Electoral Incentives and Government Transparency: Evidence from Freedom of Information Requests

Raymond Fisman, Aron Malatinszky, and Eyub Yegen*

May 21, 2025

Abstract

Government transparency is a basic requirement of government accountability, allowing citizens access to the information necessary to monitor, reward, and punish politicians and bureaucrats. Yet transparency inevitably involves discretion by the same government that is under evaluation. In this paper, we show that reelection incentives may reduce transparency, by comparing responsiveness of freedom-of-information requests made via the online platform MuckRock before versus after state and municipal elections. 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 or shortly after an election. We first document that our "revealed transparency" measure is positively correlated with a standard measure of state-level corruption. Our main analyses show that revealed transparency in state agencies is lower ahead of state-level elections in high-corruption settings, and we observe a similar pattern for municipal agencies ahead of mayoral elections; in contrast, we find the opposite relationship in lower-corruption states. Overall, our results suggest that discretion may be used to limit access to information that is relevant to elections.

JEL Classifications: D72, D73, H70

Key words: Government Transparency; Sunshine Laws; FOIA; Corruption

*We are grateful to very helpful comments from Elliott Ash and Daniel Berliner. We also thank to seminar participants at the University of Chicago, University of Toronto, Hong Kong University of Science and Technology, and the Global Integrity Workshop for invaluable feedback. Fisman: Boston University and NBER (email: rfisman@bu.edu); Malatinszky: Boston University (email: aronm@bu.edu); Yegen: Hong Kong University of Science and Technology (email: yegen@ust.hk).

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 to 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. This may account for 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] Perhaps more importantly, FOIA statutes also leave enormous room for discretion, in terms of delaying fulfillment of a FOIA request for sensitive material, or denying the request outright (Kreimer, 2007). 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

[1] Globally, 119 countries have freedom of information laws according to freedominfo.org, accessed January 15, 2025. For simplicity and to align with past work, we refer throughout to rules at both the state and federal level as FOIA laws, though it is more accurate to use FOIA to describe rules governing federal agencies and freedom of information law 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 January 15, 2025.

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

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 Virginia media outlets found that FOIA requests on topics − such as felony data or government salaries − that are clearly covered under the state's FOI statutes were nonetheless routinely ignored by government officials.[5]

Discretion can potentially lead to less transparency precisely when it is most valuable for 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. Our primary objective in this paper is to evaluate the extent to which FOIA responsiveness is affected by such electoral incentives.

Since FOIA laws are largely time-invariant, 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 the submission of FOIA requests. Details on these requests are all 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. This feature 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 constant all other aspects of the request.

[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, 2015. Available at https://www.dailypress.com/government/dp-nws-foia-project-state-responses-20151128-story.html, last accessed January 15, 2025.

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 malfeasance. We show that the likelihood of FOIA request rejection in general is highly correlated with state-level corruption, as captured by federal prosecutions for illegal use of public office (Campante and Do, 2014). This pattern is observed whether we consider the full set of 35,358 state and local FOIA requests that are no longer open, 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 ahead of an election. Alternatively, a well-intentioned public servant may be particularly well-motivated to provide timely disclosures in advance of elections, and 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 a FOIA request and those targeted by it 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

[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 Albermarle Prosecutor over FOIA Response," *Charlottesville Daily Progress*, October 29, 2015.

overall. We further argue that the ability to exercise discretion − as well as the likelihood that officials do in fact 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), who show that information on politician performance has negative electoral consequences *only* in cases in which 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 jurisdictions with higher corruption.

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; the sign is actually the reverse in the lowest-corruption states. In other words, rejection and delay both go down in low-corruption settings. We observe these patterns for the combined dataset that includes both the impact of gubernatorial elections on the responsiveness of state-level agencies as well as mayoral elections and municipal agency responsiveness: separating the data into state and municipal agencies, we observe the same patterns in both subgroups.

