Measuring Government Openness Through Freedom of Information Requests

Raymond Fisman and Eyub Yegen*

November 25, 2023

Abstract

We present a new "revealed policy" measure of government openness, based on realized responses to over 25,000 freedom-of-information requests made via the online platform MuckRock. Since many identical requests are made at the same time to many agencies across the U.S., we may assess whether an agency rejects a request, holding constant the timing, submitter, and request content. To demonstrate the validity and potential usefulness of our approach, we present two applications: we show that "revealed transparency" is positively correlated with a standard measure of state-level corruption, and explore how responsiveness to requests is affected ahead of local elections. In comparing the correlates of rejection rates to that of standard measures of legal stringency, our results also suggest that enforcement of freedom of information laws is more important than the letter of the law itself.

JEL Classifications: D72, D73, H70

Key words: Freedom of Information; Government Transparency; Sunshine Laws; 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); Yegen: Department of Finance, 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 give citizens access to the information they need to decide which politicians to reelect and which to vote out of office (Malik, 2020), limit rent-seeking (Fisman et al., 2020), improve bureaucratic performance (Honig et al., 2022), and bolster the legitimacy of government (Grimmelikhuijsen et al., 2013). That is, as Justice Louis Brandeis famously observed, "sunlight is said to be the best of disinfectants," as "publicity is...a remedy for social and industrial diseases."

Such promises of 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]

Not all sunshine is created equal, however. As observed by Cordis and Warren (2014), there is substantial variation across states in the stringency of laws. For example, as of 2021, eight states require a response to most requests within three days, while in thirteen states there is no prescribed time; the fees charged in Oregon cannot exceed \$25 without prior notification, while in Alabama the state's FOIA law allows for the charging of material costs (such as photocopying) as well as staff time, with no limit.

Previous attempts at comparing government openness across jurisdictions have focused on FOIA laws as written (Cordis and Warren, 2014; Harden and Kirkland, 2021). However, 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). And as has been noted by, in particular, Hallward-Driemeier and Lant Cordis (2015), there may be

[1]Globally, 119 countries have freedom of information laws according to freedominfo.org, accessed August 29, 2022. Also note that 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 FOIL (freedom of information law) to describe state-specific statutes.

a large gap between the stringency of written laws and their enforcement in practice.[2] 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."[3] 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.[4]

In this paper, we use a "revealed transparency" approach to measuring government openness and accountability via an analysis of over 25,000 FOIA requests made during 2010-2022, across the 50 U.S. states (for our matched-group analysis, we use a subset of just over 6,000 of these requests). The data come from MuckRock.com, a website that facilitates the submission of FOIA requests via its Application Programming Interface (API). The website then provides information on a range of outcomes, including the time to receive a response and whether the request was rejected (generally coded by MuckRock itself). The MuckRock 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 empirics, comparing the outcomes of identical FOIA requests made by the same submitter around the same date to comparable agencies, but done across various jurisdictions. By comparison, while open government advocates have conducted FOIA "audits" they have generally limited themselves to making only a few requests in each state, whereas we have quasi-experiments comprised of thousands of FOIA filings in all, which are externally valid in the sense of being made with the actual purpose of obtaining government records. This kind of apples-to-apples comparison is essential, given that differing compliance norms may themselves lead to a different

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

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

[4] "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 March 1, 2023.

composition of requests (e.g., more challenging requests may only be sent in more transparent locales).

We present two sets of results to illustrate the validity and potential usefulness of our transparency measure.

First, we show that state-level corruption – as captured by federal prosecutions for illegal use of public office (Campante and Do, 2014) – is highly correlated with the likelihood that a FOIA request is rejected. This pattern is observed whether we consider the full set of FOIA requests, or the subset of 10,892 requests that are part of a "matched group." We see this result as a basic reality check of the data, and also indicating that states that have more to hide are more apt 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).

Second, we next provide an illustrative example to show how matched-groupings of FOIA requests might be applied to generate measures of government openness at more disaggregated levels – by municipality or even individual government agencies – a possibility that is harder or perhaps impossible to implement for openness as measured by legal statutes. Specifically, we examine how local elections affect FOIA responsiveness, which serves to highlight both the application of our measure at a disaggregated level as well as the fact that "revealed openness" may be measured with relatively high temporal frequency. To do so, we look at groupings of FOIA requests that occur, for at least one city in the group, shortly before a mayoral election. The point of comparison includes identical FOIA requests made at the same time in cities that do not, at the time, have elections in the near future. To our knowledge, we are the first to examine empirically the link between electoral cycles and compliance with transparency laws.

The effect of upcoming elections on FOIA compliance is decidedly ambiguous in theory. Most obviously, incumbents may wish to avoid the release of embarrassing information ahead of an election. Alternatively, failure to comply with a request may itself be the source of campaign fodder.[5] 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 our reduced form, we will assess whether these effects collectively lead to more or less pre-election transparency overall, and whether it differs based on our state-level corruption measure that, we argue, may reflect ability to delay or reject responses.

We find that, on average, there is a modest improvement in responsiveness around elections – both rejection probability and the probability of lengthy delay are lower. However, these relationships depend on the broader institutional environment – pre-election responses are relatively slow (and more likely to result in rejection) in high-corruption states, and are relatively quicker (with lower rejection probability) in low-corruption states. We reiterate that we do not see these findings as testing any particular theory of electoral accountability but rather demonstrating the flexibility and potential uses of our measure.

Our work contributes to the broad literature in economics and political science on the causes and effects of government transparency. 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 of which also document an empirical connection between a strong media presence and government responsiveness. The potential of transparency to improve bureaucratic functioning and affect electoral outcomes has been demonstrated in Honig et al. (2022), which shows that development projects yield better outcomes after the donor agency implements an access to information policy. However, their results 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.

[5]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" amidst 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.

A pair of recent studies highlight the greater scrutiny of politicians – and resultant change in behavior – that accompanies greater transparency, by experimentally manipulating information provision. In particular, Grossman et al. (2020) 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. (2023) further show that when Delhi legislators are informed that report cards will be provided to voters, it leads to performance improvements.[6]

While these results suggest that this is in the interests of voters, it does not necessarily serve the interests of the public officials who face greater scrutiny, which naturally raises the question 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 information laws is predicted by the extent of political competition.

Our work connects most directly to two themes in prior work. First, given the focus on corruption in our first set of results, our findings are of relevance to past research that links 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

[6]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.

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 Lant Cordis, 2015), and also in our approach to identification, which exploits quasi-experiments rather than cross-state or cross-country variation in written laws.

We speak most directly to the growing body of empirical evidence on the determinants of compliance with freedom of information laws. A number of recent papers utilize audit studies to explore the characteristics associated with FOI responsiveness. Notably, 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 (as compared to as many as just over 25,000 in some of our analyses). Importantly, using 'natural' requests also opens the door to examining aspects of transparency for which it might be challenging to obtain human subjects approval.[7] Finally, we also see our approach as offering a relatively straightforward method to scaling the study of FOIA compliance.

Closest to our paper are Berliner et al. (2021) and Trautendorfer et al. (2023), which both analyze large numbers of requests to understand the nature of responsiveness. The first of these, in particular,looks at the relationship between governing party vote share and responsiveness to over 450,000 freedom of information requests in Mexico, concluding that

[7]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.

– as in our results – 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, the particular platform we deploy, given that it is often used for information gathering at scale on the identical topic by the same submitter, allows for a clear apples-to-apples comparison (our group-matched approach) across jurisdictions or circumstance more broadly.[8] Second, our test cases – which we see as a contribution that is on equal footing to our method – focuses on determinants of compliance that are largely distinct from those that these earlier papers study (the point of overlap is in how political considerations broadly defined may affect transparency, which is the focus of berliner2021political). Finally, given that measures of legal stringency in state-level FOIA statutes are well-developed, we can compare the relative importance of compliance with the law versus the laws themselves in determining government responsiveness. 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

We begin by describing the primary new dataset that we developed for the purposes of this paper, which we use to generate our "revealed transparency" measures. We utilize the FOIA requests filed via MuckRock, a non-profit which facilitates the efficient, low-cost submission of requests, potentially to many agencies or jurisdictions at one time. Its purpose is to create a "collaborative news site that brings together journalists, researchers, activists, and regular citizens to request, analyze, and share government documents, making politics more transparent and democracies more informed." In keeping with this mandate, all requests and subsequent responses made via its website are available for download via its API. Notably,

[8] 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.

given MuckRock's objective, we also observe that FOIA requests submitted via the site are reflective of actual demands made via sunshine laws.

MuckRock facilitates the submission process by generating appropriate emails or submitting requests via agency web portal as appropriate, and sending auto-reminders (in addition to any correspondence the submitter sends to an agency) until a request receives acknowledgement and/or response. MuckRock additionally 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.

2.1 Data for state-level analysis

At the end of October 2022, we scraped all requests made via MuckRock since its inception in 2010, a total of 81,180 requests.[9] In its first year, 219 requests were place; 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 81,180 total, 28,182 (34.7%) 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 52,998 requests, 36,681 (69.2%) are city-level requests and 16,317 (30.8%) are state-level requests.

MuckRock's "Essentials Team" codes the status of a given request, according to fixed

[9]This is a sufficiently modest total that we do not expect that MuckRock itself stands out as a notable source of FOIAs to officials charged 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 2019 (the last year before a COVID-induced drop) was 858,952 requests, as compared to 4,079 federal requests filed via MuckRock in 2019; the total number of FOIA requests filed via MuckRock that year were 10,118.

guidelines on the extent to which an agency has complied with the request, updating the status over time as appropriate.[10] In our main analyses, we consider the set of requests that MuckRock codes as having received a clear resolution – those that are marked by MuckRock as "completed" (21,638 of the state and local requests), and those that end in rejection (4,366 requests). