NBER WORKING PAPER SERIES

ELECTORAL INCENTIVES AND GOVERNMENT TRANSPARENCY: EVIDENCE FROM FREEDOM OF INFORMATION REQUESTS

Raymond Fisman
Aron Malatinszky
Eyub Yegen

Working Paper 35248
http://www.nber.org/papers/w35248

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
May 2026

We are grateful to Elliott Ash and Daniel Berliner for their very helpful comments. We also thank seminar participants at Harvard University, University of Chicago, University of Toronto, Hong Kong University of Science and Technology, and the Global Integrity Workshop for invaluable feedback. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Electoral Incentives and Government Transparency: Evidence from Freedom of Information Requests
Raymond Fisman, Aron Malatinszky, and Eyub Yegen
NBER Working Paper No. 35248
May 2026
JEL No. D72, D73, H70

ABSTRACT

Transparency is a basic requirement of government accountability, allowing citizens to access the information necessary to monitor, reward, and punish politicians and bureaucrats. Yet transparency inevitably involves discretion by the same government that is being evaluated. In this paper, we show that reelection incentives affect transparency, by comparing the responsiveness of freedom-of-information requests made via the online platform MuckRock just before state and municipal elections to those filed elsewhere at the same time. Since identical requests are often made simultaneously to many agencies across the U.S., we may assess whether an agency rejects a request while holding constant the timing, submitter, and request content, varying only whether the request is filed shortly before an election. We first document that FOIA rejection rates are positively correlated with a standard measure of state-level corruption. Our main analyses then show that rejections and delays of FOIA requests in state and municipal agencies are higher ahead of (state and city) elections in high-corruption settings. Overall, our results suggest that discretion may be used to limit access to information relevant to political accountability.

Raymond Fisman
Boston University
Department of Economics
and NBER
rfisman@bu.edu

Eyub Yegen
Hong Kong University of Science
and Technology (HKUST)
Department of Finance
yegen@ust.hk

Aron Malatinszky
Boston University
aronm@bu.edu

An online appendix is available at http://www.nber.org/data-appendix/w35248

1 Introduction

Transparency, it has long been argued, is a basic principle upon which government accountability rests (Hood, 2006). It has the potential to limit rent-seeking (Fisman et al., 2020), improve bureaucratic performance (Honig et al., 2023), and bolster the legitimacy of government (Grimmelikhuijsen et al., 2013). Information on government performance may further help citizens assess whether to reelect a politician or to vote them out of office (Malik, 2020). These potential benefits from government openness and accountability led to the passage of the Freedom of Information Act (FOIA) in 1966, as well as so-called "sunshine laws" at the state level, modeled on the federal legislation. Collectively, these laws proffer the right to request access to U.S. government agency records at the municipal, state, and federal levels.[1]

Very often, however, transparency requires the cooperation of precisely the same officials who will be subjected to greater scrutiny as a result. This may help to explain the weakness of some state-level FOIA statutes, and thus the substantial variation in FOIA stringency across states (Cordis and Warren, 2014). For example, as of 2024, eight states required a response to most requests within three days, while in twelve states there was no prescribed time; one comparison of fees charged to fill requests found that they averaged $2 in Washington state, and as high as $431 in Idaho.[2] FOIA statutes also leave enormous room for discretion in delaying fulfillment of a request for sensitive material or denying the request outright (Kreimer, 2008). This discretion can lead to a potentially very wide gap between the stringency of written laws and their enforcement in practice (Hallward-Driemeier and Pritchett, 2015).[3] As the *Washington Post* put it in a 2021 article on police accountability, "[n]ationwide...exemptions are carved into state public records laws, empowering police departments to deny the public access to vast amounts of information."[4] A small-scale audit by a collective of Virginia media outlets found that FOIA requests on topics – such as felony data or government salaries – that are clearly covered under the state's FOIA statutes were nonetheless routinely ignored by government officials.[5]

[1]Globally, 119 countries have freedom of information laws according to freedominfo.org, accessed March 31, 2026. For simplicity and to align with past work, we refer throughout to rules at both the state and federal level as FOIA laws, although it is more accurate to use FOIA to describe rules governing federal agencies and freedom of information law (FOIL) to describe state-specific statutes.

[2]https://www.muckrock.com/news/archives/2023/dec/20/muckrock-survey-of-foia-fees-points-to-uneven-picture-across-the-us/, last accessed March 31, 2026.

[3]Globally, there may be even wider variability in enforcement of FOI laws. See, e.g., Ackerman and Sandoval-Ballesteros (2006).

[4]"Public records laws shield police from scrutiny - and accountability," Nate Jones, *Washington Post*, July 30, 2021.

[5]"Many Virginia officials ignore state sunshine law," *The Daily Press*, November 28,

Discretion can potentially lead to less transparency precisely when it is most valuable to citizens, because scrutiny may be least desired at such times by those responding to requests. Most obviously, this may be the case in advance of elections, when officials may wish to avoid the release of potentially damaging information. There are anecdotal instances of this type of politically-motivated non-responsiveness to FOIA requests. For example, under mayor Rahm Emanuel, Chicago's police department rejected then delayed via litigation responses to FOIA requests for politically contentious material related to the shooting death of 17-year-old Laquan McDonald. The initial request was made in January 2015, a month before the Mayoral election began on February 24. The requested materials were finally released in November after a lengthy court case. It subsequently became clear that the forcefulness with which the FOIA request was rebuffed reflected a concerted effort by the mayor's office to suppress details of the case, at least until after the closely contested mayoral race. (In Section 2 we provide more detail on this case as well as a wider set of illustrative examples.)

Our primary objective in this paper is to move beyond anecdotes to evaluate more generally the extent to which FOIA responsiveness is affected by such electoral incentives.

Specifically, we examine how discretion may be differentially utilized before versus after municipal- and state-level elections, using a "revealed transparency" measure that we produce to capture government openness and accountability. We build this measure from an initial sample of 99,396 FOIA requests filed between 2010 and 2024 via MuckRock.com, a website that facilitates FOIA request submission. Details on these requests are available via MuckRock, including the time to receive a response and whether the request was rejected (generally coded by MuckRock itself). The MuckRock Application Programming Interface (API) makes it relatively straightforward to make bulk FOIA requests across a range of agencies, a feature utilized by many users. This allows us to take a "matched group" approach to our empirical analyses, comparing the outcomes of identical FOIA requests made by the same submitter to comparable agencies, but done across various jurisdictions. Crucially for our purposes, some of these jurisdictions have upcoming elections while others do not, owing to differences across municipalities and states in election cycles. We may thus capture the role of electoral timing, holding other aspects of the request constant.

Since we view our revealed transparency measure as a contribution in itself, we begin by showing the basic cross-state correlates of it. The most natural dimension to explore is state-level government malfeasance, because corrupted officials may desire less openness, and also because lack of transparency may facilitate misconduct. We show that the likelihood of FOIA request rejection in general is highly correlated with state-level corruption, as captured by

2015. Available at https://www.dailypress.com/government/dp-nws-foia-project-state-responses-20151128-story.html, last accessed March 31, 2026.

federal convictions for the illegal use of public office (Campante and Do, 2014). This pattern is observed whether we consider the full set of 35,371 state and local FOIA requests that have a clear resolution, or the subset of 21,754 requests that are part of a "matched group," in which we focus on essentially identical requests filed in different jurisdictions by the same submitter. We see this as a basic reality check of the data, and also a noteworthy finding which indicates that states with more to hide tend to use discretion in FOIA responses to avoid disclosures. Interestingly, in our matched-group approach, FOIA rejection rates are essentially uncorrelated with an overall measure of government openness based on the formal rules governing FOIA requests which has been used by previous researchers (Cordis and Warren, 2014).

In our main analyses, we examine FOIA compliance around municipal and state elections. The effect of upcoming elections on FOIA compliance is theoretically ambiguous. Most obviously, officials may wish to avoid the release of embarrassing information before an election. Alternatively, a well-intentioned public servant may be particularly motivated to provide timely disclosures in advance of elections. Moreover, failure to comply with a request may itself be the source of campaign fodder.[6] The relationship is further complicated by the fact that those tasked with filling FOIA requests and those targeted by them may not be the same person and indeed may not even be politically aligned. In essence, our analysis assesses which of these effects dominate to lead to more or less pre-election transparency overall. We further argue that the ability to exercise discretion – as well as the likelihood that officials actually have something to hide – will be greater in more corrupt environments. This is suggested, for example, by the findings of Ferraz and Finan (2008) and also Banerjee et al. (2011), which show that information on politician performance has negative electoral consequences *only* in cases where there is underlying evidence of malfeasance or dereliction of duties. We thus conjecture that we will observe a greater positive impact of pre-election timing on non-disclosure in higher-corruption jurisdictions.

Turning to our results, we find no relationship between electoral timing and FOIA responsiveness on average. This is true whether we define failure as outright rejection or delaying sufficiently that any documents are only released post-election. This average effect masks considerable heterogeneity by state-level corruption – pre-election responses are relatively slow and/or more likely to result in rejection in high-corruption states, a result that holds for the full sample of requests as well as the subset of requests with clear resolution (i.e.,

[6] For example, an incumbent county attorney in Virginia was accused of demanding unreasonable fees (over $3,000) to fill a record request that was filed as an "October surprise" amid her closely contested reelection campaign. The refusal itself then became a matter of contention. See "Suit filed against Albemarle Prosecutor over FOIA Response," *Charlottesville Daily Progress*, October 29, 2015.

rejection or completion). To provide a sense of magnitudes, the base "failure" rate in our data is 36.78%. For states in the top quartile of the corruption rate distribution, our analyses imply an increase of 11.1 percentage points (30 percent relative to the base rate) in the failure rate for pre-election requests; for lower quartiles there is, if anything, a decrease in the pre-election failure rate.

Our work sits at the intersection of two broad areas in economics and political science, one that studies the causes and effects of government transparency and a second that focuses on the causes of electoral accountability.

Theoretically, the role of transparency in fostering democratic accountability has been well-studied by Ferejohn (1999) and Besley and Prat (2006) among others, the latter documenting an empirical connection between a strong media presence and government responsiveness. The potential of transparency to improve *bureaucratic* functioning is the focus of Honig et al. (2023), which shows that development projects result in better outcomes after the donor agency implements an access to information policy. However, their findings also emphasize the importance of compliance, which is our focus − these benefits are only seen in cases where there is an appeal process for rejected information requests. Our work, by contrast, focuses on electoral rather than bureaucratic accountability, and because of our data and setting, we are able to exploit within-group variation in accountability incentives, rather than cross-project variation in accountability.

Several recent studies highlight the increased scrutiny of politicians − and resulting change in behavior − that accompanies greater transparency, by experimentally manipulating information provision. In particular, Grossman et al. (2024) and Banerjee et al. (2011) show that providing citizens with "report cards" of legislator accomplishments (e.g., meeting with the electorate, spending development funds) leads to stronger performance of incumbents in subsequent elections in Uganda and Delhi, respectively. Banerjee et al. (2024) further shows that when Delhi legislators are informed that report cards will be provided to voters, it leads to performance improvements.[7] These interventions suggest that information is useful for voters in evaluating candidates, via information provided through civil society organizations. Our findings complement these results in emphasizing the impediments to accessing such information precisely when it is helpful to the public in holding politicians accountable.

Given the potentially conflicting interests of voters and public officials who may face greater scrutiny in a more transparent system, the question naturally arises of why some governments set up transparency initiatives while others do not. The endogenous choice of

[7]Malesky et al. (2012), however, finds that the benefits of such exposure may not carry over to scrutiny of politicians in non-democratic settings, based on a similar experimental intervention conducted in the Vietnam legislative assembly.

freedom of information laws across countries has been studied by Berliner (2014), building on the insight that parties in power will be more motivated to make future governments more transparent if there is a higher probability of political turnover. In line with this prediction, Berliner finds that the strength of political opposition and frequency of turnover both predict the passage of freedom of information laws. In a similar vein, Berliner and Erlich (2015) shows that within a single country − Mexico − the passage of access to information laws is predicted by the extent of political competition. Our focus is instead on the functioning of these laws once in place, and how their efficacy may be impacted by the broader institutional environment in which they operate.