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 in which there is an appeals 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 greater 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 show 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 in accessing such information precisely when it is helpful to the public in holding politicians accountable.

Given these 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 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 political opposition (as captured by the runner-up party's vote share) and turnover (as captured by frequency of executive turnover in the previous five years) 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

[7]Malesky et al. (2012), however, suggests 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.

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 exist.

Our study also contributes to 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 are associated with an *increase* in perceived corruption. Naturally, rather than reflecting an increase in actual corruption, this pattern could very plausibly 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 prosecutions, 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 prosecutions 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 prosecutions 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 also emphasize that corruption itself may act as a mediating factor in when and how freedom of information laws are deployed.

Finally, ours is one of several studies that take 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] Finally, we also see our approach as offering a relatively straightforward method to 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, concluding 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 group-matched approach) across jurisdictions or circumstance 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. While, as we discuss below, our initial results are limited by the current stock of requests on MuckRock, we also demonstrate the rapidly-growing potential of utilizing this and other platforms as research tools for studying government transparency.

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

[8]For example, a set of requests we describe below made queries about the weapons arsenals of 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) include fixed-effects by topic-agency-year which, while ruling out many concerns of omitted variables, leaves considerable residual within-group variation.

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 MuckRock offers 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.[10] In its first year, 219 requests were placed; request volume grew steadily for the next few years, and hit a steady 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.

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 status over time as appropriate.[11] In our main analyses, we focus on the set of requests that MuckRock codes as having received a clear resolution − those that are marked by MuckRock as "completed" (29,531 of the state and local requests), and those that end in rejection (6,058 of state and local requests). 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

[10]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 via MuckRock that year was 9,862.

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

the small number of requests from the District of Columbia and U.S. territorial governments (212 and 6, respectively), yielding a sample of 35,371 requests.

Not all requests are resolved in this binary manner. The full set of designations include: withdrawn by the submitter (4,022), awaiting acknowledgment from the agency (2,422), submitted and awaiting response (12), being processed by the agency (2,282), awaiting appeal to agency (279), fix required by submitter (5,063), in litigation (104), payment required by submitter (1,622), request was only partially completed (1,020), and no responsive documents provided by the agency (15,395). Most of these categories are ambiguous in their interpretation. While about 37.2% 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 by the requester; 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." This arises most commonly 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 "no relevant documents" response often reflects some combination of willful obstruction, laziness, and incompetence (technological or otherwise).[12] MuckRock defines a request as closed if it is Completed, Partially Completed,

[12]See https://www.muckrock.com/news/archives/2018/sep/06/foia-faq-nrd-wtf/; last accessed January 15, 2025. No responsive documents can also be an interim classification if the submitter requests follow-up from the petitioned agency.

Rejected, or No Responsive Documents; all other intermediate categories are potentially still in-progress.[13]

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.[14] 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. There are also other rationales, 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 citizens of the state, 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 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

[13]We also perform a set of robustness checks in which we use all closed files, classifying "no relevant documents" as rejections and partial completions as complete, though as noted above, there are entirely legitimate reasons that an agency may not have any relevant documents available. Please see Section 3.2.1.

[14]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 is for information from the Federal Bureau of Prisons rather than a state-level department.

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.[15] 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 identical.[16] For our main analysis, which includes only completed and rejected requests, there are 13,604 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.

For this subgroup, we provide in Figure 1 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 resolved requests. We make two observations about these small- and large-group requests. First, we note that we generate very similar results when we omit relatively small groupings (e.g., fewer than five requests per group; see Appendix Table A5). Second, and perhaps more substantively, larger groupings may introduce some within-group heterogeneity, most obviously in the petitioned agency. For example, identical requests are sometimes made to district attorney's offices, public safety agencies, as well as city police

[15]We use an algorithm to construct the matched groups by matching on filer and title.

[16]We manually verify all matched groups by randomly selecting at least five FOIA requests (or all FOIA requests if the number of observations in a given group is smaller) within the matched sample. We verify that identical titles filed by the same person are FOIA requests with identical or near-identical contents.

