

Article

How Do Politicians Attribute Bureaucratic Responsibility for Performance? Negativity Bias and Interest Group Advocacy

Poul A. Nielsen,* Donald P. Moynihan†

*University of Southern Denmark; †University of Wisconsin-Madison

Address correspondence to the author at aaes@sam.sdu.dk.

Abstract

Voters reward or punish politicians by deeming them responsible for positive and negative outcomes, but how, in turn, do politicians attribute responsibility to those who actually deliver public services? Inattention to this question renders incomplete current perspectives on democratic processes of accountability, even as politicians are increasingly provided with performance data to hold bureaucrats accountable. We shed light on this issue using a survey experiment of elected officials featuring actual performance data. We find that the provision of performance data makes elected officials more willing to attribute causal responsibility to bureaucratic leaders, but only in cases of low performance, suggesting a negativity bias in public sector responsibility attribution processes. Additionally, we offer evidence that interest group advocates influence how elected officials use performance information to attribute responsibility, but contingent on ideological alignment.

Introduction

While “the concept of responsibility lies at the heart of theories of democratic accountability” (Rudolph 2003, 698) scholars have primarily focused on how voters assign responsibility to politicians, or government generically (Tilley and Hobolt 2011). Yet for democratic accountability to function, elected representatives are tasked, in turn, with making judgments about bureaucratic responsibility for the outcomes of government programs.

How do elected officials engage in responsibility attribution for the bureaucrats under their supervision? In addition to broadening the study of responsibility attribution, this question is relevant for practice because of changing patterns of governance. Responsibility attributions to public organizations or individual bureaucrats have historically been made

difficult by the complexity of public tasks, and the consequent lack of reliable indicators of effectiveness. But politicians increasingly demand that public organizations provide formal quantitative indicators of their performance (James 2011), with the expectation that elected officials will use the data to hold bureaucrats accountable.

A case in point is the education policy field we study. Here, policy reforms have explicitly tied performance management to responsibility attribution processes for bureaucrats. This is perhaps most apparent in the United States, where the No Child Left Behind Act of 2002 was, according to President Bush, “based on the fundamental notion that an enterprise works best when responsibility is placed closest to the most important activity of the enterprise, when those responsible are given greatest latitude and support, and when those responsible are held accountable for producing results.”¹ Both the No Child

We would like to thank Martin Bækgaard who helped develop the treatments. We would also like to thank Martin Bisgaard and three anonymous reviewers for their helpful comments and suggestions.

1 <http://georgewbush-whitehouse.archives.gov/news/reports/no-child-left-behind.html>

Left Behind Act and its Obama administration sequel, *Race to the Top*, institutionalized the use of performance metrics to evaluate school principals, and made the removal of school principals a first step in dealing with schools determined to be underperforming. Such policy changes are premised on the notion that it is possible and desirable to hold public managers personally responsible for outcomes that occur under their watch.

While there is growing attention to bureaucratic and even citizen use of performance information (e.g., Andersen and Moynihan 2016; James and Van Ryzin 2016; James et al. 2016; Liang and Langbein 2015; Marvel 2016; Meier, Favero, and Zhu 2015), we know little about how elected officials actually process and respond to such data for purposes of holding bureaucratic leaders responsible. Moreover, performance information does not exist in isolation. Interested stakeholders consistently advocate to politicians for particular interpretations, or contest the credibility and relevance of performance data (Moynihan 2008). Such advocacy may well influence both the willingness of elected officials to engage in attribution processes, and the manner in which they do so.

To examine these questions we mix the causal advantages of an experimental design with the external validity that comes from studying subjects and treatments that closely reflect the actual policy process. We surveyed actual local council elected officials in Denmark about a policy area they supervised, asking them to assess the degree of influence that their subordinate school principals had on performance outcomes.

We find that only among jurisdictions where performance is relatively low do elected officials—when they are provided school test score data—become more willing to attribute causal responsibility to school principals for student outcomes, suggesting a negativity bias in attribution of responsibility to bureaucrats (Kahneman and Tversky 1979; Lau 1982). This finding is consistent with a growing literature suggesting a negativity bias in the use of public sector performance data (Boyne et al. 2009; Charbonneau and Bellavance 2012; Craig, Imberman, and Purdue 2015; James and John 2007; Marvel 2016), but breaks new ground by showing how it affects elected officials' use of data to assess bureaucratic responsibility.

We also find that interest group advocacy about performance measures matter greatly to how elected officials use these data for responsibility attribution, but contingent on ideological alignment. Public sector performance data does not arrive in a depoliticized environment. Interest groups act upon their policy preferences by advocating for one set of performance data that serve their purposes, or discrediting data that does not. A treatment that featured a teachers union representative casting doubt on the validity of test

scores exacerbated partisan polarization in responsibility attribution, consistent with theories of motivated skepticism (Taber and Lodge 2006). When the interest group and politician shared the same liberal ideological space, politicians became less likely to use test scores to attribute responsibility to school principals, even as data showed relatively low performance. By contrast, conservative politicians ignored the comments of teachers unions. The result is that conservative and liberal politicians, when exposed to the type of performance data and data advocacy present in actual policy processes, arrive at very different judgments about what bureaucrats may be held accountable for.

Responsibility Attribution and Performance Data

Social psychologists point to different ways by which individuals attribute responsibility. The two most salient for democratic processes of accountability are role or functional responsibility—responsibility for fulfilling a particular set of duties—and causal responsibility—where public officials are held to account for actual outcomes, not just the duties assigned to their function (Hamilton 1978). While perceived instrumentality increases causal attribution, governments may be held causally responsible even if their actions are not clearly instrumental in contributing to outcomes (Arceneaux 2006). A growing literature has documented voters' willingness to attribute causal responsibility to elected officials for salient outcomes that elected officials have limited control over, such as the national economy, or response to disasters (Bisgaard 2015; Malhotra and Kua 2008; Rudolph 2003; Tilley and Hobolt 2011).

Absent from the literature on responsibility attribution is a consideration of the processes by which the politicians themselves engage in such judgments over bureaucrats. Democratic processes of accountability depend upon what elected officials believe bureaucrats can be held responsible for, but the context of responsibility attribution will necessarily be different for politicians than for voters. While citizens have choices over whom they allocate responsibility to, most obviously between different political parties, political principals are more constrained in their options, largely limited to the bureaucrats that work for them.

This question of responsibility attribution is also a central—though understudied—aspect of blame avoidance, since actors must be perceived to be responsible before they can be blamed (Hood 2011). Although we do not study actual blaming, we examine the necessary preceding step: responsibility attribution. The context in which elected officials engage in responsibility attribution—and by extension blame—has become increasingly saturated with performance measures that make it more tempting for elected officials to apply a

standard of causal rather than functional accountability to individual bureaucrats (Heckman et al. 2011; Moynihan 2008).

There are two causal logics for why performance data might generally trigger responsibility attribution. The first is that data increases confidence in making judgments about complex social processes (Sniezek 1992; Tsai, Klayman, and Hastie 2008). Although policymakers may have some beliefs about performance tied to organizational reputation (Carpenter and Krause 2012), they may recognize such judgments to be approximate, and therefore are less willing to make consequential decisions based on them. However, providing individuals with specific measures of performance may serve to reduce or remove such caution. As politicians believe their judgments are based on precise quantitative evidence they may enjoy increased confidence in their assessments of organizational performance and its causes, fostering a willingness to make attributions about the organization and the organizational leader.