Our study also contributes to the work that links transparency − especially via freedom of information laws − to corruption. Using cross-country data, Costa (2013) shows that the passage of FOI laws is associated with an *increase* in perceived corruption. Naturally, rather than reflecting an increase in actual corruption, this pattern could stem from increased awareness of corruption as a result of FOI-driven revelations. Cordis and Warren (2014) looks at the link between the passage of FOIA laws across states within the U.S. and corruption convictions, comparing states that pass relatively stringent laws with those that pass laxer rules. They find an intriguing non-monotonic effect of stricter FOIA laws − first convictions increase, and then decrease, a pattern they suggest may be reconciled with FOIA requests first serving to uncover corruption and then to deter it. Given this non-monotonicity, there is, unsurprisingly, a zero correlation between corruption convictions on average and the strength of FOIA regulations. Our findings are distinctive in our focus on legal compliance rather than laws as written (Hallward-Driemeier and Pritchett, 2015), and also in our approach to identification, which exploits quasi-experiments rather than cross-state or cross-country variation in written laws. In our case, we find that access to information (as captured by realized FOIA outcomes rather than legal statutes) is associated with lower corruption, and our findings also emphasize that corruption itself may act as a mediating factor in when and how freedom of information laws are deployed.

Methodologically, our paper is one of several studies that use either an audit or matched-grouping approach to measuring the extent of compliance with freedom of information laws. Lagunes and Pocasangre (2019) and Jenkins et al. (2020) both randomly vary the identity of the requester to show that, respectively, elite status in Mexico and political contributions in the U.S. have no effect on response rates. These well-identified papers provide helpful insights into FOI compliance, but are limited in the sense of building on just a few hundred requests. Furthermore, using “natural” requests, as we do, also opens the door to examining aspects of transparency for which it might be challenging to obtain human subjects approval.[8] Fi-

[8]For example, a set of requests we describe below made queries about the weapons arsenals of

nally, we also see our approach as offering a relatively straightforward method for scaling the study of FOIA compliance. Closer in methodology to our paper are Berliner et al. (2021) and Trautendorfer et al. (2023), which both analyze large numbers of requests. The first of these looks at the relationship between governing party vote share and responsiveness to over 450,000 freedom of information requests in Mexico, and concludes that transparency is affected by political circumstances. The latter paper explores how the tone and wording of requests affect responsiveness, using text analysis techniques. We see our findings as building on this earlier work in a couple of ways. First, because the particular platform we deploy is often used for information gathering at scale on the identical topic by the same submitter, it allows for a clear apples-to-apples comparison (our matched-group approach) across jurisdictions or circumstances more broadly.[9] Second, our main analysis focuses on responsiveness as a function of election timing, which is distinct from earlier work and to the best of our knowledge completely new to the literature.

Looking beyond the specific theme of freedom of information rules, we contribute to the nascent body of work that explores political control of the bureaucracy (Brierley et al., 2023). Previous work in this area considers the extent to which ostensibly independent executive agencies are politicized when they are led by elected officials, yielding variable results across different types of agencies (see, e.g., Hamel, 2025; Thompson, 2020). Our work helps to underscore that even in the absence of electoral pressures for appointed leaders of executive agencies, they may still face pressure from higher up in the political hierarchy. Existing empirical work has highlighted hierarchical control as a source of electoral cycles in bureaucratic hiring (e.g., Levitt, 1997; Enikolopov, 2014; Vlaicu and Whalley, 2016) and in dismissals of agency heads in response to publicized critical events (Camerlo and Pérez-Liñán, 2015). An older, more qualitatively oriented literature in political science focuses on the extent of bureaucratic control exercised by the executive at the national level in the U.S. (see, e.g., Wood and Waterman, 1991), and in particular government functions, as in the classic work on the link between administrative and political corruption by Wade (1982).[10]

individual police departments, which were often rejected on public safety grounds. An IRB might be wary of approving such controversial requests, and with good reason.

[9] Berliner et al. (2021) includes fixed effects by topic-agency-year which, while ruling out many concerns of omitted variables, leaves considerable residual within-group variation.

[10] It is important to note that it is not the case that bureaucrats should necessarily be fully insulated from political considerations – a point that has been made by, for example, Stephenson (2008).

2 Background and Data

2.1 Background on Political Interference in FOIA Requests

Before describing the data, we provide the reader with some further context for the phenomenon we study in this paper. There are two elements to this discussion. First, we discuss how elected officials can pressure ostensibly independent bureaucrats in general. Second, via several case studies, we emphasize the plausibility of elected officials exercising their authority to influence FOIA outcomes specifically.

The most important feature of local governance to emphasize is that governors and mayors hold formal authority over executive agencies through their power to appoint and dismiss agency heads – for instance, police chiefs are directly appointed by governors in almost every state (The Council of State Governments, 2021) and by mayors in most cities (de Benedictis-Kessner et al., 2025). This appointment power leads to two potentially complementary sources of influence. Most importantly, elected leaders may appoint officials on the basis of personal ties and loyalty (see, e.g., the classic work of Abney and Lauth, 1986). Furthermore, although limited by checks against executive powers, in many cases it is possible for elected leaders to remove agency officials at will, raising dismissal as another lever of influence.[11]

Moving from the general case to the specifics of political influence over FOIA responsiveness, we provide a series of case studies to emphasize the plausibility of the link we make in our empirical work. We provide brief summaries here, with more extended descriptions in Appendix F. We note that these examples do not necessarily pertain to corruption in the strictest sense of illegal use of office for private gain, which is what our corruption measure used in the empirical analysis captures. However, they generally reflect what voters may see as inappropriate use of public office for personal or electoral benefit. These ethical lapses – which are reflective of a broader definition of corruption – are plausibly correlated with the presence of illegal behaviors that are captured by the legal definition of corruption (see, e.g., Rose-Ackerman and Palifka, 2016 for more details).

- **Rahm Emanuel and the Laquan McDonald scandal** In January 2015, following the fatal shooting of Laquan McDonald by a Chicago police officer, journalists filed FOIA requests for investigative reports and dashcam footage. The Chicago Police

[11] For example, as of 2020, 30 states had either constitutional or administrative procedure act provisions that allowed the governor to remove agency officials. In 14 states, the governor had at-will removal power, i.e., no explanation or cause was required for dismissal. See `https://ballotpedia.org/Executive_control_of_agencies:_State_executive_removal_power_over_agency_officials`, last accessed March 31, 2026.

Department denied these requests, citing interference with an ongoing investigation, and the materials were withheld until litigation forced their release in November 2015. Subsequent disclosures, including over 3,000 internal emails, revealed that the mayor's office and city attorneys had coordinated efforts to prevent disclosure until after the closely contested mayoral election. In this case, obstruction occurred through direct involvement by the mayor's office, which sought to control the timing of disclosure for electoral advantage.

- **Buffalo Billion scandal** In July 2014, investigative journalist Jim Heaney requested records concerning Governor Andrew Cuomo's $1 billion Upstate economic development initiative, amid allegations of corruption in contract allocations. The administration responded with heavily redacted documents and subsequently resisted further compliance, arguing that disclosure might compromise ongoing investigations. In 2017, a court order compelled the governor's office to release schedules, emails, and related records. In this case, obstruction was carried out directly by the governor's office, which used claims of investigative privilege to delay compliance during an election period, thereby protecting the political interests of the executive.

- **Atlanta water bills** In March 2017, a local television station submitted FOIA requests to Atlanta's Department of Watershed Management for billing and usage records tied to properties owned by Mayor Kasim Reed and his brother, both of which were under scrutiny for delinquency or misuse. Internal communications later showed that the mayor's press secretary instructed staff to delay responses, be "as unhelpful as possible," and provide records in a confusing format. Compliance was achieved only after the threat of litigation. In this case, obstruction was enabled by the institutional hierarchy, whereby the department's communications office reported to a commissioner directly appointed by the mayor, creating a clear channel for political pressure to shape FOIA responses.

- **Control of Richmond FOIA office** In 2017, Richmond's FOIA officer, Connie Clay, reported overdue requests, including one concerning the mayor's salary, before being dismissed from her position. According to her whistleblower complaint, the Finance Director and Director of Strategic Communications, both appointed officials with authority over FOIA compliance, ignored or delayed requests and ultimately terminated her employment. Because these officials reported to the Chief Administrative Officer, who was himself appointed by the mayor with city council approval, the FOIA office was structurally embedded within a political chain of command. In this case, obstruction operated not only through discrete denials but also through institutional design, which allowed elected officials to exercise systemic control over the FOIA process.

In summary, these cases illustrate the potential influence that political leaders have over bureaucratic decision-making in general, and specifically with respect to decisions to release information in response to FOIA requests. A key feature of several of these stories is also the role of electoral concerns in these decisions.

2.2 Data

MuckRock facilitates the FOIA submission process by generating appropriate emails or submitting requests via agency web portals, and sending auto-reminders (in addition to any correspondence the submitter sends to an agency) until a request receives acknowledgment and/or response. MuckRock also provides extensive advice on how to manage submissions in order to maximize agency responsiveness and, when appropriate, provides indirect assistance in complying with local laws and/or raising the necessary funds to cover fulfillment fees. In particular, since several states require in-state filings, MuckRock offers to pair submitters with local volunteers who may file requests on their behalf (we return to this point below) and also provides crowdfunding options to submitters who need to raise money to cover fees.

In early September 2024, we scraped all requests made via MuckRock since its inception in 2010, a total of 99,396 requests.[12] In its first year, 219 requests were placed; request volume grew steadily for the next few years, and hit a rate of 10−11,000 during 2017-2020, before declining (one presumes because of COVID) to just over 8,000 in 2021. Of the 99,396 total, 31,586 (31.8%) were requests of federal agencies, which, given our interest in exploring cross-state and cross-city variation in responses, we exclude from our analysis. Of the remaining 67,810 requests, 47,254 (69.7%) are city-level requests and 20,556 (30.3%) are state-level requests. Of these, since we will be interested in the role of state-level corruption as a direct and as a mediating factor in explaining FOIA responsiveness, we further drop the small number of requests from the District of Columbia and U.S. territorial governments (400 and 11, respectively), yielding a sample of 67,399 requests.

MuckRock's "Essentials Team" codes the status of a given request according to fixed guidelines on the extent to which an agency has complied with the request, updating the

[12]This is a sufficiently modest total that we do not expect that MuckRock itself stands out as a notable source of FOIAs to officials tasked with filling requests. While it is difficult to obtain the total number of local FOIA requests, the Department of Justice reports on total FOIA filings at federal agencies. The total in fiscal year 2023 was 1,119,699 requests, as compared to 1,636 federal requests filed via MuckRock in 2023; the total number of FOIA requests filed at all levels via MuckRock that year was 9,862.

status over time as appropriate.[13]

There are two codes which indicate that a request has received a clear resolution − those that are marked by MuckRock as "completed" (29,384 of the state and local requests), and those that end in rejection (5,987 of the state and local requests). For codes not resolved in this binary manner, the additional possible outcomes include: withdrawn by the submitter (3,984), awaiting acknowledgment from the agency (2,402), submitted and awaiting response (12), being processed by the agency (2,271), awaiting appeal to agency (279), fix required by submitter (5,038), in litigation (104), payment required by submitter (1,621), request was only partially completed (1,012), and no responsive documents provided by the agency (15,305). These categories are generally ambiguous in their interpretation. While about 37.3% of the "awaiting acknowledgment" requests were relatively recent at the time the data were downloaded − from 2022-2024 − many were older, and apparently ignored despite monthly queries. Many of the other categories offer a range of explanations. For example, a fix required by submitter may be because clarifying details on the request are required, or because the responder explains that the request was directed to the wrong office or agency. Requests are generally withdrawn because they overlap with prior or concurrent requests, or may be merged with other requests. Partial completion refers to requests that, for example, do not elicit the full desired set of documents. There appear to be many reasons for this − sometimes difficulty in accessing materials that the agency directs the requester to look at; often the agency responds that the scale of the request is unreasonable and so provides only a partial set of documents. The largest "non-completion" category is "no responsive documents" (22.7 percent of all requests). This occurs most often because the agency reports that it does not have or is unable to locate the requested documents. According to MuckRock's guidance on its website, this response occurs with some frequency even when the requester knows for certain that the relevant record exists, and suggests that the response often reflects some combination of willful obstruction, laziness, and incompetence (technological or otherwise).[14]

MuckRock defines a request as closed if its status is coded as Completed, Partially Completed, Rejected, or No Responsive Documents; all other intermediate categories are potentially still in-progress. In most of our analyses, we will focus on the well-defined outcomes of rejection or completion. Given that ambiguous outcomes may also reflect official efforts to

[13]Based on correspondence with MuckRock staff. It is possible for requesters to override MuckRock's status assignment, though in practice this happens rarely.

[14]See https://www.muckrock.com/news/archives/2018/sep/06/foia-faq-nrd-wtf/; last accessed March 31, 2026. "No responsive documents" can also be an interim classification if the submitter requests follow-up from the petitioned agency, but we emphasize that it does *not* mean that the agency has not responded – indeed quite the opposite: a reply from the agency is required to receive this assignment.

stymie or limit information access, in presenting our main results, we will explore how other outcomes vary as a function of state-level attributes and timing around elections, and probe the sensitivity of our results to more inclusive definitions of failure.