Figure 1: Frequency of Identical FOIA Requests within Matched Group

This figure plots the frequency of FOIA requests that are within a given group of identical matches (i.e., same content) submitted by the same person to various agencies.

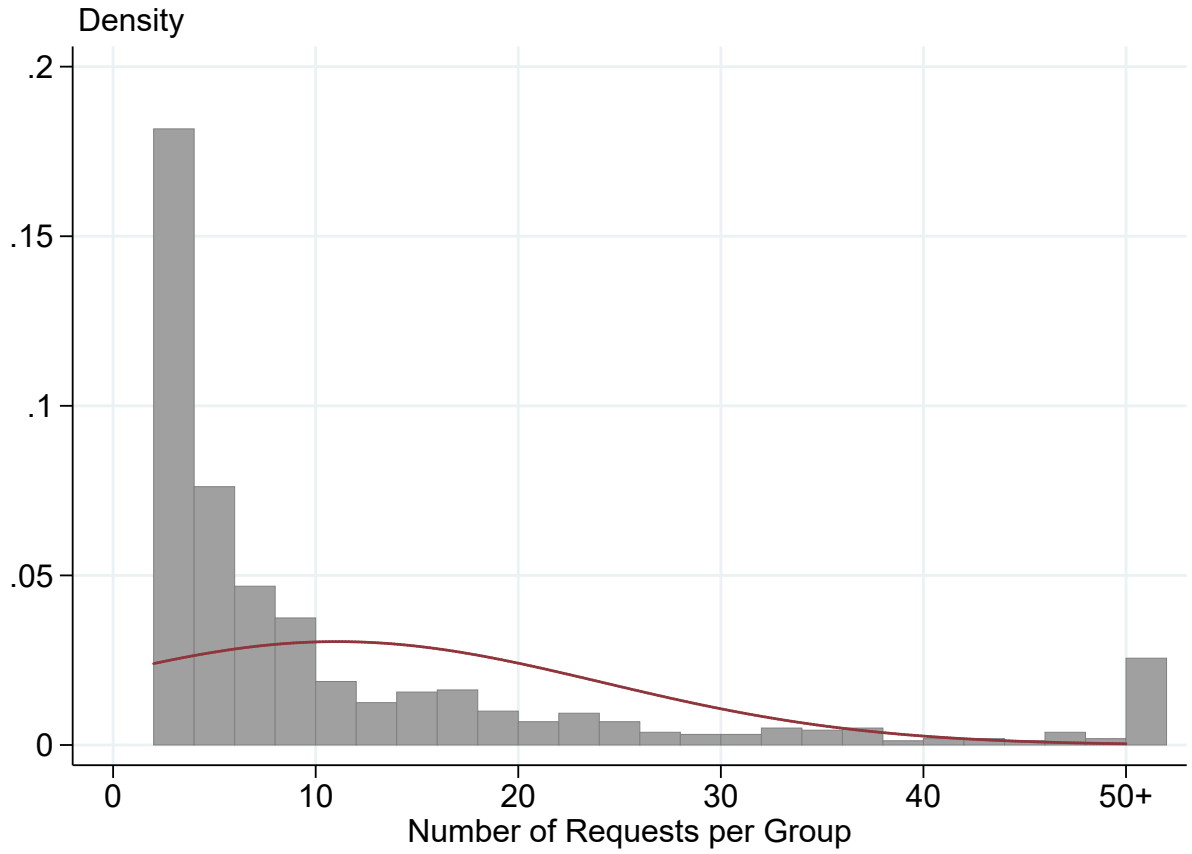

departments; or individual state agencies as well as the governor's office. While it is not obvious why respondent attributes would be correlated with any geographic or temporal variation that we explore below, we nonetheless include a robustness check in which we look exclusively at requests of policing agencies (over 40 percent of our matched-group sample), and generate very similar findings to those in our main results; see Appendix Table A4.

To measure state-level corruption, we utilize the measure from Campante and Do (2014), which captures the number of federal prosecutions per capita of state and local officials during 1976-1992. This period predates our data, so that 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. 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.[17] The Reporters Committee for Freedom of the Press compiles and publishes these data and we follow Cordis and Warren (2014) to calculate summary FOIA law scores.[18]

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 4,059 cities whose agencies appear in the MuckRock data, we observe municipal elections in 468 cities.[19] 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 2012 and 2021 recall elections in Wisconsin and California, respectively (see, e.g., Leip, 2024).

To construct the election-level dataset for the 4,275 completed or rejected matched-group filings in these 468 cities and 50 states holding 1,658 elections, we use the following steps.

[17]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.

[18]It is also natural to control for the size of government, though in practice these proxies vary little across states, in particular the number of government employees by state. 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.

[19]ALGED provides results from contests decided by direct election and 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.

First, from our list of matched-sample groupings of FOIA requests, we select groups for which at least one request took place within six months 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*, that denotes whether there is a gubernatorial or mayoral election in the request's state or city in the next six months. Since our data on elections extends through 2021, we restrict to the set of FOIA requests that were submitted at least six months before the end of 2021. Obviously, there is nothing special about the six-month window, so will consider alternate windows ranging from four to twelve months in our empirical analysis.

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. (A natural concern is that the volume of citizen queries and complaints may differ around elections (Dipoppa and Grossman, 2020). In practice in our data, 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. Please see Appendix B for these analyses.) 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 months between a FOIA request filing date and an election in the same locale, and we identify a request as $Unfilled$ if it is unresolved within the longest election window. In the example above, we would thus assign $Unfilled = 1$ 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 $Unfilled$ (in the above example, three months).

We combine our rejection and failure measures to capture whether a request received timely and satisfactory response to generate the variable $Failure$, which denotes requests for which $Unfilled = 1$ or $Rejected = 1$. We will also look at $Unfilled$ and $Rejected$ separately as outcomes.[20] The sample for this set of analyses comprises 2,668 requests submitted to city agencies and 1,607 requests submitted to state agencies in 271 matched groupings.

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 some summary statistics for the FOIA data; Panel B provides summary statistics for our corruption measure and for the control variables. Across all agencies, among requests that received unambiguous resolution as either rejected or fully fulfilled, 21% were rejected.[21] There is considerable heterogeneity across levels of government. The rejection rate for federal agency FOIAs is nearly double that of municipal agencies.

[20]Note that $Unfilled$ and $Rejected$ 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 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.

[21]When we define rejection as either outright rejection or the agency reporting that no documents are available, and include also partially fulfilled 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

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 local government agencies with variable exposure to gubernatorial and mayoral elections used in our main analysis. Requests are marked as completed if MuckRock records them as having received a clear resolution (either coded by MuckRock as "completed" or "rejected"). For state-level variables, we use the per capita rate of corruption convictions during 1976-1992 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						
Completed	0.502	0.500	99,396	1.000	0.000	4,275
Rejected if completed	0.207	0.405	49,892	0.173	0.378	4,275
No responsive documents	0.228	0.420	99,396	0.000	0.000	4,275
Filing per journalist	23.71	168.87	4,193	32.88	78.25	130
Federal Agencies						
Completed	0.453	0.498	31,586			
Rejected if completed	0.297	0.457	14,303			
No responsive documents	0.231	0.422	31,586			
Filing per journalist	15.64	111.17	2,020			
State Agencies						
Completed	0.515	0.500	20,556	1.000	0.000	1,607
Rejected if completed	0.213	0.409	10,591	0.200	0.400	1,607
No responsive documents	0.240	0.427	20,556	0.000	0.000	1,607
Filing per journalist	11.81	54.60	1,740	17.66	35.61	91
City Agencies						
Completed	0.529	0.499	47,254	1.000	0.000	2,668
Rejected if completed	0.152	0.359	24,998	0.156	0.363	2,668
No responsive documents	0.221	0.415	47,254	0.000	0.000	2,668
Filing per journalist	18.11	118.05	2,609	24.70	61.94	108
Panel B: State-Level Variables						
Corruption score	0.275	0.132	50			
State FOIA law 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, the 10 states with the highest rejection rates to the 10 with the lowest rejection rates, the latter has notably higher corruption as captured by per capita corruption convictions (0.32 versus 0.19). We will explore these patterns further in a regression framework − also including city agencies within each state − in the next section.