The second reason why elected officials might use performance data to attribute responsibility to bureaucratic actors is because they themselves have reason to fear data-based attributions. A number of studies find that performance ratings for schools (Berry and Howell 2007) or local government (Boyne et al 2009; James and John 2007) affect the vote share for incumbent school board members and local politicians in the next election. Holbein (2016) provides evidence that performance data increased voter turnout and competitiveness in local school board elections. James and Moseley (2014) moreover find that the provision of performance data can increase citizen perceptions that local governments are responsible for outcomes in a recycling program. James et al. (2016) provide evidence of a strong political incentive for elected officials to shift blame onto managers, showing that providing citizens information that emphasizes political oversight of failing services increased citizen attribution of blame to local politicians, but that providing information about bureaucratic oversight of the same services reduced blame toward politicians.

While the literature of blame avoidance does not consider performance data explicitly, it does consider the asymmetric nature of risk and credit. Both Weaver (1986, 1987) and Hood (2011) portray public bureaucracies as inherently at greater risk of blame than credit-claiming opportunities. Rational politicians are therefore blame minimizers rather than credit-claiming maximizers (Weaver 1987). Politicians have a variety of strategies to deflect blame, which often centers on avoiding risky policies, or putting distance between themselves and implementation processes that could go wrong. But for some services, these options may

not be feasible, whether because the policy is perceived as an inherently public function, or due to strong stakeholder preferences (Moynihan 2012). In such conditions, elected officials have fewer options, and the temptation to “deflect blame by blaming others” (Weaver 1986, 385) or “delegate responsibility down the line” (Hood 2011) becomes stronger.

Therefore, if elected officials sense that performance data will be used to judge them, they may become more likely to engage in judgments of bureaucratic leaders. The blame avoidance perspective proposes that elected officials should be especially motivated to attribute responsibility elsewhere in cases of low performance. Indeed, the empirical study of blame avoidance and reputation portrays public organizations as more motivated to avoid portrayals of failure than to take risks for success (Gilad, Maor, and Bloom 2013; Mortensen 2012). This negativity bias is also evidenced in individual-level psychological work on attitude formation (Skowronski and Carlston 1989) and loss aversion (Kahneman and Tversky 1979), and is a basic psychological trait that has been observed in a variety of settings (Rozin and Royzman 2001).

Survey experiments demonstrate that negativity bias shapes how members of the public process performance data. For example, citizen support for programs increases when the same performance statistics are presented in terms of levels of success (e.g., client satisfaction) rather than failure (rates of dissatisfaction) (Olsen 2015a). More detailed analyses find that citizens respond to neutral and positive framing of performance information in similar ways, but are much more responsive to negative information in evaluating service quality and in engaging in spontaneous responsibility attribution (Olsen 2015b). Simply providing citizens with comparative performance data has, on aggregate, the effect of depressing citizen evaluations of their local schools (Barrows et al. 2016). The tendency to be critical of public outcomes may also be shaped by ideological beliefs. For example, Baekgaard and Serritzlew (2016) offer evidence that citizens who have a pre-existing preferences for private provision of goods tend to process performance data about equivalent public and private performance more negatively for public organizations. Marvel (2016) offers evidence that citizens bombarded with a steady stream of critical messages about government hold a negative implicit bias that prevents them from giving credit even when public sector performance is high.

Citizens become more critical of political incumbents when provided with data showing poor government performance, but are unwilling to assign equivalent credit when performance is good (James 2011; James and Moseley 2014). The risks of performance data can extend to the ballot box. Studies of the link between English

local government performance metrics and incumbent vote shares show that low performance scores lead to a loss of support for local elected officials in elections, but positive performance scores were not equivalently rewarded (Boyne et al 2009; James and John 2007). Low performance scores also limit the capacity to win popular support for more revenue. Kogan, Lavertu, and Peskowitz (2016) show that districts whose schools are labeled as failing under No Child Left Behind Act provisions have a lower probability of winning new resources via referenda on local school tax levies. An examination of the effects of No Child Left Behind on turnout found that it increased both turnout and exit from schools but only in conditions where schools were labeled as failing (Holbein 2016).

There is some evidence that policymakers also exhibit a concern with the negative performance assessments in their decisions. One study of Texas school administrators used a regression discontinuity design to illustrate the risks of characterizing essentially equivalent schools as high or low performing: schools just marginally below the cutoff line for ratings received significantly lower budgets than those just above the rating line (Craig, Imberman, and Purdue 2015). Public managers may be aware of the risk of negative performance scores, and take pains to protect themselves. For example, low performance scores in English local government increases the risk of turnover for management teams, though not for the chief executive, which Boyne et al. (2010) suggest indicates patterns of blame-shifting. A study of Canadian municipalities showed that those with lower scores tended to present more detailed justifications in performance reports to deflect blame (Charbonneau and Bellavance 2012).

Collectively then, while performance data might enable both changes for credit claiming and blame in public organizations, the existing empirical literature on public organizations suggests that the risk of blame for bureaucratic leaders outweighs the opportunity for credit. However, we lack direct evidence about whether elected officials' willingness to attribute responsibility—and blame—is triggered by the provision of negative information. Therefore, we propose the following negativity bias hypothesis:

- H1. When provided with performance data, elected officials' attribution of responsibility to public sector leaders will be most pronounced in cases of low performance scores.

Interest Groups and Data Advocacy

Interest groups routinely attempt to influence how elected officials interpret information (Mahoney and Baumgartner 2008; Moe 2011) though there is little attention to how

they fulfill this role for performance data or attributions of bureaucratic responsibility. Broad theoretical considerations of blame avoidance have not addressed how the rise of performance regimes coincides with the role of advocates in affecting attributions of responsibility. Such advocates obviously exist and seek influence in the policy and implementation processes. Moynihan's (2008) interactive dialogue model considers how they may play a role in the context of the growing turn to performance data. The model proposes that the ambiguous nature of performance data creates an opportunity of advocates to spin data in a way that is supportive or detrimental to an agency or program. In a context where elected officials are exposed to performance data, interest groups may seek to influence decisions by promoting some performance metrics, contesting the relevance or validity of others, and offering interpretations about the meaning and implications of data (Moynihan 2008).

Here, we focus specifically on whether interest groups casting doubt on performance data alters the willingness of elected officials to use such data to engage in responsibility attribution. In debates about the performance of public organizations, a frequently invoked argument is that performance data in general, or the specific performance data being relied upon, is unreliable, incomplete, or misleading (Dixit 2002; Moynihan 2008). For example, Berry and Howell's (2007) analysis of the relationship between school test scores and electoral support of school board candidates found a significant relationship in one time period, not in another. They speculate that the relationship weakened partly because "interest groups and the public education sector attempted to discredit the state's testing regime" resulting in more skeptical media coverage, lowering voter's willingness to employ performance data in making attributions (Berry and Howell 2007).

If, as hypothesis 1 suggests, the provision of performance data might increase the willingness to engage in responsibility attribution, how do interest group efforts to engage in data advocacy affect the attribution process? While interest groups could potentially be generally persuasive in influencing how performance data is used, the relationship between the audience and the advocate likely matters. A central assumption in the mapping of ideological spaces is that political actors are attracted to those who share the same ideological beliefs (e.g., Bonica 2013), finding them more credible and trustworthy (Cohen 2003; Slothuus and de Vreese 2010). This is also consistent with theory (Moynihan 2008) and empirical evidence (Baekgaard and Serritzlew 2016; James and Van Ryzin 2016) that political ideology triggers motivated reasoning in the processing of performance data.

The literature on voter attribution of responsibility points to the importance of partisan or ideological relationships in responsibility attribution. As ideological

distance from the current government grows, partisans attribute greater responsibility to government for poor performance and deny credit for positive outcomes, whereas partisan-aligned voters tend to absolve government of responsibility for failure and attribute credit for success (Bisgaard 2015; Malhotra and Kua 2008; Tilley and Hobolt 2011).

Being more generous to those with shared ideological beliefs is a form of motivated reasoning operating at the group level, inducing in-group members to attribute positive outcomes to the deliberate actions of their own group, whereas negative outcomes are attributed to the actions of out-group members or to external causes beyond their control (Tilley and Hobolt 2011). The process of responding to cues of shared ideological beliefs has been convincingly demonstrated in work on public opinion that shows that party frames are important to attitude formation as voters engage in partisan-based motivated reasoning, and that party effects are more prominent for contested political issues and for the more politically aware (Slothuus and de Vreese 2010). For example, although partisan voters came to agree on the seriousness of the United Kingdom recession between 2004 and 2010, they polarized into opposite camps when it came to attributing blame to political parties (Bisgaard 2015).

Motivated reasoning may result in new information having a polarizing effect, driven by what Taber and Lodge (2006) characterize as motivated skepticism. Taber and Lodge show that political actors do not just seek out and uncritically accept information that fits with prior beliefs (a confirmation bias), but also actively refute information that is contrary to their prior beliefs (a disconfirmation bias). Motivated skepticism seems especially likely for elected officials relative to members of the general public. Taber and Lodge (2006) argue that actors with strong political beliefs and some measure of political sophistication will become more polarized as a result of engaging in information processing. More sophisticated political actors have greater knowledge to engage in disconfirmation with sources of information they disagree with, and greater motivation to disagree with a perspective at odds with their initial beliefs. Given that elected officials are relatively sophisticated, have generally fixed political beliefs, and have a high level of policy knowledge, we expect that the potential for motivated reasoning and resulting polarization to be high in our setting. In framing our hypothesis on the effects of interest group advocacy, we therefore consider whether the advocate and audience share ideological beliefs, with the expectation that liberal interest groups will be better able to discourage liberal politicians to use performance data to engage in responsibility attribution, but that conservatives will be motivated to refute

interest group criticisms, and may even become more committed to using performance data as a result.

- H2. Interest group statements about the quality of the data result in political polarization in the use of performance data for responsibility attribution.

In our experiment the comments of the advocate come from the teachers union, a group that traditionally has a closer ideological relationship with liberal rather than conservative political parties in Denmark and elsewhere (Moe 2011). This ideological alignment is clear in recurring debates over standardized testing of student academic performance, and accordingly we propose that in this context liberal party members are more receptive to the concerns expressed by the advocate about the performance data.

Approximately 1 year prior to this study, Denmark experienced a major work conflict between the teachers union and their employers, potentially resulting in more polarized attitudes toward the teachers union. Moreover, performance measurement in public education is a highly salient and contested issue in many countries, including in the Danish context examined here (Nielsen and Baekgaard 2015). This combination of politically contested performance measurement and differences in alignment between partisans and the advocate should allow us to detect motivated responses to advocacy if they exist (Slothuus and de Vreese 2010).

Research Design and Data

We examine our hypotheses in the context of Danish local government. Local elected officials oversee the funding and organization of local government services, which make up roughly half of all public spending in Denmark (Blom-Hansen and Heeager 2011). In particular, we focus on local elected officials' responses to performance data about the public schools they oversee.

In Denmark, the authority of bureaucratic leaders over their organizations and employees has gradually been expanded in exchange for increased performance oversight, with public education being one of the more striking examples (Nielsen 2014). Reinforcing this pattern, public managers have also been subject to more flexible and often incentive-based employment contracts, increasingly shifting risk to public managers both in terms of employment security and pay. Consequently, the assignment of credit and blame to bureaucratic leaders has become more important in political oversight and of greater consequence to public managers.

Danish schools are held accountable through a relatively sophisticated value-added performance

measurement system to assess school performance. These scores, made publicly available from the Ministry of Education, are widely disseminated: think-tanks have published performance rankings of schools and municipalities, which frequently make front-page news of both national and local newspapers. Local elected officials are ultimately charged with the management of public services, and are therefore subject to credit and blame for public service delivery from voters. By virtue of not just their official responsibilities but also deeper knowledge of local schools, local elected officials are also central actors in assigning responsibility for school performance, making them the appropriate political audience of performance data for our study.

Our experiment was conducted by embedding municipality-specific performance information treatments in a survey sent to all local elected officials in the 98 Danish municipalities during March and April 2014. With 1,016 valid responses, the response rate was 42%. Liberal party members were slightly more likely to respond to the survey (44%) than conservatives (40%), but conservatives still make up a marginally larger proportion of the sample (48% vs. 47%). The remaining 5% consists of councilors from locally based parties and independents. On other observable factors measured outside the survey (gender, member of Mayoral coalition, municipality characteristics) we found no significant differences between respondents and nonrespondents, though councilors from larger municipalities tended to respond slightly less frequently (see [Appendix Table A1](#)).

Experimental Treatments

To design our performance information treatment we used a publicly available ranking of Danish municipalities based on value-added performance, that is, student performance after taking into account a number of highly detailed register-based socioeconomic controls, most important of which are parents' education level, income, and immigrant status. This relatively high-quality value-added measure lends credibility to the performance data as a measure of academic performance, although political parties disagree about how important academic performance should be relative to other goals of public education. The high public salience of public education makes it ripe for battles in responsibility attribution and blaming ([Weaver 1986](#); [Hood 2011](#)). Considering the amount of media and public attention given to school rankings, local elected officials will be familiar with these rankings, and the performance data we use is no different from that which is already publicly available. In this respect, the officials we study should already have some knowledge of the performance of schools in their region, meaning that our treatments might be seen as a conservative test of the effects of performance data relative to

a context where no information is available. The particular performance data used here was gathered from a municipality election web site of the nation-wide newspaper *Berlingske*, showing both the exact value-added score and national ranking of each municipality.

Based on this ranking, we divided the municipalities into three performance groups of equal size: the lowest, middle, and highest performing third of the municipalities. This division allows us to examine the three basic types of performance signals (low, neutral, and high) while still maintaining sufficiently large subsamples. This division does not correspond to how the data is officially reported, but it does mirror the basic performance categories that the media and other actors typically focus on. Within each of the three performance groups respondents were randomly assigned to the control or treatment groups. Following the approach of [James \(2011\)](#) in providing performance data to citizens, we thus employed no deception in our experiment due to ethical concerns. As only the receipt of performance data—and not the performance score itself—was randomized, we effectively obtain three similar sub-experiments, that is, one for each performance group, totaling nine different conditions. All respondents were allocated to one of these conditions, and no other related treatments were included in the survey. The exact performance information treatment is shown in [table 1](#).

Alongside the performance information treatment we provided a second treatment group with an advocacy statement about the data. Given that the field we study is education, the treatment includes criticisms of student test scores that reflect actual scenarios. As shown in [table 1](#), we referred to actual policy statements made by the Danish teachers union. In Denmark, the teachers union more closely aligns with liberal than conservative parties, which is important to the test of H2. We chose not to personalize the advocate to limit the potential that personal characteristics (such as gender, or some unobserved familiarity with an actual union representative) might bias the results.

An identical introduction was given to all three groups to ensure that only the performance cue varies between the treatment and control groups and not the general cue to think about performance measurement as such. The two treatments were not fully crossed, that is, we did not include a pure advocacy frame treatment. This was partly a result of the study focus, but it was also a reflection of the limited population size that comes from elected officials and therefore the sample size that can be obtained with this case. As we discuss below, the findings clearly indicate that the advocacy frame effects arise in interaction with the performance cue. We found no significant differences between the treatment groups on observable characteristics, indicating a successful randomization.

Table 1. Experimental Treatments

Control Group	Performance Information Treatment	Performance Information + Data Advocacy Treatment
The quality of public services is frequently debated. For instance, there has been a distinct focus on measuring teaching quality in Danish public schools.	The quality of public services is frequently debated. For instance, there has been a distinct focus on measuring teaching quality in Danish public schools. An assessment from 2012 shows that the grade point average in the final public school exam in (name of municipality) ^a was placed in the (best, middle, worst) ^b third among the Danish municipalities when taking into account the social composition of the students.	The quality of public services is frequently debated. For instance, there has been a distinct focus on measuring teaching quality in Danish public schools. An assessment from 2012 shows that the grade point average in the final public school exam in (name of municipality) ^a was placed in the (best, middle, worst) ^b third among the Danish municipalities when taking into account the social composition of the students. The use of quality measurements has also been criticized, however. For instance, the Danish Teachers Union has argued that national test scores provide an incorrect picture of the quality of public schools.

^aThe name was inserted for the respondent's municipality.

^bBest, middle, or worst was inserted here depending on the performance of the respondent's municipality.

Measures

To measure our dependent variable, attribution of responsibility to bureaucratic leaders, we included the following two Likert-scaled items in the survey which were summed to produce an index ranging from 0 to 8: "The school principals do not have much influence on how well their schools are performing" (reversed) and "The individual school principals are to a great extent responsible for how well their students are doing academically." Our measures therefore reflect an assumption of causal responsibility, consistent with [Arceneaux \(2006\)](#). The second item explicitly refers to student performance, a key element of school performance, whereas the first item inquires about school performance more generally to also allow for a more explicit value-added conception of school performance. The two ordinal-scaled items are highly correlated ($\gamma = 0.53, p < .001$), suggesting they reflect aspects of the same underlying construct.² In order to avoid raising the respondents' awareness of the relevance of the experimental treatments to the dependent variable, three question batteries were placed between the treatment and response items.

Party affiliation was coded from the web sites of the municipalities. Our sample includes members from the eight political parties that also formed the 2014 multi-party system at the national level. The full sample

includes an additional 52 members (or 5%) from a number of different locally based parties. These are difficult to categorize in clear partisan terms, as they are often less ideologically motivated, and they were therefore left out of the interaction analyses ([Mortensen and Seeberg 2016](#)). As the experimental design consists of nine groups it is not feasible to examine partisanship patterns at the party level. Fortunately, the Danish multi-party system consists of relatively stable coalition patterns at the national level, which are mostly emulated at the local government level. The few exceptions typically consist of one smaller party, or just a few party members, allying with the other side to obtain the office of mayor or significant committee chairs; there are no indications that these deviations reflect alternative policy or ideologically motivated coalition patterns. The Social Liberal party is the only existing party that has traditionally been described as a centrist party, but since 1993 it has consistently been in coalition with left-wing parties, and is positioned clearly to the left of the Social Democrats on most noneconomic policy issues. Moreover, concerning the specific policy area we study, public education, the social liberals have historically been regarded as the closest political ally of teachers. Accordingly, we code members of the Unity List, the Socialist People's Party, Social Democrats, and Social Liberals as liberal parties; whereas the Danish People's Party, Conservatives, Liberals, and Liberal Alliance are all coded as conservative parties. Descriptive statistics are provided in the [Appendix Table A2](#).

² Analyzing the two items separately yields results that are highly similar to those reported here.

To take into account the clustered nature of the data, with local elected officials nested within municipalities, we follow Nielsen and Baekgaard's (2015) analysis of a similar experimental setup and estimate the treatment effects in a regression framework using municipality fixed effects and cluster-robust standard errors. Both fixed effects and random effects models would yield unbiased estimates, but the fixed effects approach is sometimes more efficient, using only within-municipality variation to compare differences between treated and nontreated respondents within the same municipality, thus reducing error variance resulting from cross-municipal differences between elected officials (Allison 2009). Accordingly, the inclusion of municipality fixed effects is used as a means of poststratification rather than post hoc covariate adjustment (Bowers 2011).³ Except from a preliminary analysis, we estimate separate models for each of the three performance groups (low, average, and high), as they receive different performance information treatments. Thus, treated respondents from high-performing municipalities are compared only to control respondents from the same municipalities.

Results

The result of our analysis for hypothesis 1 is provided in table 2.⁴ Table 2 shows that the general provision of performance data, regardless of performance signal displayed in the data, increases the willingness of the subjects to attribute bureaucratic leadership control over outcomes. Relative to the control group that did not receive performance data, the first column of table 2 shows that on average those exposed to performance information appear more willing to conclude that school principals control student performance. However, consistent with hypothesis 1, table 2 also shows that the relationship between the provision of

performance data and attribution of leadership control is concentrated among local elected officials whose schools are in the lowest category of performance. Treated respondents receiving a signal of average or high performance show no differences in responsibility attribution to leaders compared to their nontreated counterparts, suggesting that performance data does not in and of itself result in higher attribution of leadership responsibility for these groups.

Accordingly, this finding lends support to hypothesis 1, offering evidence of a negativity bias in how elected officials use performance data to make judgments about leadership responsibility. The statistical significance of the effect features a *p*-value slightly above the .05 level using a two-tailed test. As with any test relying on a limited population of elected officials, the sample size poses a challenge to obtaining a significant finding. However, our hypothesis is directional and would fall below the 0.05 threshold using a one-tailed test.⁵ The estimated effect corresponds to a 33% standard deviation (SD) change in responsibility attribution, suggesting that this is not a trivial effect. However, the size of the standard error also suggests that the effect varies somewhat across the subjects. On average, the estimated treatment effect moves respondents 7% on our responsibility attribution scale, though we note that only few observations are placed in the bottom half of scale range. To provide a more intuitive presentation of the findings from table 2, the treatment and control group averages are illustrated in figure 1.

Since the performance data represents actual performance, we did not randomize the content of the performance data (low, average, high) across all subjects, but only whether the performance data treatment was received. It is therefore important to consider whether the apparent negativity bias might instead be the result of inherent characteristics of the low-performance group, and thus not necessarily caused by the specific

Table 2. Impact of Performance Information Treatments on Responsibility Attribution, Split by Performance Group

	All	Low Performance	Average Performance	High Performance
Performance information	0.270 ⁺ (0.147)	0.544 ⁺ (0.283)	0.107 (0.259)	0.161 (0.223)
Constant	5.947 ^{**} (0.0726)	5.894 ^{**} (0.121)	5.910 ^{**} (0.135)	6.059 ^{**} (0.118)
N (politicians)	707	229	229	249
N (municipalities)	98	33	32	33

Note: Municipality fixed effects. Cluster-robust standard errors in parentheses.
⁺*p* < .10; ^{*}*p* < .05; ^{**}*p* < .01 (two-tailed test).

³ To ensure transparency, we also report the findings obtained from alternative tests and we note that simple *t*-tests of differences in means provides slightly lower *p* values across the treatment effect models reported below.

⁴ The Stata code used to construct the tables and figures is available from the authors on request.

⁵ The following alternative estimators yield only slightly different *p* values and fairly similar effect sizes: *t*-test of differences in means (*p* = .056); OLS (*p* = .056); OLS with cluster-robust standard errors (*p* = .13); municipality fixed effects (*p* = .020); municipality fixed effects with cluster-robust standard errors (*p* = .063); random effects (*p* = .055).

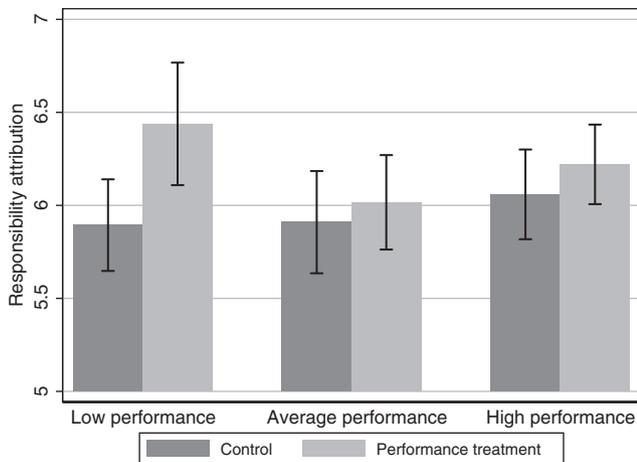


Figure 1. Impact of Performance Information Treatments on Responsibility Attribution. *Note:* Comparison of treatment and control group means, including 95% confidence intervals, based on the results reported in Table 2.

low performance signal.⁶ However, figure 1 illustrates that in the absence of performance data there is no significant or substantive difference between the three control groups in their baseline attribution of responsibility to leaders. A further indication of the similarities between the three groups is that a balance test comparing the low performance respondents to the rest of the sample finds no significant differences in the probability of belonging to the low performance group (Appendix Table A3). The only variable close to significance is membership of a liberal party ($p = .068$), but further analyses show no evidence of the liberal-conservative dummy moderating the impact of the performance information treatments.⁷ Finally, as shown below in the analysis of the political moderating effects for hypothesis 2, which control for party membership, we also find effects only for the low-performance group, lending further support to the negativity bias suggested in hypothesis 1. The balance of evidence therefore suggests that the provision of performance

data to elected officials responsible for low-performing municipalities encourages greater responsibility attribution to bureaucratic leaders.

Hypothesis 2 focuses on whether interest groups can alter how elected officials use performance data to engage in responsibility attribution. In contrast to the results shown in table 2, where the performance information treatment was associated with greater responsibility attribution, table 3 shows that when the negative advocacy statement is added to the performance information treatment, the difference with the control groups disappears for all performance groups, including the low performance group. This appears to indicate that elected officials have adjusted their patterns of responsibility attribution in a way consistent with the claims of the advocate, that is, becoming more reluctant to use performance data to attribute responsibility. However, examining the aggregate result of the advocacy treatment provides only a partial picture if subgroups of the treatment groups are responding to the data advocacy in opposite ways. This is essentially the logic behind hypothesis 2, which proposed that the advocacy treatment would generate a polarizing effect, causing liberals to become less likely to engage in responsibility attribution, but maintaining or increasing the likelihood that conservatives would use the data.

The interaction models in table 3 provide results for hypothesis 2. Again the average and high performance groups show no indications of moderating effects or of any significant marginal effects. As mentioned earlier these null findings for the average and high performance interactions corroborate hypothesis 1 about a negativity bias in the response to performance data: to the degree that elected officials appear to be engaged in and responsive to performance data for responsibility attribution, and attentive to arguments about that data, it is on the low end of the performance distribution.

The model for the low-performance group in table 3 shows clear signs of a moderating effect of party affiliation on the impact of the performance information-advocacy treatment. In line with hypothesis 2, treated members of liberal parties—those that are more ideologically aligned with the teachers union advocate—do not respond with the same increase in responsibility attribution to school principals as do members of conservative parties, as shown by the significant negative interaction term.⁸

Moreover, as illustrated by the marginal effects plot in figure 2, when liberal party members are provided the advocacy treatment they even end up attributing significantly less responsibility for outcomes to

6 Power tests based on the current sample characteristics show that the sample sizes for the average and high performance groups were sufficiently large to detect medium or large effects (effects above 25% of SD in the dependent variable). Accordingly, we cannot rule out that smaller true effects would go unnoticed with the current sample size. Although such smaller effects could still be consequential, they would be smaller than the treatment effect estimated for the low performance treatment.

7 We tested gender, age, education level, party membership, and membership of the mayoral coalition, and we also included four additional municipality-level variables: net public school expenditure per student, average school size in the municipality, municipality size, and tax base. As a robustness test we also ran a model that added a partisan control and an interaction effect between the performance information treatment and political party, but did not find a result for these additional variables, or any meaningful change to the effect of the performance data treatment moderated by party affiliation (results available from authors upon request).

⁸The interaction remains highly significant across all the regression estimators mentioned in footnote 5.

Table 3. Impact of Performance Information–Advocacy Treatments on Responsibility Attribution Moderated by Party Affiliation, Split by Performance Group

	All Groups	All Groups	Low Performance	Low Performance	Low Performance	Average Performance	Average Performance	Average Performance	High Performance	High Performance
Performance information–advocacy treatment	-0.170 (0.149)	0.262 (0.218)	-0.168 (0.295)	0.936* (0.408)	-0.0342 (0.259)	0.105 (0.343)	-0.293 (0.229)	-0.174 (0.342)		
Party dummy (Liberal = 1)		-0.747** (0.212)		-0.311 (0.409)		-1.082* (0.409)		-0.847** (0.273)		
Performance information–advocacy treatment × party		-0.669* (0.309)		-1.687** (0.597)		-0.347 (0.521)		-0.0161 (0.457)		
Constant	6.004** (0.0692)	6.359** (0.125)	5.990** (0.128)	6.145** (0.256)	5.890** (0.125)	6.435** (0.213)	6.121** (0.109)	6.481** (0.166)		
N (politicians)	667	628	231	218	213	205	223	205		
N (municipalities)	98	98	33	33	32	32	33	33		

Note: Municipality fixed effects. In line with the reported findings, using the same reduced number of observations in the main effect models as in the interaction models yields slightly smaller coefficients and similar standard errors. Cluster-robust standard errors in parentheses.
+ $p < .10$; * $p < .05$; ** $p < .01$ (two-tailed test).