Requests can be rejected for many reasons, and even criteria that would seem to be objectively defined end up involving substantial discretion. This is important for understanding why responsiveness may deviate from what one might expect based on laws as written. An agency may determine that a request involves confidential or sensitive information; or it may have failed to access the requested documents after undergoing a "reasonable" search, where the definition of reasonable is a matter of interpretation.

We provide an illustrative example of a grouping of 29 requests filed in early August 2015, which queried state-level Departments of Correction about their death penalty procedures.[15] 26 of these received a determination of Completed or Rejected. The 10 completed ones included detailed manuals, often dozens of pages in length. The most common reason for rejection was confidentiality of Department of Corrections records, or records related to execution specifically. Other rationales were also provided, however. The responder in Wyoming, for example, wrote that, "the Wyoming Department of Corrections does not release its Capitol [sic] Punishment policy in its entirety because it contains significant safety and security information that is protected under the Wyoming Public Records Act." Finally, the request was rejected in Tennessee because agencies in that state only respond to requests from Tennessee residents, and the filer was a resident of Massachusetts. (Note, however, that the response from Delaware did not mention the submitter's out-of-state residency despite a similar in-state requirement, and was rejected in that state on the basis of confidentiality concerns.) Throughout, we will be sure to account for out-of-state restrictions which, as noted earlier, are recognized by MuckRock as a common reason for rejection − laws in Arkansas, Delaware, Kentucky, Tennessee, and Virginia limit requests to state citizens, while Alabama has more ambiguous rules. (Yet requests made without proof of citizenship in all of these states are nonetheless sometimes approved.)

To be part of a "matched grouping" (as in the above case) we require at least two identical requests, which we define as having the same content requested by the same submitter. While it is possible to tailor the content of each message in a group, in practice these requests are always "batched" so that the full content of the request is also near-identical, with minor variations according to pertinent statutes, names, and terminology specific to

[15]The grouping may be found via title "Department of Corrections Death Penalty Procedures" submitted by Emily Hopkins during July 1, 2015 − September 30, 2015. Note that this search returns 30 requests, but one was for information from the Federal Bureau of Prisons rather than a state-level department.

each jurisdiction. To identify these groups in the data, we therefore match requests using cleaned submitter usernames and request titles.[16] For our main analysis, which includes only completed and rejected requests, there are 13,617 singleton requests, leaving 21,754 observations that are in matched groupings. Since we will be using these data to look at how state-level corruption is correlated with responsiveness, we focus on groupings for which there is cross-state variation, thus limiting our sample to groupings with at least two requests in different states. This restriction further reduces our sample to 12,603.

In Appendix B, we characterize the requests in our sample along a number of dimensions. We first provide in Appendix Figure B1 a histogram showing the number of requests per matched group. There is a considerable range in the scale of requests. 183 groupings include only two requests with a clear outcome (often part of a larger grouping that included other requests for which there was no clear resolution as of September 2024), and 86 groupings include over 30 requests with a clear resolution. We make two observations about these small- and large-group requests. First, we generate very similar results when we omit relatively small groupings (e.g., fewer than five requests per group; see Appendix Table A2). Second, larger groupings may introduce some within-group heterogeneity, most obviously in the petitioned agency. For example, identical requests are sometimes made to district attorneys' offices, public safety agencies, as well as city police departments; or individual state agencies as well as the governor's office. Thus, as a robustness check, in Section 4, we show that our results are not driven by any particular geography or agency.

We next show the types of agencies that are targeted with requests. In Appendix Figure B2, we show the distribution of requests by agency category. Police is, by a wide margin, the most common category, constituting 35.99% of all requests; a number of other areas related to criminal justice are also prominent. Turning to the types of submitters, we categorize individuals into three categories – journalists, researchers, and others – based on a combination of information obtained directly from their MuckRock profiles, as well as web searches (see Appendix D for details). We present these results in Appendix Figure B3, where we see that the dominant category is journalist. While it is possible that journalists are overrepresented among MuckRock submitters relative to the universe of FOIA requesters, they are also a

[16] We manually verify all matched groups in our elections analysis sample by randomly selecting at least five FOIA requests (or all requests if a group contains fewer observations) within each group. We confirm that identical titles filed by the same person correspond to FOIA requests with identical or near-identical content. In almost all instances, the underlying content is identical; in a very small number of cases, the subject matter is the same but the breakdown differs slightly (e.g., stating "all investments" rather than explicitly listing investment types). We identified only two cases in which a matched group had a request that differed from the others in what it was filed for, implying an error rate of about 0.2%. Excluding these rare cases leaves the results unchanged.

particularly relevant group for monitoring and disciplining elected officials, a point we return to in our results below.

Finally, to provide a sense of the main topics motivating the requests in our main analysis sample, we present a set of word clouds based on the request titles in our matched-group sample. To construct these word clouds, we first remove "junk" terms – specifically agency names and/or geographies from each title, since both are usually referenced in the request. Otherwise, the word clouds would be dominated by terms such as "New York," "California," and "Police," which reveal nothing about the content of the request (see Appendix E for more details on the word cloud construction, including the full list of junk terms). We aggregate all words to the username level, to avoid over-weighting the influence of large groups of requests and/or requesters who tend to focus on the same topic across requests.

Several notable patterns emerge. Focusing first on individual words (Appendix Figure E1), requests very often mention communications or emails, which is as expected, and consistent with the tenor of the examples provided earlier. Marginally smaller terms include "complaints" which very often relate to the use of force by law enforcement and "contract" which generally involves queries about outsourcing contracts. Turning to bigrams (Appendix Figure E2), we see frequent requests involving "use force," and "police report" or "incident report." For non-police requests, "incident report" and "autopsy report" are notably large – the latter a reflection of the large number of requests on this theme related to homicides. Note, though, that there is a very wide diversity of topics in both the single-word and bigram word clouds.[17]

Turning to the other data we use in our analysis, to measure state-level corruption, we utilize the measure from Campante and Do (2014), which captures the number of federal convictions per capita of state and local officials during 1976-2002. This period predates our data, so there is no mechanical link from FOIA requests to the emergence of corruption cases. We interpret the corruption measure as a broad indicator of government probity, and one that is a largely fixed characteristic over the time span we study.[18] Finally, we include the log of income per capita and population at the state-year level as basic controls, and stringency of a state's FOIA laws from the Open Government Guide.[19] The Reporters Committee for Freedom of the Press compiles and publishes these data and we follow Cordis and Warren

[17]We repeat the same analysis for the matched-group sample in Appendix Figures E3 and E4.

[18]In Appendix Figure B4, we show the relationship between average corruption levels from the first and second halves of the Campante and Do (2014) data. There is a strong positive correlation between the two ($\rho = 0.52$), which is broadly consistent with this measure serving essentially as a fixed state-level attribute.

[19]We obtain real income data at the state-year level from the Bureau of Economic Analysis for the years 2010 through 2023, and we use annualized Q3 real income data for 2024.

(2014) to calculate summary FOIA law scores.[20]

In our main analyses, we focus on FOIA responses around gubernatorial and mayoral elections. We determine the timing of municipal elections using the American Local Government Elections Database (ALGED), which provides local election returns from 1989 to 2021 in most U.S. cities, counties, and districts with populations over 50,000 (de Benedictis-Kessner et al., 2023). Of the 4,969 cities whose agencies appear in the MuckRock data, we observe municipal elections in 475 of them.[21] The ALGED data include only the month of election. To obtain the precise day, we provided ChatGPT with the spreadsheet of city-by-month observations, and prompted it to provide the day-of-the-month in each case. We were able to validate this approach in a 10 percent random sample of election dates, finding that ChatGPT's accuracy was 100 percent.[22] Gubernatorial elections are more straightforward, and were held at regular four-year intervals in all 50 states throughout our sample period, with the exception of the 2011 special election in West Virginia and the 2012 and 2021 recall elections in Wisconsin and California, respectively (see, e.g., Leip, 2024).

To construct the election-level dataset for the 4,269 completed or rejected matched-group filings in these 475 cities and 50 states holding 1,681 elections, we use the following steps. First, from our list of matched-sample groupings of FOIA requests, we select groups for which at least one request took place within 180 days prior to a gubernatorial or mayoral election in the state or city where the FOIA request was filed. Within each grouping, we define an indicator variable, *Election*, which denotes whether there is a gubernatorial or mayoral election in the request's state or city in the next 180 days. Since our data on elections extend through 2021, we restrict the sample to the set of FOIA requests that were submitted at least 180 days before the end of 2021. We will also consider alternate windows ranging from

[20] It is also natural to control for the size of government, e.g., via the number of government employees by state, though in practice such proxies vary little across states. The inclusion of this control in our analyses − whether it enters linearly or interacted with our pre-election measure − has no effect on our estimates.

[21] ALGED provides results from contests decided by direct election as well as those decided by city council appointment. We focus on direct elections by eliminating elections in which ten or fewer total votes were cast. For cases in which mayors are elected in multiple rounds of voting − whether through preliminary and general elections or through runoff elections held because no candidate secured a majority in an initial round − ALGED typically includes only returns from final rounds of voting. The database does not identify whether the 23 cases in which we observe multiple elections within a year in the same city are multiple rounds from the same contest or separate contests (occurring, for example, as a result of resignations). Therefore, to maintain consistency, in the 8 cases in which we observe data on multiple elections within a four-month span in the same city, we focus on final elections, and in the remaining 15 cases in which we observe data on multiple elections within a year in the same city, we focus on first elections.

[22] We used an automated search with OpenAI's `gpt-4o-search-preview` model, which is specialized for web searches in the Chat Completions API, on October 13, 2025.

120 to 360 days as robustness checks.

There are two ways that a pre-election request can "fail": it can take so long to fulfill that information is revealed only after the election, or it may be rejected outright.[23] Note that in most of our main analyses, we limit the sample to requests that end in completion or rejection, so that we focus on the delay of requests that nonetheless have a clear resolution. We also show that the results are robust to including "closed" requests with more ambiguous outcomes (i.e., requests coded as "partially completed" or "no responsive documents"), as well as maintaining the full set of matched group requests.

Rejection is straightforward to define, as outlined in the previous part of this section. It is more complicated to evaluate whether a request takes "too long." Consider, for example, a group of requests for which there is an election in one location that is three months in the future and a second in which there is an election five months in the future, as well as several requests in cities or states with no upcoming election. We cannot straightforwardly define a deadline as the actual time-to-election for the $Election = 1$ cases, since this does not allow for a clearly defined deadline for the non-election (control) observations. If we were to define the control deadline based on three or five months, it would create a mechanical correlation between deadline and treatment status. Since the median number of elections in a grouping is four, this issue arises for a majority of cases, and is especially relevant for the larger groups of requests.

Our approach is to apply a uniform deadline across all observations within a grouping. In each matched group, we define the longest election window as the maximum number of days between a FOIA request filing date and an election in the same locale, and we identify a request as having a *Delayed Response* if it is unresolved within the longest election window. In the example above, we would thus assign $\mathbb{1}(Delayed\,Response) = 1$ for all requests that are not resolved within five months. As a robustness check, we also look at the shortest pre-election window to define $\mathbb{1}(Delayed\,Response)$ (in the above example, three months).

We combine our rejection and failure measures to capture whether a request received a timely and satisfactory response. Specifically, we construct the variable *Failure*, which equals 1 if either $\mathbb{1}(Delayed\,Response)= 1$ or $\mathbb{1}(Rejected) = 1$. We examine $\mathbb{1}(Delayed\,Response)$ and $\mathbb{1}(Rejected)$ separately as outcomes as well.[24] The sample for this set of analyses com-

[23] A natural concern is that the volume of citizen queries and complaints may differ around elections (Dipoppa and Grossman, 2020). We address this issue in Section 4.4.2 below, where we show that, in practice, the monthly volume of FOIA requests received by jurisdictions in the six months prior to elections is statistically indistinguishable from the volume received during other months.

[24] Note that $\mathbb{1}(Delayed\,Response)$ and $\mathbb{1}(Rejected)$ are not mutually exclusive outcomes − if a request is rejected after the election, then both variables are equal to 1. We see this as an appropriate

prises 2,674 requests submitted to city agencies and 1,595 requests submitted to state agencies in 269 matched groupings.