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 the 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 2,128 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 our inability to compare our sample to the universe of FOIA requests may limit the extent to which we may make generalization based on our results, it is less clearly an issue for the internal validity of our findings, which are based on matched-groupings of identical requests. That is, within the corpus of documents we focus on, we are making comparisons among requests that are the same apart from the agency with which they are filed. In our first set of results, we identify the role of jurisdiction via a cross-state comparison of otherwise identical requests; in our second set of results, identification is based on a combination of differences in the timing of elections (for the direct effect) as well as cross-jurisdictional differences (for whether timing matters differently as a function of state-level corruption).

A further 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.[22] 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.

[22]Another matched-group example makes this clear. A 2013 request from George LeVines, at the time a Massachusetts-based journalist, 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 did provide such inventories were in fact being *too* transparent. Note that because all of these requests were made within Massachusetts, it is not included in our cross-jurisdictional analyses.

3 Results

We present our findings in two parts. First, we show the relationship between FOIA compliance and state-level corruption. We see this as a potentially interesting finding in itself, and one that will also allow us to contrast how our "revealed transparency" measure versus formal FOIA rules correlate with corruption. We will then present our main findings on electoral timing and FOIA compliance, linking responsiveness to concerns of political accountability.

3.1 Relationship between state-level corruption and revealed transparency

To examine the correlation between corruption and revealed transparency, we begin by presenting 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:

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

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 1976-1992. 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

Table 2: State Level Corruption and FOIA Responsiveness

This table reports the relationship between state level corruption and the average FOIA rejection rate by the given state. Columns 1 to 4 report the results using state-year level data, whereas columns 5 and 6 report the results using data at the state level. We require that the given FOIA filing is either rejected or accepted (i.e., removing, for instance, ongoing or appealed cases). 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

requests we have found to be a very common reason for rejection.[23] 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 quite stable around 0.25-0.35 (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

[23]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.

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%.

Figure 2: 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.