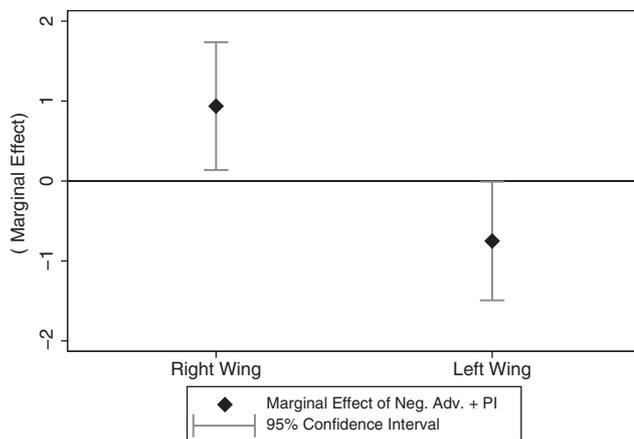


Figure 2. Marginal Effects of the Performance Information-Advocacy Treatment for Liberal and Conservative Parties (Low Performance Group).

bureaucratic leaders than does the similar liberal control group (the whole confidence interval is below 0, with a p value of .047 for the marginal effect). Here, liberal officials appear to be engaging in motivated reasoning, actively engaging in judgments at odds with the performance data if given a reason to do so by a trusted source.

Conservatives who received the performance information and advocacy treatments are significantly more likely to attribute responsibility to school principals in the low performance group relative to peers in the control group who received no treatments. Unlike their liberal peers, therefore, conservative officials do not appear to be influenced by the advocacy of union interest groups. Similar to conservative peers who received only the performance information treatment, conservatives continue to use performance data to engage in responsibility attribution for schools in the low performing category even if also provided the data advocacy treatment.

Consistent with hypothesis 2, we thus see clear signs of the performance information-advocacy treatment leading to greater polarization among political parties in terms of how they attribute responsibility to bureaucratic leaders, ranging from a significantly negative effect for liberals to a significantly positive effect for conservatives. Indeed the size of the difference in the estimated marginal effects for liberals and conservatives of the performance information-advocacy treatment is just over 1 SD of the dependent variable. This pattern is indicative of a pattern of motivated skepticism, under which actors initiate active reasoning for or against an advocacy message depending on their relation to the advocate (Taber and Lodge 2006). Accordingly, the presence of both performance data and interest group advocacy potentially becomes a mechanism for more politically polarized decision

processes rather than a force for establishing consensus about policy outcomes.

Limitations

Our study comes with a number of limitations that are largely a function of the research design choices made, and which point to avenues for future research. We sought to study actual elected officials with working knowledge over the bureaucrats we asked them to make judgments on. This makes their assessments more plausible and externally valid than evaluations made in hypothetical scenarios but it should be recognized that officials do not come to these judgments with a blank slate, and may have preexisting beliefs about the bureaucrats in question. However, the use of randomization means prior beliefs do not explain the difference between treatment and control groups, and the fact that elected officials have prior beliefs should dampen the size of the effects of the treatments, rather than generate a false positive. There are obvious limitations to studying a particular function in one country, but education is an almost universally important public function, and the business of holding bureaucrats accountable is a basic task of elected officials, even if the tools available to them to apply punishments and rewards vary by setting. Future work could examine if similar patterns of responsibility attributions hold in other settings and policy areas, consider a wider array of experimental treatments beyond those employed here, or seek to link assignments of responsibility to actual decisions.

Conclusion

This article expands the study of responsibility attribution, examining how elected officials attribute causal responsibility to bureaucrats for outcomes if faced with two policy prompts: the provision of evidence on the outcomes of public services, and interest group arguments that seek to influence the interpretation of that data. To our knowledge this is the first study that examines responsibility attribution to bureaucrats from the perspective of elected officials while using an experimental approach to assess causal connections. Overall, we find that elected officials become more willing to assign responsibility to bureaucrats if given performance data, but only in cases where data suggests low performance.

The results contribute to what some have termed behavioral public administration (Grimmelikhuijsen et al. forthcoming) in that it addresses how cognitive processes and biases operate in the construction of public accountability. Indeed, the results also reinforce recent evidence that shows that comparative presentations of performance data is especially powerful in

triggering cognitive processes of evaluating public services (Barrows et al. 2016; Olsen forthcoming).

The findings also offer implications for governance and research on political processes of responsibility attribution. One is that the general increase in performance data available to elected officials encourages their attribution of causal responsibility to bureaucrats. This does not mean that performance data necessarily generates better judgments: responsibility attribution research has shown that biased beliefs are no barrier to engaging in judgments (Hobolt et al. 2013; Tilley and Hobolt 2011). Indeed, studies of decision making suggest that information might lead to increased confidence even if it does not improve, or objectively reduces, the quality of decisions (Hall, Ariss, and Todorov 2007; Tsai, Klayman, and Hastie 2008).

While voter confidence in objectively poor attribution judgments may be an inevitable part of democracy, systems of representative government depend upon elected officials engaging in more careful assessments of the factors that result in public outcomes. Simply because politicians feel emboldened by the receipt of performance data to make responsibility attribution judgments, it does not follow that such judgments are always justified. Performance data in and of itself does not describe the causes of that performance, or point to the actors responsible for organizational outcomes (Moynihan 2008). Put another way, performance data may provide more information, and consequently more confidence in judgments made, without necessarily providing the causal clarity to justify that confidence.

Another implication of our analysis is that elected officials are most attentive to, and willing to engage in responsibility attribution for, low performance. While there have not been definitive studies on whether the negativity bias is more prominent in public organizations than private organizations, the empirical patterns on performance information use by political principals certainly fit with patterns of asymmetric attention to negative information generally (Kahneman and Tversky 1979), and blame-avoidance activities in public settings in particular (Hood 2011; Weaver 1987). The blame-avoidance literature is premised on the notion that negativity bias is a prominent force in the consideration of how responsibility—and by extension blame—is attributed, and we offer evidence on the underlying cognitive processes among politicians. As the evidence accumulates, it is hard not to conclude that performance management often becomes a negativity game that public managers are forced to play.

Our research also provides a window into how performance management relates to public sector innovation. Although one of the logics for performance data is to encourage bureaucratic innovation that fosters better outcomes (Heckman et al. 2011), bureaucrats

may avoid risks if they have reason to fear a disproportionate punishment from failure. A related question for future research is how the risks from responsibility attribution may vary between individual bureaucrats and their organization. Nielsen and Baekgaard (2015) use a similar experiment to show that elected officials were more inclined to provide organizational resources to schools that were labeled as performing poorly (see also George et al. forthcoming). But even as their organizations might benefit budgetarily from low performance scores, our analysis suggests that individual leaders may be put at risk. As bureaucrats come to understand these dynamics, their incentives are not to foster high performance, but to avoid negative low performance scores.

We also find that data advocacy can alter how performance data influences attribution processes, with interest groups most effective in altering the responsibility attribution patterns of politicians that share their ideological views. By contrast, elected officials not aligned with the interest group were unreceptive to its arguments. While our results demonstrate the potential effect of interest groups engaging in data advocacy, it also shows the limits of such efforts when made across the partisan divide. Practical guidelines and training on performance management emphasize technical skills, but neglect the ways by which politics infuses choices about both the selection of performance measures and their interpretation (Hatry 2006; or Moynihan (2013) on performance reforms of the George W. Bush administration). Practical guidelines are necessary, but just as human resource professionals are sometimes trained to be aware of their implicit biases, policymakers and bureaucrats should be made aware of how negativity biases and ideological preferences affect how they and others use performance data.

Part of the appeal of performance data for government is the promise to offer an objective and neutral account of public sector outcomes. But our findings suggest that the cumulative effect of performance data is to encourage polarization among elected officials. Prior evidence shows that conservatives are generally more willing than liberals to attribute responsibility to bureaucratic leaders (Nielsen and Moynihan forthcoming), and we show that data advocacy increases the distance between partisans. This finding arises even under relatively simple conditions of politicians examining a single performance metric. In reality, the ambiguity inherent in performance data arises not just from competing beliefs about how to interpret measures, but also on which measures to pay attention to in conditions where many are vying for attention (Moynihan 2008). As competing measures contradict one another, political principals may become less definitive in their judgments (Moynihan 2015). Our work extends these

observations about the relationship between advocacy and performance data by also considering how ideology affects the credibility of the advocate to the intended audience member. The results imply that ideology will play an even greater role in shaping the interpretation of data in more realistic conditions of multiple competing measures; conversely, ideology may be less important in conditions where there is relatively less polarization around a policy issue, or where performance data points overwhelmingly in a single direction. As elected officials examine data, they bring their own assumptions and beliefs to the process,

which colors how they use it. But they also work in an environment where interest groups will offer explanations on the meaning of data that will serve to further polarize how elected officials make judgments. Liberal and conservative elected officials arrived at significantly different positions in making responsibility attribution judgments as a result of being exposed to data advocacy. Whereas our study only included a liberal advocate, in real life conservative interest groups will try with equal effort to influence politicians, potentially resulting in even greater polarization than we have shown here.

Appendix

Table A1. Balance Test of Differences between Respondents and Nonrespondents

	Full Population	Only Liberals and Conservatives
Liberal versus conservative (liberal = 1)		0.231** (0.0870)
Gender (male = 1)	0.139 (0.0910)	0.121 (0.0933)
Member of mayor coalition	0.0611 (0.0833)	0.118 (0.0973)
Tax base	0.000497 (0.00142)	0.000173 (0.00149)
Net public school expenditure per student	0.00174 (0.00563)	0.00233 (0.00590)
Average school size in municipality	0.000219 (0.000174)	0.000206 (0.000178)
Municipality size (inhabitants)	-0.000830+ (0.000491)	-0.000961+ (0.000496)
Constant	-0.699 (0.442)	-0.796+ (0.459)
N (politicians)	2,417	2,303
N (municipalities)	98	98

Note: Standard errors in parentheses. Entries are random effects logistic regression coefficients predicting the likelihood of having responded to the survey.

+ $p < .10$; * $p < .05$; ** $p < .01$ (two-tailed test).

Table A2. Descriptive Statistics

Variable	Obs	Mean	SD	Min	Max
Leadership attribution	1,016	6.022	1.626	0	8
Liberal versus conservative (liberal = 1)	944	0.485	0.500	0	1
Member of mayor coalition	944	0.569	0.495	0	1
Gender (male = 1)	944	0.718	.450	0	1
Age (year of birth)	944	1,961	11.58	1,937	1,995
Length of education in years	944	5.019	2.383	0	8

Table A3. Balance Test of Differences between the Low-Performance Group and the Other Groups

	Low Performance
Proportion of males	-0.908 (1.697)
Average age (year of birth)	-0.0888 (0.0651)
Average length of education in years	0.295 (0.307)
Proportion of liberals [versus conservatives]	2.548+ (1.398)
Proportion members of mayor coalition	1.951 (1.370)
Tax base	-0.0132 (0.00971)
Net public school expenditure per student	0.0109 (0.0289)
Average school size in municipality	0.000720 (0.000909)
Municipality size (inhabitants)	-0.00165 (0.00399)
Constant	171.7 (127.4)
N (municipalities)	98

Note: Standard errors in parentheses. Entries are random effects logistic regression coefficients predicting the likelihood of being in the low performance group.

+ $p < .10$; * $p < .05$; ** $p < .01$ (two-tailed test).

References

- Andersen, Simon Calmar, and Donald P. Moynihan. 2016. How leaders respond to diversity: The moderating role of organizational culture on performance information use. *Journal of Public Administration Research and Theory* 26:448–60.
- Arceneaux, Kevin. 2006. The federal face of voting: Are elected officials held accountable for the functions relevant to their office? *Political Psychology* 27:731–54.
- Allison, Paul D. 2009. *Fixed effects regression models*. Thousand Oaks, CA: Sage.
- Baekgaard, Martin, and Søren Serritzlew. 2016. Interpreting performance information: Motivated reasoning or unbiased comprehension. *Public Administration Review* 76:73–82.
- Barrows, Samuel, Michael Henderson, Paul E. Peterson, and Martin R. West. 2016. Relative Performance information and perceptions of public service quality: Evidence from American School Districts. *Journal of Public Administration Research and Theory* 26:571–83.
- Berry, Christopher R., and William G. Howell. 2007. Accountability and local elections: Rethinking retrospective voting. *Journal of Politics* 69:844–58.
- Bisgaard, Martin. 2015. Bias will find a way: Economic perceptions, attributions of blame, and partisan-motivated reasoning during crisis. *The Journal of Politics* 77:849–60.
- Blom-Hansen, Jens, and Anne Heeager. 2011. *Denmark: Between local democracy and implementing agency of the welfare state?* In *The Oxford handbook of local and regional democracy in Europe*, eds. John Loughlin, Frank Hendriks, and Anders Lidström, 221–41. Oxford, UK: Oxford Univ. Press.
- Bonica, Adam. 2013. Ideology and interests in the political marketplace. *American Journal of Political Science* 57:294–311.
- Bowers, Jake. 2011. Making effects manifest in randomized experiments. In *Cambridge Handbook of Experimental Political Science*, eds. James N. Druckman, Donald P. Green, James H. Kuklinski, and Arthur Lupia, 459–80. Cambridge: Cambridge Univ. Press.
- Boyne, George, Oliver James, Peter John, and Nicolai Petrovsky. 2009. Democracy and government performance: Holding incumbents accountable in English local governments. *Journal of Politics* 71:1273–84.
- Boyne, George A., Oliver James, Peter John, and Nicolai Petrovsky. 2010. Does public service performance affect top management turnover? *Journal of Public Administration Research and Theory* 20:261–79.
- Carpenter, Daniel P., and George A. Krause. 2012. Reputation and public administration. *Public Administration Review* 72:26–32.
- Charbonneau, Etienne, and François Bellavance. 2012. Blame avoidance in public reporting. *Public Performance & Management Review* 35:399–421.
- Cohen, Geoffrey L. 2003. Party over policy: The dominating impact of group influence on political beliefs. *Journal of Personality and Social Psychology* 85:808–822.
- Craig, Steven G., Scott Imberman and Adam Perdue. 2015. Do administrators respond to their accountability ratings? The response of school budgets to accountability grades. *Economics of Education Review* 49:55–68.
- Dixit, Avinash. 2002. Incentives and organizations in the public sector: An interpretative review. *Journal of Human Resources* 37:696–727.
- George, Bert, Sebastian Desmidt, Poul A. Nielsen, and Martin Baekgaard. forthcoming. Rational planning and politicians' preferences for spending and reform: Replication and extension of a survey experiment. *Public Management Review*. doi:10.1080/14719037.2016.1210905
- Gilad, Sharon, Moshe Maor, and Pazit Ben-Nun Bloom. 2013. Organizational reputation, the content of public allegations, and regulatory communication. *Journal of Public Administration Research and Theory* 23:581–608.
- Grimmelikhuijsen, Stephan G., Sebastian Jilke, Asmus Leth Olsen, and Lars G. Tummers. Forthcoming. Behavioral public administration: Combining insights from public administration and psychology. *Public Administration Review*. doi:10.1111/puar.12609
- Hall, Crystal C., Lynn Ariss, and Alexander Todorov. 2007. The illusion of knowledge: When more information reduces accuracy and increases confidence. *Organizational Behavior and Human Decision Processes* 103:277–90.
- Hamilton, V. Lee. 1978. Who is responsible? Toward a social psychology of responsibility attribution. *Social Psychology* 41:316–28.
- Hatry, Harry P. 2006. *Performance measurement: Getting results*. Washington, DC: The Urban Institute.
- Heckman, James J., Carolyn Heinrich, Pascal Courty, Gerald Marschke, and Jeffery Smith (eds.). 2011. *The performance of performance standards*. Kalamazoo, MI: Upjohn Institute.
- Hobolt, Sara B., James Tilley, and Jill Wittrock. 2013. Listening to the government: How information shapes responsibility attributions. *Political Behavior* 35:153–74.
- Holbein, John. 2016. Left behind? Citizen responsiveness to government performance information. *American Political Science Review* 110:353–68.
- Hood, Christopher. 2011. *The blame game: Spin, bureaucracy, and self-preservation in government*. Princeton, NJ: Princeton Univ. Press.
- James, Oliver. 2011. Performance measures and democracy: Information effects on citizens in field and laboratory experiments. *Journal of Public Administration Research and Theory* 21:399–418.
- James, Oliver and Gregg Van Ryzin. 2016. Motivated reasoning about public performance: An experimental study of how citizens judge the affordable care act. *Journal of Public Administration Research and Theory*. doi: 10.1093/jopart/muw049