2.2.1 Summary statistics

We next provide an overview of the FOIA response data, even the portions of it that we do not necessarily utilize in our analysis, to illustrate the breadth of requests that are made via the website, and also to consider potential selection issues that exist for MuckRock-generated requests. We begin in Panel A of Table 1 with summary statistics for the FOIA data; Panel B provides summary statistics for our corruption measure and for the control variables that vary at the state level. Across all agencies, among requests that received unambiguous resolution as either rejected or completed, 21% were rejected.[25] There is considerable heterogeneity across levels of government. The rejection rate for federal agency FOIAs is nearly double that of municipal agencies.

Table 1: Summary Statistics

This table summarizes the outcomes of FOIA requests made to different levels of government, along with state-level covariates. The full sample of requests includes all FOIA requests made via MuckRock, while our elections sample comprises all matched groups of requests submitted to state and city government agencies with variation in exposure to gubernatorial and mayoral elections used in our main analysis, which receive a clear resolution of "completed" or "rejected." For state-level variables, we use the per capita rate of corruption convictions during 1976-2002 from Campante and Do (2014), a measure proxying for the stringency of FOIA laws from the Open Government Guide, 2013 state population, and 2013 state income per capita (in 2017 dollars).

	Full sample			Elections sample		
	Mean	SD	Obs.	Mean	SD	Obs.
Panel A: FOIA Responses						
All Agencies						
Clear resolution	0.502	0.500	99,396	1.000	0.000	4,269
Among clearly resolved requests:						
Completed	0.793	0.405	49,892	0.824	0.381	4,269
Rejected	0.207	0.405	49,892	0.176	0.381	4,269
Filings per submitter	23.71	168.87	4,193	33.35	79.24	128

continues on the next page

coding to the extent that a rejection itself could be seen in a negative light, but in practice this distinction is irrelevant for our main analyses that focus on failure along either dimension.

[25] When we define rejection as either outright rejection or the agency reporting that "no relevant documents" are available, and include also partially completed requests as completed, the rejection rate is far higher – around 45% – simply because of the many requests that receive a "no relevant documents" response.

Table 1: Summary Statistics (cont.)

	Full sample			Elections sample		
	Mean	SD	Obs.	Mean	SD	Obs.
Federal Agencies						
Clear resolution	0.453	0.498	31,586			
Among clearly resolved requests:						
Completed	0.703	0.457	14,303			
Rejected	0.297	0.457	14,303			
Filings per submitter	15.64	111.17	2,020			
State Agencies						
Clear resolution	0.515	0.500	20,556	1.000	0.000	1,595
Among clearly resolved requests:						
Completed	0.787	0.409	10,591	0.797	0.402	1,595
Rejected	0.213	0.409	10,591	0.203	0.402	1,595
Filings per submitter	11.81	54.60	1,740	17.53	36.00	91
City Agencies						
Clear resolution	0.529	0.499	47,254	1.000	0.000	2,674
Among clearly resolved requests:						
Completed	0.848	0.359	24,998	0.841	0.366	2,674
Rejected	0.152	0.359	24,998	0.159	0.366	2,674
Filings per submitter	18.11	118.05	2,609	24.99	62.29	107
Panel B: State-Level Variables						
Corruption rate	0.275	0.132	50			
Average FOIA score	6.036	2.287	50			
State population (2013)	6,306,863	7,055,386	50			
State income (2013)	54,297	10,197	50			

In Appendix Table A1, we show the rejection rates (conditional on being completed or rejected) for the 10 federal agencies with the largest number of FOIA requests, all 50 states, and the 15 cities with the largest number of FOIA requests. One can easily discern where the high federal rejection rate comes from: Enforcement agencies – most notably the Central Intelligence Agency and National Security Agency – have rejection rates well above 60 percent. The federal agency with the most FOIA requests by far is the Federal Bureau of Investigation, which has a rejection rate of nearly 35 percent, comparable to other enforcement entities (e.g., Customs and Border Protection, Drug Enforcement Agency).

State-level variation plausibly offers our first window into the extent to which rejections reveal something about government openness. If we compare, say, states with above-median corruption to those with below-median corruption, the rejection rates are 0.24 versus 0.15

respectively; the rank correlation between the two is 0.38. We will explore these patterns further in a regression framework − also including city agencies within each state − in the next section.

2.2.2 Concerns, limitations, and applicability of the MuckRock data

Before proceeding to our results, we discuss the potential concerns and limitations that arise from our focus on MuckRock-based requests, and whether rejection is an appropriate measure of transparency.

As noted above, MuckRock constitutes a very small fraction of FOIA requests that are made to federal agencies. And because we cannot access this universe of requests, we cannot assess the extent to which the requests in our data reflect the broader set of requests made of federal government agencies (for city and state agencies we have not found consistent statistics for even the total number of requests). This limits our ability to evaluate the extent to which requests filed via MuckRock are representative of the broader population of requests. We can nonetheless claim coverage of a broad and diverse set of agencies that are petitioned and an array of topics. Beyond the range suggested by the federal requests shown in Appendix Table A1, we also record requests of, for example, police departments in 3,109 municipalities and state-level departments in 50 states; topics range from queries on salaries of city officials in various departments to use of force by police to information on air quality.

While it is not possible to compare our sample to the universe of FOIA requests, we can compare overall rejection rates at federal agencies (which are not part of our analysis sample) to rejection rates in the MuckRock data. We do so using official data obtained from `FOIA.gov`, where statistics are published on FOIA requests at the level of individual agencies and bureaus. In Figure B5, we show the relationship between federal agency rejection rates based on administrative data versus rejection rates calculated in our MuckRock data. A clear positive relationship emerges − the raw correlation between the two is $\rho = 0.57$ while the bivariate regression yields a coefficient of 0.37 ($p < 0.01$). This gives us some confidence that MuckRock requests are reflective of the broader set of requests along this important dimension.

Returning to the particular agencies targeted by MuckRock requests, it is notable that over a third of our matched-group sample is focused on police departments, which have relatively high rejection rates. The high rejection rate may result from the contentious nature of many of these requests. This is, on the one hand, an important caveat in thinking about the extent to which one may generalize from our sample to FOIA requests more generally. On the other hand, particularly given our focus on information and political accountability, the requests we study may be particularly well-suited to evaluating the role of FOIA laws in facilitating transparency around elections.

Finally, focusing on the MuckRock data specifically, there is the further concern that the set of matched-group requests – which we focus on for reasons of internal validity – may differ systematically from those in the larger set of MuckRock requests. At least in terms of rejection rates, they are remarkably similar. For example, Appendix Figure B6 shows that, at the state level, there is an extremely close correspondence between matched-sample rejection rates and those from the full set of MuckRock requests.

A final question of interpretation is that we base our measure of transparency in part on request rejections which, as discussed above, often are presented by responding agencies as reflecting confidentiality concerns. This naturally raises the issue of whether states that have lower rejection rates are "too open" relative to some optimal benchmark.[26] The question of optimal transparency is beyond the scope of our analysis; however, the fact that, as we show below, rejection rates are positively correlated with corruption suggests at a minimum that differences in responsiveness reflect something beyond differences in confidentiality standards.

3 Results

3.1 Relationship between state-level corruption and FOIA rejections

In our first set of analyses, we show the correlation between corruption and FOIA rejection rates. We first present analyses based on state-year and state level aggregates of all FOIA requests made of state and city agencies, and the measure of state-level corruption from Campante and Do (2014). We then provide a parallel set of results based on our set of matched groupings.

In our first analysis, we collapse the data to the state-year level, for the years 2013-2024 (due to the very small number of requests via MuckRock in 2010-2012), and consider the following specification:

[26] A 2013 request from George LeVines, at the time a Massachusetts-based journalist, makes this clear. LeVines asked the state's police departments to provide, "any lists, databases and inventory rosters containing equipment used in the field of duty (i.e., firearms, protective gear, surveillance equipment, tactical and defense equipment, vehicles, etc)." Many departments declined to respond, citing concerns that such disclosures could be useful for those planning to attack the department or the public at large. That is, one can question whether the minority of departments that provided such inventories were in fact being *too* transparent.

$$Rejection\ Rate_{st} = \alpha + \beta \times Corruption\ Rate_s + X_{st} + \epsilon_{st} \tag{1}$$

The dependent variable is the fraction of requests that are rejected, among those that end in either full completion or rejection in state s in year t, and the sample includes all state-year observations for which there was at least one request that was either completed or rejected (a total of 590 state-year observations out of a possible 600). The main independent variable, *Corruption Rate*, is the per capita rate of corruption convictions in state s during 1976-2002. We include as control variables X the log of population and log of income per capita, as well as whether a state limits FOIA requests to state residents (as is the case in Arkansas, Delaware, Kentucky, Tennessee, and Virginia), which in our inspection of individual FOIA requests we have found to be a very common reason for rejection.[27] Standard errors are clustered by state.

Results based on equation (1) appear in Table 2, columns 1−4, with progressively more controls added. Across all specifications, the coefficient on *Corruption Rate* is stable around 0.23-0.31 (significant at the 5% level). Note that, when we include a variable in column 4 that proxies for stringency of FOIA laws from the Open Government Guide, it has a negative coefficient, though it is not statistically significant at conventional levels, and as we will see shortly, the negative relationship does not hold for our favored within-group analysis.

The coefficient on *Corruption Rate* is very similar when we collapse all variables to the state level in columns 5 and 6 (again significant at the 5% level). To provide some sense of magnitude in these cross-sectional relationships, both *Rejection Rate* and *Corruption Rate* have standard deviations of about 0.13, so that a standard deviation increase in corruption rate is associated with a 21 percent increase in the average rate at which FOIA requests are rejected, relative to the baseline average rejection rate of 20%.

In Figure 1 we illustrate the relationship between *Corruption Rate* and *Rejection Rate*. No individual outlier is driving the result, and the relationship is approximately linear.

As discussed in the introduction, one useful aspect of the MuckRock platform is that it facilitates the filing of identical requests across jurisdictions. From an identification perspective, this feature allows us to compare the success and failure of essentially identical requests submitted to various (high versus low corruption) jurisdictions. We define a group g of requests as those filed by the same submitter with the same content. We amend the above specification to account for the different data structure:

[27]We have confirmed that our results are not sensitive to the treatment of Alabama, which has ambiguous rules on residency. Additionally, if we include the number of government employees per capita, it has no impact on the estimated coefficients; we do not include this as a control because, as noted in Section 2, there is little variation in this ratio across states.

Table 2: State-Level Corruption and FOIA Responsiveness

This table reports the relationship between state-level corruption and the average FOIA rejection rate by state and city agencies in the given state. Columns 1 to 4 report the results using state-year level data, and columns 5 and 6 report the results using data at the state level. We require that the given FOIA filing is either rejected or completed. We use the corruption measure of Campante and Do (2014); the outcome variable, *Rejection Rate*, is the average FOIA rejection rate by the given state. The sample mean of the dependent variable, *Rejection Rate*, is 0.20, and the standard deviation of *Corruption Rate* is 0.13. Standard errors are clustered at the state level and reported in parentheses. * indicates significance at the 10% level, ** at the 5% level, and *** at the 1% level.

	Dependent Variable: *Rejection Rate*					
	State-Year Level				State Level	
	(1)	(2)	(3)	(4)	(5)	(6)
Corruption Rate	0.308**	0.309**	0.290**	0.226**	0.309**	0.256**
	(0.125)	(0.127)	(0.124)	(0.101)	(0.127)	(0.114)
Log(Income)			-0.106	-0.079		-0.078
			(0.112)	(0.089)		(0.096)
Log(Population)			-0.003	0.001		-0.003
			(0.016)	(0.013)		(0.013)
Average FOIA Score				-0.013		
				(0.008)		
$\mathbb{1}$*(Residents Only)*				0.193***		0.187***
				(0.050)		(0.052)
Fixed Effects						
Year		X	X	X		
N	590	590	590	590	50	50
R^2	0.046	0.103	0.114	0.224	0.108	0.335

Figure 1: State-Level Corruption and FOIA Responsiveness

This figure plots the relationship between the state-level corruption rate from Campante and Do (2014) and the average state-level FOIA rejection rate. The 95% confidence interval (CI) is also plotted.

$$Rejected_{rgst} = \alpha + \beta \times Corruption\ Rate_s + X_{st} + \gamma_g + \eta_{q(t)} + \epsilon_{rgst} \tag{2}$$

where *Rejected* is an indicator variable denoting that request r from matched group g filed in state s at time t was rejected. In all specifications, we use two-way clustering by state and matched-group. In our most stringent specifications, we include 800 group fixed effects. When we use this approach, we identify the relationship between corruption rate and FOIA rejections holding constant the characteristics of FOIA requests and submitters.