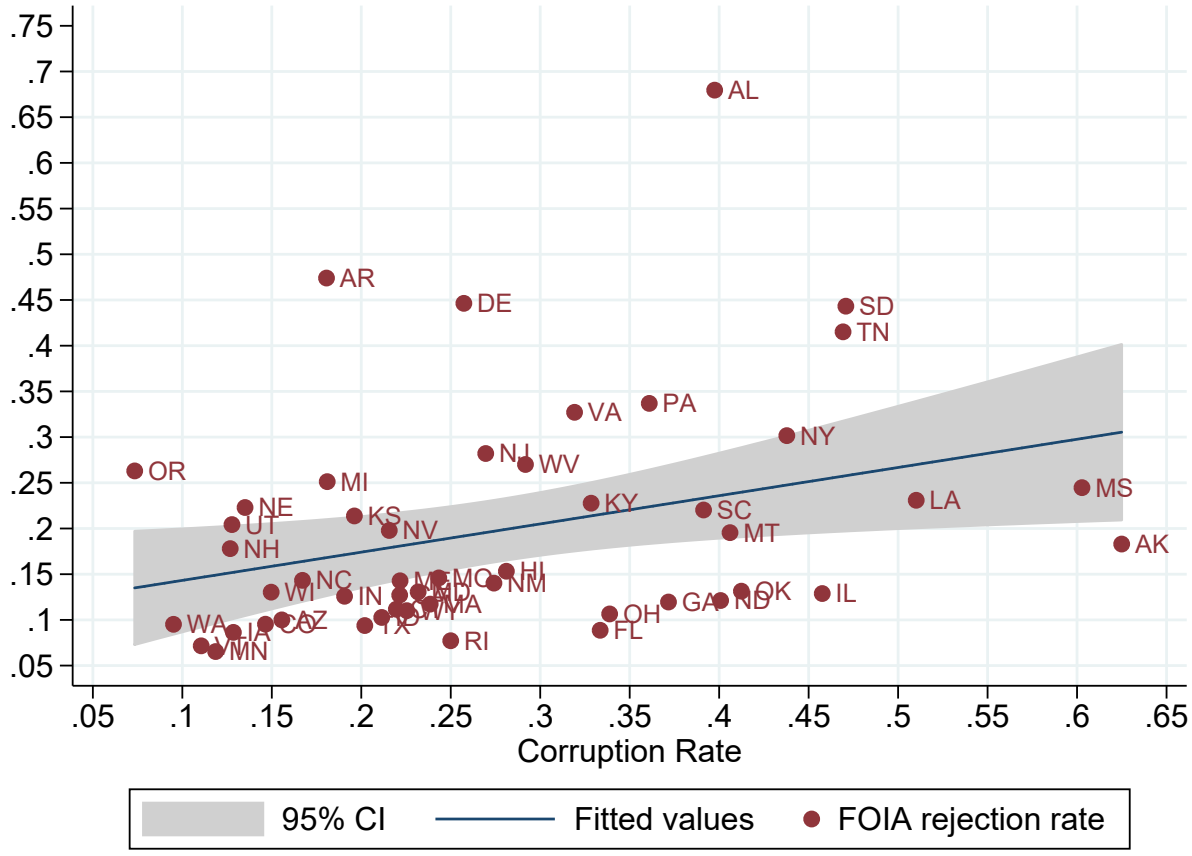

In Figure 2 we illustrate the relationship between *Corruption Rate* and *Rejection Rate*. No individual outlier is driving the result, and the relationship is approximately linear. In Appendix Table A2, we use an alternative measure of rejection rate, in which we include partial completions as completed requests, and "no document" responses as rejections. We obtain similar results to those in our main analysis.

Finally, in Appendix Table A3 we repeat our analysis based on equation (1) using instead the natural log of average response time as the outcome. We observe no significant relationship between the time to fill a request and any state-level attribute.[24]

Naturally, there are many differences across states that might influence both FOIA success rates as well as the types of requests that are filed. 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. Let us 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:

$$Rejection_{rgst} = \alpha + \beta \times Corruption\ Rate_s + X_{st} + \gamma_g + \epsilon_{rgst} \tag{2}$$

where *Rejection* is an indicator variable denoting that request r from matched group g filed in state s in year t was rejected. In all specifications, we use two-way clustering − again by state, and also by grouping of matched requests.[25] In our most stringent specifications, we include 800 group fixed effects. We emphasize that when we use this approach, we identify the relationship between corruption rate and FOIA rejections, holding constant (by construction) the characteristics of FOIA requests and submitters.

[24]While we focus on the subset of requests that we use in Table 2, if we include all requests, we similarly observe no relationship between fulfillment time and any state-level variable.

[25]We obtain essentially the same results when we use three-way clustering by state, by grouping of matched requests, and by quarter of submission.

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 *Rejection*, an indicator variable denoting that a request was rejected. The sample mean of *Rejection* 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: *Rejection Rate*					
	(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

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 of Table 2. The patterns are comparable to those based on the preceding cross-sectional analyses, though the point estimates are marginally higher in magnitude. In column 4 we repeat the specification of column 3, but include only 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. It is noteworthy that *Average FOIA Score* is uncorrelated with rejection rates and, indeed, we note is uncorrelated with *Corruption Rate* as well.