- James, Oliver, Sebastian Jilke, Carolyn Petersen and Steven Van de Walle. 2016. Citizens' blame of politicians for public service failure: Experimental evidence about blame reduction through delegation and contracting. *Public Administration Review* 76:83–93.
- James, Oliver, and Peter John. 2007. Public management at the ballot box: Performance information and electoral support for incumbent English local governments. *Journal of Public Administration Research and Theory* 17:567–80.
- James, Oliver, and Alice Moseley. 2014. Does performance information about public services affect citizens' perceptions, satisfaction, and voice behaviour? Field experiments with absolute and relative performance information. *Public Administration* 92:493–511.
- Kahneman, Daniel, and Amos Tversky. 1979. Prospect theory: An analysis of decision under risk. *Econometrica* 47:263–92.
- Kogan, Vladimir, Stéphane Lavertu, and Zachary Peskowitz. 2016. Performance federalism and local democracy: Theory and evidence from school tax referenda. *American Journal of Political Science*. 60:418–35.
- Liang, Jiaqi, and Laura Langbein. 2015. Performance management, high-powered incentives, and environmental policies in China. *International Public Management Journal* 18:346–85.
- Lau, Richard R. 1982. Negativity in political perception. *Political Behavior* 4:353–77.
- Mahoney, Christine, and Frank Baumgartner. 2008. Converging perspectives on interest group research in Europe and America. *West European Politics* 31:1253–3.
- Marvel, John D. 2016. Unconscious bias in citizens' evaluations of public sector performance. *Journal of Public Administration Research and Theory* 26:143–58.
- Malhotra, Neil, and Alexander G. Kuo. 2008. Attributing blame: The public's response to Hurricane Katrina. *Journal of Politics* 70:120–35.
- Meier, Kenneth J., Nathan Favero, and Ling Zhu. 2015. Performance gaps and managerial decisions: A Bayesian decision theory of managerial action. *Journal of Public Administration Research and Theory* 25:1221–46.
- Moe, Terry M. 2011. *Special interest: Teachers unions and America's public schools*. Washington, DC: Brookings Institution Press.
- Mortensen, Peter B. 2012. 'It's the central governments fault': A study of elected officials use of blame shifting rhetoric. *Governance* 25:439–61.
- Mortensen, Peter B., and Henrik B. Seeberg. 2016. Why are some policy agendas larger than others? *Policy Studies Journal* 44:156–75.
- Moynihan, Donald P. 2008. *The dynamics of performance management: Constructing information and reform*. Washington, DC: Georgetown University Press.
- . 2012. Extra-Network organizational reputation and blame avoidance in networks: The Hurricane Katrina example. *Governance* 25:567–88.
- . 2013. Advancing the empirical study of performance management: What we learned from the Program Assessment Rating Tool. *American Review of Public Administration* 43:497–515.
- . 2015. Uncovering the circumstances of performance information use: Findings from an experiment. *Public Performance and Management Review* 39:33–57.
- Nielsen, Poul A. 2014. Performance management, managerial authority, and public service performance. *Journal of Public Administration Research and Theory* 24:431–58.
- Nielsen, Poul A., and Martin Baekgaard. 2015. Performance information, blame avoidance, and politicians' attitudes to spending and reform: Evidence from an experiment. *Journal of Public Administration Research and Theory* 25:545–69.
- Nielsen, Poul A., and Donald P. Moynihan. Forthcoming. Romanticizing bureaucratic leadership? The politics of how elected officials attribute responsibility for performance. *Governance*.
- Olsen, Asmus Leth. 2015a. Citizen (dis)satisfaction: An experimental equivalence framing study. *Public Administration Review* 75:469–78.
- . 2015b. Negative performance information causes asymmetrical evaluations and elicits strong responsibility attributions. *Paper presented at the annual meeting of the American Political Science Association*, San Francisco, CA.
- . Forthcoming. Compared to what? Experimental evidence on social and historical reference points in performance evaluation. *Journal of Public Administration Research and Theory*.
- Rozin, Paul, and Edward B. Royzman. 2001. Negativity bias, negativity dominance, and contagion. *Personality and Social Psychology Review* 5:296–320.
- Rudolph, Thomas. 2003. Who's responsible for the economy? The formation and consequences of responsibility attributions. *American Journal of Political Science* 47:698–713.
- Skowronski, John J., and Donald E. Carlston. 1989. Negativity and extremity biases in impression formation: A review of explanations. *Psychological Bulletin* 105:131–42.
- Slothuus, Rune, and Claes H. de Vreese. 2010. Political parties, motivated reasoning, and issue framing effects. *The Journal of Politics* 72:630–645.
- Sniezek, Janet A. 1992. Groups under uncertainty: An examination of confidence in group decision making. *Organizational Behavior and Human Decision Processes* 52:124–55.
- Taber, Charles S., and Milton Lodge. 2006. Motivated skepticism in the evaluation of political beliefs. *American Journal of Political Science* 50:755–69.
- Tilley, James, and Sara B. Hobolt. 2011. Is the government to blame? An experimental test of how partisanship shapes perceptions of performance and responsibility. *The Journal of Politics* 73:316–330.
- Tsai, Claire I., Joshua Klayman, and Reid Hastie. 2008. Effects of amount of information on judgment accuracy and confidence. *Organizational Behavior and Human Decision Processes* 107:97–105.
- Weaver, R. Kent. 1986. The politics of blame avoidance. *Journal of Public Policy* 6:371–98.
- . 1987. *The politics of blame avoidance*. Cambridge: Cambridge University Press.