We present results based on this matched grouping approach in Table 3. In the first three columns, we present specifications that are comparable to those in Table 2. The patterns are similar to those based on the preceding cross-sectional analyses, though the point estimates are marginally higher. Column 4 repeats the specification of column 3, but includes only

Table 3: Matched Sample Analysis of Corruption and FOIA Response

This table presents the association between the state-level corruption rate from Campante and Do (2014) and the responses of FOIA requests filed in that state, focused on "matched groups" of requests (see text for details). The dependent variable is $\mathbb{1}$(Rejected), an indicator variable denoting that a request was rejected. The sample mean of $\mathbb{1}$(Rejected) is 0.16; the standard deviation of *Corruption Rate* is 0.12. In columns 4, 5, and 6, we require that the standard deviation of corruption within a "matched group" should be greater than 0. Standard errors are double clustered at the state and group level.

	Dependent Variable: $\mathbb{1}$(Rejected)					
	(1)	(2)	(3)	(4)	(5)	(6)
Corruption Rate	0.419***	0.440***	0.438***	0.422***	0.343**	0.277**
	(0.146)	(0.144)	(0.143)	(0.149)	(0.129)	(0.110)
Log(Income)			-0.041	-0.075	-0.106	-0.068
			(0.089)	(0.103)	(0.092)	(0.077)
Log(Population)			-0.033**	-0.027**	-0.018	-0.012
			(0.014)	(0.013)	(0.011)	(0.008)
Average FOIA Score						-0.005
						(0.006)
$\mathbb{1}$*(Residents Only)*						0.244***
						(0.047)
Fixed Effects						
Quarter of Submission		X	X	X	X	X
Matched-Group					X	X
Condition						
Matched-Group-level σ(*Corruption Rate*)>0				X	X	X
N	21,754	21,754	21,754	12,603	12,603	12,603
R^2	0.018	0.055	0.065	0.061	0.292	0.316

groupings for which there is within-group variation in the state of request, since this is the variation we will exploit in our within-group analysis. In the final two columns, we present results for our preferred specifications that include matched group fixed effects. The point estimate on *Corruption Rate* is 0.34, significant at the 5% level. The final column also includes a measure of FOIA *legal* stringency (*Average FOIA Score*) and also a control for whether filing a FOIA requires in-state residency. Notably, *Average FOIA Score* is uncorrelated with rejection rates and, indeed, is uncorrelated with *Corruption Rate* as well.

Overall, we draw two main lessons from the results we present in this section. First, our findings suggest that agencies' responses to informational requests are highly correlated with an institutional feature – state-level corruption – that one might expect ex ante to be associated with a government's willingness to be held up to scrutiny. We see this as an interesting fact in itself which also serves as a basic reality check on the data. Our matched grouping results bolster the credibility of the basic relationship between corruption and revealed transparency, and suggests that we may use the structure of the data to explore the determinants of transparency across any source of variation that exists within our groups. Motivated by this observation, we now turn to our main analysis, which explores how FOIA responsiveness varies with electoral pressures, exploiting within-group variation in FOIA outcomes across local governments and over time.

3.2 Electoral pressure and FOIA responsiveness

In this section, we examine responsiveness to FOIA requests before versus after city- and state-level elections. As we explained in the introduction, the relationship between electoral pressures and FOIA responsiveness is theoretically ambiguous. Officials may be less responsive to FOIA requests if the resulting revelations risk political embarrassment. Alternatively, a well-meaning civil servant may be particularly attentive in providing timely disclosures to voters ahead of elections, so that electoral pressures may lead to greater responsiveness.

Our specification focuses on groupings for which at least one request is to an agency in a state or municipal government which has a gubernatorial or mayoral election in the upcoming 180 days, and at least one request for a different state or municipal agency where there is *not* an election in the next 180 days as a benchmark.

As discussed in Section 2, there are two ways of avoiding a pre-election FOIA response: delay or rejection. Our primary measure, *Failure*, combines both.

Our main specification explores the direct relationship between an upcoming election and *Failure*, and also considers how this relationship might vary as a function of the institutional environment. Recall that we include both state and municipal elections; let us denote $b \in$

$\{c, s\}$ as the level of government (city or state) in which a given FOIA request was made. Our analyses are then based on the following specification:

$$\begin{aligned} Failure_{rgbt} &= \alpha + \beta_1 \times Election_{rgbt} + \beta_2 \times Corruption\ Rate_{s(b)} \\ &\quad + \beta_3 \times Election_{rgbt} \times Corruption\ Rate_{s(b)} + \gamma_g + \eta_{q(t)} + \epsilon_{rgbt} \end{aligned} \tag{3}$$

for request r that is a part of group g, submitted at an agency in government b at time t. We include a set of matched group fixed effects, γ_g, for requests with the same submitter and content, as well as quarter of submission fixed effects, $\eta_{q(t)}$, for requests submitted in the same calendar quarter. In some specifications, we will also include state fixed effects, $\lambda_{s(b)}$, which absorbs the level effect of corruption. We cluster standard errors by matched group and by state.[28]

Figure 2: Timeline of a Matched Group of Five Requests

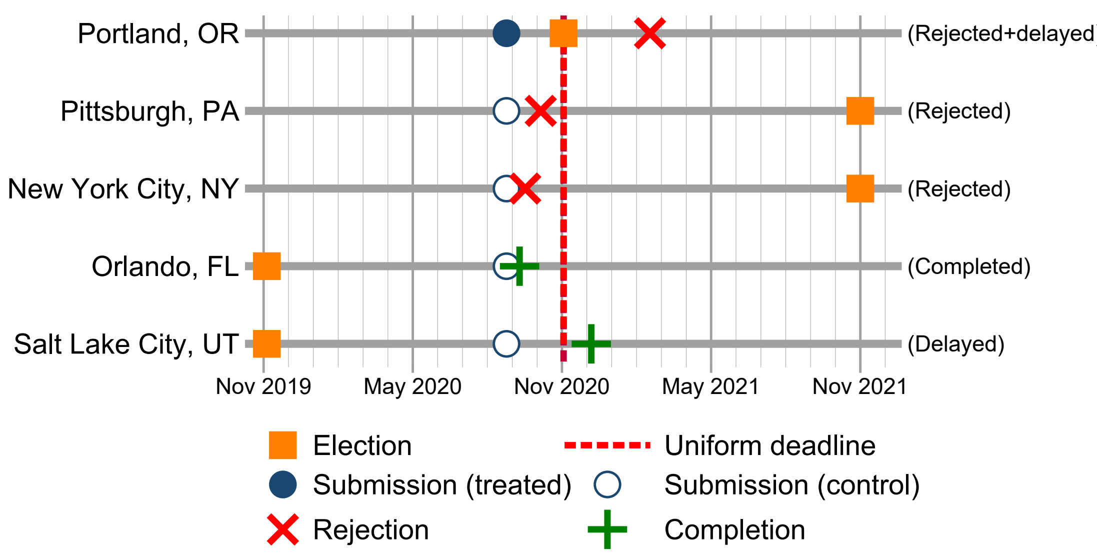

We illustrate our methodology in Figure 2, using a matched group of five requests made on August 25, 2020.[29] The "treated" request in this case was submitted to the police department

[28]All results are robust to three-way clustering standard errors by matched group, state, and quarter of submission.

[29]There were 22 requests in total made by the submitter on that date, all with identical titles. However, the other 17 were not marked as "rejected" or "completed."

in Portland, Oregon, 70 days before the city's mayoral election on November 3. The same request (from the same submitter) was made to police departments in four other cities, none of which had an election in the 180 days after the request, and thus serve as control requests. Because the only treated request is from Portland, we use its election date as the uniform deadline to define whether a response was delayed (the dotted red vertical line). The timing and nature of responses are captured by red Xs (rejections) and green pluses (completions). So, for example, Portland's response was both delayed and rejected; Salt Lake City's was completed, but late, while Pittsburgh's was on time but rejected. Of the five requests *Failure* = 0 only for Orlando, where it was completed, and done so before the deadline.

Before presenting the results based on our selected sample of requests with clear resolutions, we first provide, in Table 4, results from specifications similar to that of Equation (3), but for the full set of matched group requests, looking separately at each potential outcome. In the first three sets of columns, we consider the outcomes that constitute 95% of the sample, i.e., rejection, completion, no documents, and in the fourth group of columns we consider the residual set of "other outcomes." Each grouping thus considers the frequency of a given outcome relative to all other outcomes in a given matched group.[30] Focusing first on the direct relationships among *Corruption Rate*, $\mathbb{1}(\textit{Election})$, and the various possible outcomes, we make two observations: first, in this matched group election sample, we again observe a strong positive relationship between *Corruption Rate* and rejection, with a corresponding negative relationship with completion; there is no significant (or economically meaningful) relationship to $\mathbb{1}(\textit{No Documents})$ or $\mathbb{1}(\textit{Other})$. Turning to the relationship between pre-election submission and outcomes, in the absence of the interaction term we observe no significant (or economically large) coefficient on $\mathbb{1}(\textit{Election})$ in any specification, indicating that, on average, election timing is unrelated to request outcomes. Finally, looking at the coefficients on the interaction of *Corruption Rate* and $\mathbb{1}(\textit{Election})$, the most prominent pattern we observe is a statistically significant ($p < 0.01$) and large positive coefficient in predicting rejection (columns 2 and 3), indicating that, in higher *Corruption Rate* states,

[30] Conceptually, the analysis is better-suited to a multinomial regression. However, even if we limit ourselves to just four categories, the most straightforward adaptation of our empirical strategy into a multinomial logistic regression requires us to estimate 3×410 parameters once we include matched-group and time fixed effects. Many of these category-specific parameters cannot be identified due to the lack of variation in the corresponding category within a given matched group, which prevents the convergence of a maximum likelihood estimator. In Appendix C, we describe and present results based on an empirical strategy that is adjusted to allow for estimation in our data, in particular omitting the matched-group fixed effects from our specification and clustering standard errors on matched groups *or* states rather than two-way clustering. These results provide a near-identical message to that of the simpler OLS-based approach we provide in the main text, regardless of our choice of clustering.

there is a relative increase in rejections ahead of elections. The corresponding coefficients for other outcomes are uniformly negative and, while large in magnitude in some cases − especially for $\mathbb{1}(No\,Documents)$ − none is significant even at the 10% level.

We now turn to focus on what we see as our main analyses, which incorporate potential failure via rejection *or* delay in response past the election date. These results appear in Table 5. We provide two sets of results − in columns 1-4, we consider rejection or delay as indicating failure, for the full sample of requests, so that we effectively consider any outcome that is provided before the election that is *not* rejection as the benchmark "non-failure." In columns 5-8, we present analogous results, but focus only on requests that end in clear resolution, i.e., ones that are ultimately rejected or completed. Finally, in columns 9-12 we repeat these specifications, but instead of classifying only rejection and responses delayed past the election deadline as failure, we consider any outcome *except* completion before the election deadline as indicating an unsatisfactory resolution. That is, for the full sample of requests, we use the indicator variable *Completed on Time* as the outcome.[31]

Looking first on the full sample results, in the first column, we present the relationship that includes just the direct effects of *Election* and *Corruption Rate*, as well as group fixed effects. The point estimate on *Election* is −0.016, and does not approach statistical significance. The standard error of 0.022 indicates that we can reject, at the 5% level, an effect size smaller than -0.059 or larger than 0.028. Thus, on average there is no substantial difference in request success as a function of election timing. We again find a positive coefficient on *Corruption Rate*, indicating a higher failure rate in states with more corruption convictions per capita. Finally, we look at *Completed on Time* as an outcome in column 9. Since the dependent variable now reflects success rather than failure, the negative sign (significant at the 5% level) on *Corruption Rate* implies a higher completion rate in low-corruption states; we continue to observe a very small and statistically insignificant coefficient on the direct effect of $\mathbb{1}(Election)$.

Our main interest is whether the effect of electoral pressures varies with the broader institutional environment: in a setting in which government malfeasance is minimal, there is less incentive and/or willingness to withhold information from the voting public, while the opposite is true in a more corrupt environment. We thus assert that the prediction of less responsiveness around elections is more likely in less accountable political environments, while the opposite may be true for more accountable jurisdictions, i.e., a positive coefficient

[31] Note that this is a very conservative definition, since it even defines partial completion as unsatisfactory. Please see Appendix Table A9, columns 4-6, for a variant that treats partial completion as success. Also observe that $\mathbb{1}$(Failure) and $1-$ *Completed on Time* are identical for the set of clear resolution requests, hence our focus on the full sample.