In the appendix material we present one further robustness check. To address concerns that our matched groups may contain heterogeneity across types of agencies, we focus on the subsample of requests that were made to policing agencies. They constitute approximately 40 percent of our overall sample; we see this as a simple and transparent way of focusing on cases with minimal within-group heterogeneity in agency type. These results appear in Appendix Table A4 for our main sample. The results are quite similar to those in our results that include all agencies.

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 is associated with a government's willingness to be held up to scrutiny. We see this as an interesting fact in itself, but also as a basic reality check on the data. Our matched grouping results bolster the credibility of the basic relationship between corruption and revealed transparency. Additionally, the matched grouping approach hints at the possibility that we may use the structure of the data to explore the determinants of transparency across any source of vari-

ation that exists within our groups − not simply across-state variation. 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 cities and over time.

3.2 Electoral pressure and revealed transparency

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. Most obviously, officials may be less apt to respond to FOIA requests if the resulting revelations risk political embarrassment. In the other direction, a well-meaning civil servant may be particularly attentive in providing timely disclosures to voters ahead of elections; and it is also possible electoral pressures may lead to greater responsiveness, lest opacity become an election issue in itself.

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 six months, and at least one request for a different state or municipal agency where there is *not* an election in the next six months as a benchmark; we will also consider shorter and longer pre-election windows, from four months to twelve months.

As discussed in Section 2, there are two ways of avoiding a pre-election FOIA response: delay ($Unfilled$), or outright rejection ($Rejection$). Our primary measure, $Failure$, combines both of these.

In our main specification, we explore the direct relationship between an upcoming election and $Failure$, and also consider 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:

$$Failure_{rgbt} = \alpha + \beta_1 \times Election_{rgbt} + \beta_2 \times Corruption\ Rate_{s(b)} + \beta_3 \times Election_{rgbt} \times Corruption\ Rate_{s(b)} + \gamma_g + \epsilon_{rgbt} \quad (3)$$

for request r that is a part of group g, submitted an an agency in government b at time t. As before, 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 also by state.[26]

Results based on equation (3) appear in Table 4. We first focus on our combined measure of failure, capturing whether a request is rejected or was not resolved before the pertinent election. 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 -0.015, and does not approach statistical significance. The standard error of 0.027 indicates that we can reject, at the 5% level, an effect size smaller than -0.07 or larger than 0.04. Thus, on average there is no substantial difference in request success as a function of election timing. Comfortingly, for this subsample of requests we again find a positive coefficient on *Corruption Rate*, indicating a higher failure rate in states with more corruption prosecutions per government employee.

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. To explore this prediction, we add the term *Election*×*Corruption Rate* in column 2. The coefficient on this interaction is

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

0.56 (significant at the 1% level), and implies a "crossing point" of *Corruption Rate* $= 0.28$, where the *Election* effect is zero, at approximately the 60th percentile of the distribution of *Corruption Rate*. The third column of Table 4 adds state fixed effects. We can no longer identify a direct effect of *Corruption Rate*, but the main coefficient of interest, on *Election* $\times$ *Corruption Rate*, is relatively unaffected by this inclusion.

