completely Violates	-0.171*** (0.015)	0.116*** (0.015)	0.115*** (0.017)	0.062*** (0.015)	0.046*** (0.014)
Mostly Disagrees	-0.096*** (0.015)	0.056*** (0.016)	0.079*** (0.017)	0.042*** (0.014)	0.039*** (0.014)
Reinforces	-0.013 (0.015)	0.025 (0.016)	0.022 (0.017)	0.024* (0.013)	0.017 (0.014)
Unrelated	-0.067*** (0.016)	0.011 (0.016)	0.039** (0.017)	0.031** (0.014)	0.017 (0.014)
Person Making Request (Baseline = Immediate Supervisor)					
A Coworker	-0.104*** (0.014)	0.023* (0.013)	0.038** (0.015)	-0.011 (0.012)	-0.006 (0.012)
An Agency Appointee	-0.002 (0.014)	-0.045*** (0.013)	-0.006 (0.015)	0.021* (0.012)	0.027** (0.012)
The President	0.056*** (0.014)	-0.112*** (0.015)	-0.050*** (0.015)	0.030** (0.013)	0.020 (0.012)
Constant	0.616*** (0.022)	0.577*** (0.023)	0.403*** (0.024)	0.094*** (0.019)	0.126*** (0.019)
Number of Observations	4361	4361	4361	4361	4361
R ²	0.119	0.080	0.059	0.020	0.017
Adjusted R ²	0.115	0.076	0.055	0.016	0.013
F-Test	30.754***	19.846***	14.258***	4.729***	3.939***

Note: Coefficients are OLS estimates recovered from regressing the continuous rating responses on the indicators for the references component treatments. Per Hainmueller, Hopkins, and Yamamoto (2013), the coefficients are equivalent to the AMCEs. Each respondent contributed up to two observations to the data, since each respondent was asked to rate two different profiles; however, some respondents only rated one profile. As such, we cluster our standard errors (in parentheses) on the

respondent. Two-tailed tests: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Online Appendix References

- Ansolabehere, Stephen, and Brian F. Schaffner. 2018. "CCES Common Content, 2016." In: Harvard Dataverse.
- Berinsky, Adam J, Gregory A Huber, and Gabriel S Lenz. 2012. "Evaluating Online Labor Markets for Experimental Research: Amazon.com's Mechanical Turk." *Political Analysis* 20(3):351-68.
- Johnston, Richard, Kathleen Jameison, and Diana Carole Mutz. 2008. "National Annenberg Election Survey." ed. A.P.P. Center.
<https://services.annenbergpublicpolicycenter.org/naes08/phone/index.html> (accessed December 21, 2017).
- Mullinix, Kevin J., Thomas J. Leeper, James N. Druckman, and Jeremy Freese. 2015. "The Generalizability of Survey Experiments." *Journal of Experimental Political Science* 2 (2):109-38.