Table 4: Matched Sample Analysis of Request Outcomes Around Elections

This table presents the test of whether jurisdictions with an election have a different propensity to issue FOIA request rejection, completion, "No Documents", or other responses during the 180 days prior to the jurisdiction's mayoral or gubernatorial election when we match the given FOIA request to other identical FOIA requests that were filed to a government agency in a different jurisdiction by the same submitter. The variable $\mathbb{1}$(Election) takes the value of one if the FOIA request was filed to a government agency in a jurisdiction that had an election within 180 days following the initial request date. The mean of the dependent variable $\mathbb{1}$(Rejected), is 0.11; for dependent variables $\mathbb{1}$(Completed), $\mathbb{1}$(No Documents), and $\mathbb{1}$(Other Outcome), the means are 0.52, 0.31, and 0.05, respectively. Standard errors are double clustered at the matched group and state level, and are reported in parentheses. * indicates significance at the 10% level, ** at the 5% level, and *** at the 1% level.

	Dependent Variable											
	$\mathbb{1}$(Rejected)			$\mathbb{1}$(Completed)			$\mathbb{1}$(No Documents)			$\mathbb{1}$(Other Outcome)		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
$\mathbb{1}$(Election)	-0.003	-0.074**	-0.084***	-0.000	0.015	0.029	0.010	0.058*	0.060**	-0.007	0.000	-0.006
	(0.014)	(0.031)	(0.026)	(0.009)	(0.029)	(0.025)	(0.012)	(0.031)	(0.030)	(0.010)	(0.021)	(0.024)
Corruption Rate	0.293***	0.239**		-0.332***	-0.321***		0.088	0.125*		-0.048	-0.043	
	(0.104)	(0.102)		(0.111)	(0.111)		(0.064)	(0.074)		(0.035)	(0.045)	
$\mathbb{1}$(Election) × *Corruption Rate*		0.275***	0.316***		-0.059	-0.110		-0.189	-0.193		-0.028	-0.013
		(0.100)	(0.090)		(0.097)	(0.086)		(0.126)	(0.120)		(0.074)	(0.081)
Fixed Effects												
Matched-Group	X	X	X	X	X	X	X	X	X	X	X	X
Quarter of Submission	X	X	X	X	X	X	X	X	X	X	X	X
State			X			X			X			X
N	7,232	7,232	7,232	7,232	7,232	7,232	7,232	7,232	7,232	7,232	7,232	7,232
R^2	0.202	0.204	0.260	0.343	0.343	0.366	0.351	0.351	0.361	0.117	0.117	0.129

Table 5: Matched Sample Analysis of Failure to Respond Around Elections

This table examines whether jurisdictions with an upcoming election have a different rate of failure to respond to FOIA requests, defined as either a rejection or failure to respond before election day, during the 180 days prior to the jurisdiction's mayoral or gubernatorial election when we match the given FOIA request to other identical FOIA requests that were filed to a government agency in a different jurisdiction by the same submitter. The dependent variable $\mathbb{1}$(Failure) takes the value of one if the FOIA request received a rejection or a "Delayed Response." The dependent variable $\mathbb{1}$(Completed on Time) takes the value of one if the FOIA request was fully completed before the election deadline. The variable $\mathbb{1}$(Election) takes the value of one if the FOIA request was filed to a government agency in a jurisdiction that had an election within 180 days following the initial request date. The sample of clear resolutions is limited to requests that received a clear resolution, i.e., rejection or completion. The mean of the dependent variable $\mathbb{1}$(Failure) is 0.30 for the sample of all requests and 0.37 for the sample of clear resolutions only, and the mean of the dependent variable $\mathbb{1}$(Completed on Time) is 0.40 for the sample of all requests. Standard errors are double clustered at the matched group and state level, and are reported in parentheses. * indicates significance at the 10% level, ** at the 5% level, and *** at the 1% level.

	Dependent Variable											
	$\mathbb{1}$(Failure)								$\mathbb{1}$(Completed on Time)			
	All Requests				Clear Resolutions Only				All Requests			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
$\mathbb{1}$(Election)	-0.016	-0.122***	-0.132***	-0.054	-0.012	-0.136**	-0.140***	0.410	0.003	0.055	0.064**	-0.301
	(0.022)	(0.043)	(0.040)	(1.116)	(0.027)	(0.053)	(0.048)	(1.217)	(0.013)	(0.034)	(0.030)	(0.747)
Corruption Rate	0.260*	0.178			0.320	0.212			-0.274**	-0.234*		
	(0.142)	(0.136)			(0.192)	(0.185)			(0.124)	(0.123)		
$\mathbb{1}$(Election) × *Corruption Rate*		0.411***	0.442***	0.455***		0.487***	0.489***	0.483***		-0.199*	-0.227**	-0.228**
		(0.143)	(0.148)	(0.143)		(0.152)	(0.145)	(0.138)		(0.104)	(0.097)	(0.102)
$\mathbb{1}$(Election) × *Log(Income)*				0.031				-0.017				0.029
				(0.102)				(0.110)				(0.073)
$\mathbb{1}$(Election) × *Log(Population)*				-0.025*				-0.022				0.003
				(0.015)				(0.017)				(0.011)
$\mathbb{1}$(Election) × *Average FOIA Score*				-0.003				-0.001				0.000
				(0.006)				(0.007)				(0.004)
$\mathbb{1}$(Election) × $\mathbb{1}$(Residents Only)				-0.034				0.002				0.007
				(0.036)				(0.056)				(0.041)
Fixed Effects												
Matched-Group	X	X	X	X	X	X	X	X	X	X	X	X
Quarter of Submission	X	X	X	X	X	X	X	X	X	X	X	X
State			X	X			X	X			X	X
N	7,232	7,232	7,232	7,232	4,267	4,267	4,267	4,267	7,232	7,232	7,232	7,232
R^2	0.236	0.237	0.270	0.270	0.296	0.299	0.345	0.345	0.301	0.302	0.324	0.324

on the term $\mathbb{1}(Election) \times Corruption\ Rate$. In column 2, we find that the coefficient on this interaction is 0.41 (significant at the 1% level), and implies a "crossing point" of *Corruption Rate* = 0.30, where the *Election* effect is zero, at the 64th percentile of the distribution of *Corruption Rate*. The third column of Table 5 adds state fixed effects. We can no longer identify a direct effect of *Corruption Rate*, but the main coefficient of interest, on $\mathbb{1}(Election) \times Corruption\ Rate$, is relatively unaffected by this inclusion. To probe the sensitivity of our results to omitted variable concerns, in column 4 we include as covariates the interaction of $\mathbb{1}(Election)$ with controls for population, income, FOIA stringency scores, and in-state FOIA filing requirements. The coefficient on $\mathbb{1}(Election) \times Corruption\ Rate$ is unchanged. In columns 5-8, we repeat these analyses, limiting the sample to requests that received a clear resolution, i.e., rejection or completion. The patterns are broadly similar to those reported in the full-sample results, though the point estimates are larger throughout, as we might anticipate if the inclusion of ambiguous categories as non-failures effectively treats them all as zeros when in fact the true outcome is less clear.[32] The results using $Completed\ on\ Time$ yield a very similar message – the coefficient on the interaction $\mathbb{1}(Election) \times Corruption\ Rate$ is negative and significant at the 5% level. Overall, the results indicate that a combination of changes in rejection and completion times – rather than substitution to ambiguous categories – drive our main findings.

For the remainder of this section, we will focus on the set of requests with clear resolutions to avoid this sort of misclassification of ambiguous resolutions; we provide further robustness checks based on more expansive samples and/or more inclusive definitions of failure in Section 4.2.

In Appendix Table A3, we look separately at the two components of *Failure*, i.e., *Rejected* and *Delayed Response*. While the coefficients are of consistent signs for both sets of results, it is clear that our main results are driven primarily by rejection.

We next explore potential non-linearities in how elections impact FOIA responses as a function of different corruption levels in Appendix Table A4. Instead of the linear interaction, we include $\mathbb{1}(Election)$ interacted with four indicator variables for $Corruption\ Rate$ quartiles. The highest quartile has a positive and significant coefficient; the coefficients for all other quartiles are negative, but insignificant. This specification also provides the most straightforwardly interpretable coefficient estimates: the coefficient on $\mathbb{1}(Election) \times \mathbb{1}(Top\,Quartile\,Corruption\,Rate)$ of 0.11 implies a 30 percent increase in the probability of failure for pre-election requests in the set of highest-corruption states. The other corruption quartile interactions are all negative $-$ broadly consistent with our interpretation in the linear model of lower pre-election failure rates in less corrupt environments $-$ though none

[32]Note that two observations drop out of the sample because they are fixed-effect singletons.

approach statistical significance.

We also provide a more general illustration of the patterns in the data, based on a regression with $\mathbb{1}(Failure)$ as the outcome, and state fixed effects as well as these state fixed effects interacted with $\mathbb{1}(Election)$ as covariates (in addition to matched-group and quarter-of-submission fixed effects). We present the results of this specification in Figure 3, in which we show each state's corruption rate on the x-axis and each state's pair of coefficient estimates on the vertical axis. In effect, this graph captures the relationship between corruption and state-level failure rates pre-election versus other times. We believe this is the most transparent way of presenting our main result. We include lines of best fit for pre-election requests versus others, and observe a steeper gradient for the pre-election sample. We also observe that the steeper slope is driven at least in part by the relatively large gaps in pre-election versus other failure rates of the highest-corruption states.

Figure 3: Relationship between state-level corruption and rejection rates: pre-election versus other requests

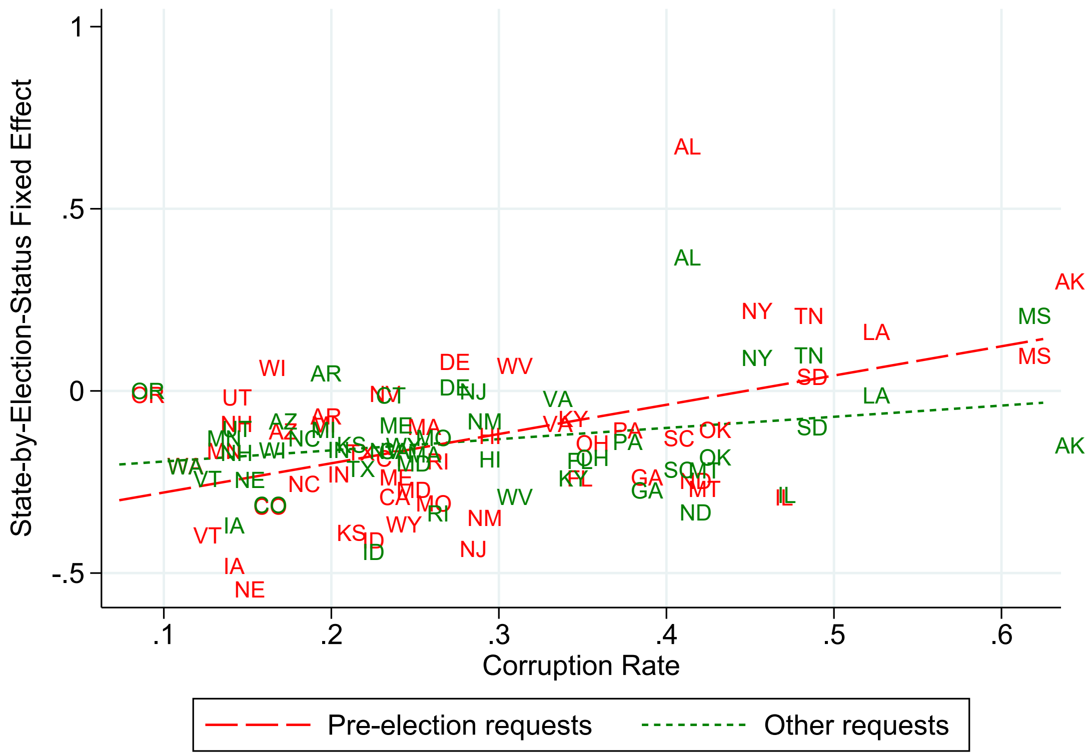

4 Robustness and Extensions

In this section, we explore the robustness of our results, and also examine whether the patterns we document in the preceding section vary by geographic or political characteristics. The main takeaway from these analyses is that our findings are quite stable across a range of sample splits, alternative specifications, definitions of failure, and sample selection criteria.

4.1 Alternative estimation approach: focusing on within-jurisdiction pre- versus post-election requests

In our main analysis, we use cross-jurisdiction variation within matched groups to identify the effect of elections (and its interaction with corruption). An alternative approach is to focus on within-jurisdiction variation around elections. While we cannot hold the features of a request constant, it allows us to consider the differential responses of a particular government pre- versus post-election.[33]

In Panel A of Appendix Figure A1, we provide rejection rates for high- versus low-corruption jurisdictions for requests filed in the 180 days before and 180 days after an election.[34] As suggested above, selection into this sample requires the presence of at least one request in the 180 days before an election and one request in the 180 days that follow. Additionally, to parallel our main results, we split the sample using the results from Table 5 to sort states into those with a positive versus negative impact of elections on failure – this crossing point is at *Corruption Rate* = 0.28; just over a third of the 3,191 requests in this sample are above this threshold.[35] Consistent with the presentation of our main results, more corrupt states have higher rejection rates for all requests, but this gap is considerably wider for pre-election requests. The gap narrows largely because rejection rates in high-corruption states decrease after elections. This approach also allows us to plot out rejection rates for the high- versus low-corruption sample over 30-day periods around election dates, which we provide in Panel B of the figure. Given the relatively small samples for each 30-day period, these data exhibit considerable noise. Notably, rejection rates are particularly high in the 30 days immediately preceding an election.[36]

[33]The biggest concern that this raises is that the types of request that are filed before elections are systematically different from those filed after elections. As noted earlier, we do not find any evidence that this is the case. See Appendix Table A5.

[34]We focus on rejection rather than delay in these analyses because of the difficulty in appropriately defining delay for monthly windows around the election date.

[35]In Appendix Figure A2 we split the data at the sample median. While less well-motivated conceptually, this split of the data yields similar patterns.

[36]In Appendix Figure A3, we also provide a placebo graph to assess the presence of seasonality,

This approach also can be conceptualized as a variant on our main estimating equation, in which we focus on the set of jurisdictions as defined above, and instead of addressing concerns of request heterogeneity with matched-group fixed effects, we address a distinct set of heterogeneity concerns with the inclusion of jurisdiction fixed effects. These results appear in Appendix Table A6, and yield point estimates that are broadly consistent with those based on identification that emphasizes within-matched-group variation.[37]

4.2 Are the results affected by the definition of rejection?

In our preferred analysis, we focus on the set of requests with clear resolutions (i.e., rejection or completion), showing that the results are robust to the inclusion of the full set of matched group requests. In those analyses, we classify ambiguous cases (e.g., no responsive documents or partial completion) as non-rejection. In this subsection, we explore whether our main results are sensitive to classifying these outcomes as non-failures, as well as the inclusion versus exclusion of "in process" requests in the sample. We consider three permutations of our main results:

1. Full set of matched group requests; failure defined as rejection, no responsive documents, or delay in response past election day.
2. Full set of matched group requests; failure defined as neither completed nor partially completed, or delay in response past election day.
3. Sample of "closed" requests (completed, partially completed, rejection, no responsive documents); failure defined as rejection, no responsive documents, or delay in response

given that elections occur most commonly on the first Tuesday in November. To generate an event plot that is seasonally comparable to Appendix Figure A1, but without any election, we take the full set of requests, and consider any first Tuesday of November for which there is no local election 360 days prior to or following this date in the city or state where it was filed. For this placebo election date, we then look at responses to any requests in the preceding and following 180 days. As in our main election analyses, the rejection rate is higher in high-corruption states. But the narrowing of the gap after the first Tuesday in November is modest to non-existent for this set of placebo dates. The accompanying Appendix Table A7 shows that the pre-versus-post difference is statistically indistinguishable from zero in all specifications.

[37]As a further alternative, in Appendix Table A8, we include all matched groups conditional on having either a request that takes place in the 180 days preceding an election *or* a request that occurs in the 180 days following an election; in both cases we require at least one request within the group for which there is no election in the preceding or following 180 days. We then augment our main estimating equation to include an indicator variable for post-election requests as well as its interaction with corruption. This serves essentially as a placebo against which to benchmark our results on pre-election requests. The point estimates on pre-election requests are similar to those we report in our main results, while those on post-election requests are consistently negative, though generally not statistically significant.

past election day.

Results are presented in Appendix Table A9. In all cases, the coefficient of primary interest on the term $Corruption\,Rate \times \mathbb{1}(Election)$ is significant at least at the 10% level (at least at the 5% level in our preferred specification that includes state fixed effects). The point estimates are smaller as well, which is consistent with our earlier findings, and our interpretation that our main results are driven primarily by unambiguous cases in the data, i.e., outright rejection or delay.

4.3 Are the results driven by particular subsets of the data?

It would be of concern for the breadth of interpretation of our results if they were driven by a relatively narrow segment of the data. In this section, we consider the more prominent dimensions of our data to explore whether our main results are sensitive to the inclusion/exclusion of particular groups of observations.

4.3.1 Are results driven by particular geographies?

To explore whether our findings are driven by particular states or regions of the U.S., we provide two "leave-out" analyses based on Equation (3), in which we re-estimate our results leaving out one state or one census region at a time. Appendix Figure A4 provides the resulting set of point estimates for the coefficients on $\mathbb{1}(Election)$ and $\mathbb{1}(Election) \times Corruption\,Rate$. To a surprising degree, we observe very stable point estimates in both the state-by-state (Panel A) and region-by-region (Panel B) results. This reflects in part the broad geographic variation in the number of requests in each state in our data. The state with the most requests in our elections dataset is California, with a total of 654 requests (15.32% of the sample), while for census regions it is West, with a total of 1,397 requests (32.72% of the sample).

4.3.2 Are the results driven by particular agencies?

We take a similar leave-out approach to examining whether our results are driven by any particular agency type. We divide agencies into seven broad groupings:

- Police, 55.77% of all requests.
- Other law enforcement (e.g., public safety departments and investigative units), 3.51% of all requests.
- Corrections, 6.89% of all requests.
- Education (i.e., schools and universities), 5.88% of all requests.

- Executive (i.e., mayors' and governors' offices), 4.01% of all requests.
- Record keeping (city clerk and secretary of state), 3.47% of all requests.
- Chief legal offices (district attorneys and attorneys general), 2.46% of all requests.

These results appear in Appendix Figure A5. Again, the results are very stable across all subsamples, indicating that no individual group of requests is driving our results.

4.3.3 Are results driven by particular submitters?

There exist some particularly heavy MuckRock users. The most prolific submitted a total of 585 requests (31 matched groups), or 13.70% of the total; four others made over 200 requests each. In Appendix Figure A6, we show that the point estimates are insensitive to the exclusion of requests from any of these submitters.

We also check whether our results are driven by submissions of journalists – a substantial majority of submissions on MuckRock – versus all others, by adding the third-order term $\mathbb{1}(Election) \times Corruption\,Rate \times Journalist$ to our main specification. We present these results in Appendix Table A10, which shows no evidence of any differential failure rate for requests submitted by journalists.

Overall, the results in this section indicate that our findings are not driven by any particular subset of requests.

4.4 Do pre-election requests systematically differ from post-election ones?

Our matched-group approach absorbs some potential omitted variables that could be correlated with requests made pre- versus post-election. Yet within-group variation remains which, importantly, could also be correlated with state-level corruption. We thus examine whether there are other request characteristics that are not held constant within a group that are predicted by $\mathbb{1}(Election) \times Corruption\,Rate$, and also whether the overall quantity of pre-election requests is higher, which could also – as a result of limited bureaucratic capacity – impact responsiveness.

4.4.1 Do pre-election requests have different characteristics from post-election requests?

It is possible that some features of requests – such as the agency they are directed toward or the identity of the submitter – could have a *differential* impact on responsiveness ahead of elections. For example, agencies that generally receive heavy request loads may be disproportionately affected by elections. To the extent that these characteristics are correlated with state-level corruption, they may serve as confounds in our election timing analysis. We consider some of the more prominent possible cases of this type of omitted variable problem here.

To examine whether pre-election rejection rates are driven by high-rejection agencies, we calculate the average rejection rate over our entire sample for each of 29 agency types (police; corrections; mayoral office; etc.), and add its interaction with $\mathbb{1}(Election)$ as a control in Appendix Table A11, along with agency-type fixed effects. The coefficient on this term is small and statistically insignificant; it has no effect on the coefficient on $\mathbb{1}(Election) \times Corruption\,Rate$.

A second concern is that a submitter may be motivated to file a request primarily based on concerns in a particular locale and may make similar requests elsewhere to, for example, form a benchmark or control for the response to their main request. If these "muckraking" requests tend to cluster in corrupt states, it could again create an omitted variable problem. As a proxy for a submitter's main geographic focus, we use their state of residence to define $\mathbb{1}(Home\,State)$, and include its interaction with $\mathbb{1}(Election)$ as a control. As shown in column 2 of Appendix Table A11, its coefficient is small and insignificant and does not affect the coefficient on $\mathbb{1}(Election) \times Corruption\,Rate$.[38]

4.4.2 Could the increase in pre-election rejection come from increased bureaucratic workload?

If request volume increases in advance of elections, rejection and/or response times could increase because agencies have difficulty dealing with requests in a timely manner due to limited bureaucratic capacity.

To examine whether the volume of FOIA requests submitted to city and state governments varies pre- versus post-election, we take the set of all city-months from January 2010 through June 2021 in the cities for which we have data on mayoral elections from the ALGED and data on FOIA requests from MuckRock, as well as the set of all state-months during the same period. We then count the number of submitted FOIA requests to city government agencies in each city-month and to state government agencies in each state-month. Using the sample of 65,550 city-months across the 475 cities for which we observe both elections and any FOIA requests between 2010 and 2021 and 6,900 state-months across all 50 states,

[38] We also show that these request characteristics are balanced between pre- and post-election requests, by using $\mathbb{1}(Home\,State)$ and $Agency\,Rejection\,Rate$ as outcomes in analyses based on Equation (3). See Appendix Table A12.

we take the number of FOIA submissions $Requests_{jt}$ in jurisdiction j at (monthly) time t, and we estimate

$$Requests_{jt} = \alpha + \beta \times Election_{jt} + X_{jt} + \epsilon_{jt} \quad (4)$$

where $Election_{jt}$ indicates whether the month t is one of the six months preceding an election in jurisdiction j, and X_{jt} includes controls for the level of the jurisdiction (i.e., state or local), state fixed effects, and year fixed effects. In cases where two elections occur within six months of each other in the same jurisdiction, we use only the first election.

The results appear in Appendix Table A13. Across all specifications, we find precisely estimated zeros for the coefficient on *Election*. The average jurisdiction in the sample receives 0.28 FOIA requests in a given month; comparing within states and years, an increase of more than 20% in monthly FOIA request submission volumes during pre-election periods would fall well outside of our estimated 99% confidence interval. Overall, we do not find any evidence supporting the view that high pre-election request volume (combined with differential bureaucratic capacity across states) accounts for our results.

4.4.3 Are the results driven by selection into the matched group sample?

We showed in Appendix B that the rejection rate for our matched-group sample was similar to that of the overall set of MuckRock requests. Here, we assess whether our election results are robust to including all MuckRock requests, i.e., not limiting ourselves to matched-groups with at least one pre-election request. To conduct this analysis, we constructed the sample of all 13,399 clearly resolved (i.e., either fully completed or rejected) requests in the 50 states and 468 cities appearing in both the MuckRock data and elections data filed between 2010 and the first half of 2021 (and so can be matched to the ALGED elections database). We then estimate our main elections analysis (without matched-group fixed effects) on this sample. These results appear in Appendix Table A14. The first three columns show the results for the unrestricted sample, while columns 4-6 show the results for the sample restricted to matched groups in which at least one request is filed in the 180 days prior to an election and at least one request is not (i.e., our matched-group sample), but we still do not include matched-group fixed effects. The results across all columns are qualitatively similar to those we report in our main results in Table 5: throughout, the coefficient on $Corruption\,Rate$ is positive and significant at least at the 5% level. The coefficient on $\mathbb{1}(Election) \times Corruption\,Rate$ is consistently positive, though not statistically significant in the full sample without state fixed effects.

4.4.4 Are pre-election requests more likely to receive a clear resolution?

In our main analyses, we focus on whether a request is rejected or completed, given the ambiguities in other classifications. But this raises a concern that a clear resolution (and hence selection into the sample) could itself be correlated with electoral timing. This could be the case if, for example, officials have insufficient time to chase down a response to every request (and hence may only partially complete them, or report that they could not find responsive documents).

In Appendix Table A15, we thus present our main specification without limiting the sample to only rejected or completed requests. In the first three columns, we repeat our main analysis, but with $\mathbb{1}(Rejected\,or\,Completed)$ as the dependent variable to capture selection into our main analysis sample. We find that this is predicted by the interaction $\mathbb{1}(Election) \times Corruption\,Rate$. In the next two sets of columns, we repeat this analysis for $\mathbb{1}(Rejected)$ and $\mathbb{1}(Completed)$ separately. Focusing first on the term of primary interest, we find that the significant coefficient on $\mathbb{1}(Election) \times Corruption\,Rate$ for clear resolution is driven entirely by rejections (columns 4-6); in fact, while not statistically significant, the coefficient is negative in predicting whether the response to a request is fully completed (columns 7-9). Also of note, the direct effect of $Corruption\,Rate$ is positive and significant in predicting rejection in columns 4 and 5, while it is of the opposite sign and significant in predicting completion (columns 7-8). Thus, our earlier result linking corruption to rejections in general is also not driven by selection into the sample of resolved requests.

As a further related robustness check, we also look at $\mathbb{1}(Delayed\,Response)$ as an outcome using the full set of requests. This addresses the concern that ambiguous categories may be used as a delay tactic. It is important to note that two patterns in the data indicate that this is unlikely to be a substantial concern. First, the preceding results show that, if anything, these ambiguous categories appear *less* frequently in response to pre-election queries. Second, in our data, these categories tend to receive, if anything, speedier resolutions on average – in particular, the largest ambiguous category by far, "no responsive documents," has a median response time of 16 days, as compared to 28 days for rejections and 27 days for completions. In practice, when we repeat our analysis with $\mathbb{1}(Delayed\,Response)$ as an outcome in columns 10-12 of Table A15, the patterns are very similar to those we observe in our selected sample of requests with a clear resolution in our main analysis. In particular, the point estimate on $\mathbb{1}(Election) \times Corruption\,Rate$ is 0.13; it is not significant, as is the case for our main analysis sample, where we found that pre-election "failure" was caused primarily by rejection.

4.5 Robustness to additional empirical considerations

We begin by examining the sensitivity of our main results to definitions of the pre-election window. We consider windows of as little as 120 days and as much as 360 days. The short-window results appear in Appendix Tables A16 and A17, while the long-window results are in Appendix Tables A18 and A19. Reassuringly, focusing on our $\mathbb{1}(\textit{Rejected})$ outcome, which is not a function of the length of the pre-election window, the patterns we observe based on a 180-day window are largely unaffected by using a shorter or longer window; in particular, the coefficient on $\textit{Election} \times \textit{Corruption Rate}$ is significant at least at the 10% level in all specifications that include state fixed effects, and consistently in the range of $0.28 - 0.53$. (The $\mathbb{1}(\textit{Delayed Response})$ results are more sensitive to shortening the window; these results were insignificant even in our main analyses.)

The other modeling choice in our main results relates to how we define whether a request is delayed beyond the election date. In the baseline, we use the longest election window within each matched group to set a uniform deadline. In Appendix Table A20, we instead use the shortest election window. As noted above, this change can affect the *Delayed Response* outcome, since many groups contain multiple election dates within the six-month window and an earlier cutoff naturally results in more requests being classified as delayed. Conversely, this modeling choice has no impact on our *Rejected* outcome, which drives the results we observe in our baseline results in Table 5. The results in Appendix Table A20 confirm this pattern: while the *Rejected* outcome remains unaffected, the *Failure* outcome becomes somewhat sensitive to this assumption. In our preferred specification in column 3, the coefficient on the interaction $\mathbb{1}(\textit{Election}) \times \textit{Corruption Rate}$ remains positive and is statistically significant at the 10% level.

While we show in Section 4.4.2 above that request volume does not systematically vary pre- versus post-election, we may still be concerned that requests tend to go to higher-volume agencies which, in advance of elections, are particularly prone to rejection. In Appendix Table A21, we therefore add total agency request volume and its interaction with $\mathbb{1}(\textit{Election})$ as controls. Whether we measure volume based on total requests over the whole sample period or requests in a given year, neither the direct effect nor its interaction with the election indicator variable approaches statistical significance. Importantly, the inclusion of this term also has no effect on the estimated coefficient on $\mathbb{1}(\textit{Election}) \times \textit{Corruption Rate}$.

In Appendix Table A22, we distinguish between state versus municipal elections. While we do not have any ex ante expectations about any differential effect of elections on FOIA outcomes between these two groups, the point estimate on $\mathbb{1}(\textit{Election}) \times \textit{Corruption Rate}$ is considerably larger for state-level agencies and elections (though in our preferred specifications in columns 3 and 6, the coefficients are both significant at least at the 5% level).

We next examine whether the effects differ by political party of the incumbent. There is no reason to expect such differences, so in a sense this analysis reflects a check of whether the results are stable across subsamples. The findings in Appendix Table A23 show that the coefficient of primary interest is directionally similar for the two groups, and we cannot reject that they are equal across parties.

It is natural to consider whether the results are more pronounced in more competitive elections. However, this is problematic for several reasons. First, one must find a credible proxy for competitiveness, which is a measurement challenge that the field of political economy has yet to resolve (Shaukat, 2019). A more immediate concern is that there are relatively few close elections − a total of 96 in our sample − which reduces the sample size by nearly 80% even with a vote margin of 10%. We nonetheless present the results for close elections using realized vote margins of 5% and 10% in Appendix Table A24. None of the coefficients in any specification approaches statistical significance in these results.

5 Conclusion

In this paper we provide what is, to our knowledge, the first empirical analysis of how electoral incentives affect government transparency. To do so, we collect a set of "natural experiments" based on groups of identical FOIA requests submitted via the website MuckRock, which are externally relevant in the sense of capturing the requests of actual filers. Because we use a within-group design, differences in response rates cannot be attributed to differences in the types of requests that are made across, say, more or less transparent jurisdictions.

We show that FOIA rejection rates predict state-level corruption convictions, which we take as a rough proxy for government probity. To explore how election incentives impact transparency, we compare responsiveness to identical requests across states and cities, exploiting variation in election timing across jurisdictions. Our main finding is that, in high-corruption settings, transparency − as captured by both rejections and delaying responses past election day − is lower in the months preceding an election.

We see many potential directions, both in the use of the ever-expanding data available of matched-group requests and in the analysis of incentives to minimize disclosure. For example, it may be possible to use tags associated with particular requests, or the text of the request itself to explore responsiveness to requests that risk embarrassment or bad PR; it may similarly be possible to classify the submitter as a journalist or everyday citizen on the basis of such text, which might similarly result in different treatment of comparable requests, particularly in advance of elections. We may also explore how particular events impact openness − the 2020 Black Lives Matter protests occurred near the end of our sample period, and they may have had a differential effect on police openness, particularly in cities where protesters were most active.

References

Abney, G. and T. P. Lauth (1986). *The politics of state and city administration.* SUNY Press.

Ackerman, J. M. and I. E. Sandoval-Ballesteros (2006). The global explosion of freedom of information laws. *Administrative Law Review 58*, 85–130.

Banerjee, A., N. Enevoldsen, R. Pande, and M. Walton (2024). Public information is an incentive for politicians: Experimental evidence from delhi elections. *American Economic Journal: Applied Economics 16*(3), 323–353.

Banerjee, A., S. Kumar, R. Pande, and F. Su (2011). Do informed voters make better choices? experimental evidence from urban india. *Unpublished Manuscript, Harvard University*.

Berliner, D. (2014). The political origins of transparency. *The Journal of Politics 76*(2), 479–491.

Berliner, D., B. E. Bagozzi, B. Palmer-Rubin, and A. Erlich (2021). The political logic of government disclosure: Evidence from information requests in mexico. *The Journal of Politics 83*(1), 229–245.

Berliner, D. and A. Erlich (2015). Competing for transparency: political competition and institutional reform in mexican states. *American Political Science Review 109*(1), 110–128.

Besley, T. and A. Prat (2006). Handcuffs for the grabbing hand? media capture and government accountability. *The American Economic Review 96*(3), 720–736.

Brierley, S., K. Lowande, R. A. Potter, and G. Toral (2023). Bureaucratic politics: Blind spots and opportunities in political science. *Annual Review of Political Science 26*(1), 271–290.

Camerlo, M. and A. Pérez-Liñán (2015). Minister turnover, critical events, and the electoral calendar in presidential democracies. *The Journal of Politics 77*(3), 608–619.

Campante, F. R. and Q.-A. Do (2014). Isolated capital cities, accountability, and corruption: Evidence from us states. *The American Economic Review 104*(8), 2456–81.

Cordis, A. S. and P. L. Warren (2014). Sunshine as disinfectant: The effect of state freedom of information act laws on public corruption. *Journal of Public Economics 115*, 18–36.

Costa, S. (2013). Do freedom of information laws decrease corruption? *The Journal of Law, Economics, & Organization 29*(6), 1317–1343.

de Benedictis-Kessner, J., M. Harvey, D. Jones, and C. Warshaw (2025). The partisanship of mayors has no detectable effect on police spending, police employment, crime, or arrests. *Science Advances 11*(3).

de Benedictis-Kessner, J., D. D. I. Lee, Y. R. Velez, and C. Warshaw (2023). American local government elections database. *Scientific Data 10*(1), 912.

Dipoppa, G. and G. Grossman (2020). The effect of election proximity on government responsiveness and citizens' participation: Evidence from english local elections. *Comparative Political Studies 53*(14), 2183–2212.

Enikolopov, R. (2014). Politicians, bureaucrats and targeted redistribution. *Journal of Public Economics 120*, 74–83.

Ferejohn, J. (1999). Accountability and authority: toward a theory of political accountability. *Democracy, Accountability, and Representation 2*, 131–150.

Ferraz, C. and F. Finan (2008). Exposing corrupt politicians: The effects of brazil's publicly released audits on electoral outcomes. *The Quarterly Journal of Economics 123*(2), 703–745.

Fisman, R., F. Schulz, and V. Vig (2020). Financial disclosure and political selection: Evidence from india. *Unpublished Manuscript, Boston University*.

Grimmelikhuijsen, S., G. Porumbescu, B. Hong, and T. Im (2013). The effect of transparency on trust in government: A cross-national comparative experiment. *Public Administration Review 73*(4), 575–586.

Grossman, G., K. Michelitch, and C. Prato (2024). The effect of sustained transparency on electoral accountability. *American Journal of Political Science 68*(3), 1022–1040.

Hallward-Driemeier, M. and L. Pritchett (2015). How business is done in the developing world: Deals versus rules. *Journal of Economic Perspectives 29*(3), 121–140.

Hamel, B. T. (2025). Inequality in the classroom: Electoral incentives and the distribution of local education spending. *The Journal of Politics 87*(2), 527–543.

Honig, D., R. Lall, and B. C. Parks (2023). When does transparency improve institutional performance? evidence from 20,000 projects in 183 countries. *American Journal of Political Science 67*(4), 1096–1116.

Hood, C. (2006). Transparency in Historical Perspective. In *Transparency: The Key to Better Governance?* British Academy.

Jenkins, N. R., M. Landgrave, and G. E. Martinez (2020). Do political donors have greater access to government officials? evidence from a foia field experiment with us municipalities. *Journal of Behavioral Public Administration 3*(2).

Kreimer, S. F. (2008). The freedom of information act and the ecology of transparency. *University of Pennsylvania Journal of Constitutional Law 10*, 1011.

Lagunes, P. and O. Pocasangre (2019). Dynamic transparency: An audit of mexico's freedom of information act. *Public Administration 97*(1), 162–176.

Leip, D. (2024). United states governor election data.

Levitt, S. D. (1997). Using electoral cycles in police hiring to estimate the effect of police on crime. *The American Economic Review 87*(3), 270–290.

Malesky, E., P. Schuler, and A. Tran (2012). The adverse effects of sunshine: a field experiment on legislative transparency in an authoritarian assembly. *American Political Science Review 106*(4), 762–786.

Malik, R. (2020). Transparency, elections, and pakistani politicians' tax compliance. *Comparative Political Studies 53*(7), 1060–1091.

Rose-Ackerman, S. and B. J. Palifka (2016). *Corruption and Government: Causes, Consequences, and Reform* (2 ed.). Cambridge: Cambridge University Press.

Shaukat, M. (2019). Too close to call: Electoral competition and politician behavior in india. *Massachusetts Institute of Technology Working Paper*.

Stephenson, M. C. (2008). Optimal political control of the bureaucracy. *Mich. L. Rev. 107*, 53.

The Council of State Governments (2021). *The Book of the States* (53 ed.). Lexington, KY: The Council of State Governments.

Thompson, D. M. (2020). How partisan is local law enforcement? evidence from sheriff cooperation with immigration authorities. *American Political Science Review 114*(1), 222–236.

Trautendorfer, J., L. Schmidthuber, and D. Hilgers (2023). Are the answers all out there? investigating citizen information requests in the haze of bureaucratic responsiveness. *Governance*.

Vlaicu, R. and A. Whalley (2016). Hierarchical accountability in government. *Journal of Public Economics 134*, 85–99.

Wade, R. (1982). The system of administrative and political corruption: Canal irrigation in south india. *The Journal of Development Studies 18*(3), 287–328.

Wood, B. D. and R. W. Waterman (1991). The dynamics of political control of the bureaucracy. *American Political Science Review 85*(3), 801–828.