Table 4: 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 six months 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 𝟙(Election) takes the value of one if the FOIA request was filed to a government agency in a jurisdiction that had an election six months prior to the election date. The mean of the dependent variable, 𝟙(Failure), is 0.32, and the mean of 𝟙(Rejected) and 𝟙(No Decision) are 0.17 and 0.19, 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								
	𝟙(Failure)			𝟙(Rejected)			𝟙(No Decision)		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
𝟙(Election)	-0.015	-0.158***	-0.175***	-0.006	-0.089**	-0.115***	-0.011	-0.054	-0.058
	(0.027)	(0.052)	(0.045)	(0.017)	(0.038)	(0.029)	(0.021)	(0.052)	(0.045)
Corruption Rate	0.376*	0.251		0.414**	0.342**		0.034	-0.003	
	(0.195)	(0.192)		(0.165)	(0.163)		(0.130)	(0.137)	
𝟙(Election) × *Corruption Rate*		0.557***	0.604***		0.323***	0.409***		0.165	0.166
		(0.139)	(0.123)		(0.113)	(0.089)		(0.164)	(0.136)
Fixed Effects									
Matched-Group	X	X	X	X	X	X	X	X	X
Quarter of Submission	X	X	X	X	X	X	X	X	X
State			X			X			X
N	4,273	4,273	4,273	4,273	4,273	4,273	4,273	4,273	4,273
R^2	0.247	0.250	0.302	0.306	0.308	0.388	0.167	0.168	0.209

The remaining columns of Table 4 look separately at the two components of $Failure$, i.e., $Rejection$ and $Unfilled$. While the coefficients are of consistent signs for both sets of results, it is clear that our main results are driven primarily by $Rejection$. If we were to focus on this more readily-interpretable outcome, it provides a very similar message to what we see for the combined $Failure$ variable.

3.2.1 Heterogeneity and Robustness of Election Results

We conclude our analysis by exploring the robustness of our results and also whether the patterns we document in the preceding section vary by geographic or political characteristics. Throughout, we focus primarily on $Unfilled$ as our summary measure of FOIA responsiveness. The main takeaway from these analyses will be that our findings are remarkably stable across a range of alternative specifications and sample splits.

We begin by examining the sensitivity of our main results to definitions of the pre-election window. We consider somewhat shorter windows of as little as four months, as well as longer windows of up to twelve months. The short-window results appear in Appendix Table A6 while the long-window results are in Appendix Table A7. Reassuringly, the patterns we observe based on a six-month window are largely unaffected by using a shorter or longer window; in particular, the coefficient on $Election \times Corruption\ Rate$ is significant at least at the 10% level in all specifications, and consistently in the range of $0.30 - 0.55$.

The other modeling choice we make in generating our main results is using the longest election window in each group to set the uniform deadline in defining whether a request is filled before the election. In Appendix Table A8, we instead use the shortest election window. As we observed above, this difference has a potentially large impact on our results, since we have many groupings with multiple election dates within the six-month window. However, as we see in Appendix Table A8, our findings are essentially unchanged when we use the shortest election window. (The fact that our results are unaffected is perhaps less surprising ex post, since we saw in Table 4 that our main results are driven primarily by rejections

rather than delays, and the definition of rejection is not affected by whether we use later or earlier election dates to define pre-election fulfillment.)

In Appendix Table A9, 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 $Election \times Corruption\ Rate$ 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).

In a similar spirit to our earlier robustness check to examine whether agency heterogeneity might account for our results, in Appendix Table A10 we distinguish between police versus non-police requests. We observe a significant interaction in both groups.

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 A11 show that the coefficient of primary interest is quite similar for the two groups.

As a final heterogeneity test, 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 fully resolve (Shaukat, 2019). A more immediate concern is that there are relatively few close elections − a total of 94 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 A12. None of the coefficients in any specification approaches statistical significance in these results.

Finally, in Appendix Table A13, we show our main results for the more expansive definitions of rejection (including "no responsive documents") and completion (including "partially completed"). The patterns are very similar to those reported in Table 4, though the point estimates are smaller, perhaps as expected if, for example, many requests now coded as

rejections are false positives.

4 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 exploit the collection 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. And 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 our "revealed transparency" measure predicts state-level corruption prosecutions, 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 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, as well as 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, as a way of classifying requests as more or less sensitive, 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 in which protesters were most active.

References

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

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.

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., 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.

Ferejohn, J. (1999). Accountability and authority: toward a theory of political accountability. *Democracy, Accountability, and Representation 131*, 133.

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.

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, G. E. Martinez, et al. (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. (2007). 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.

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.

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

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*.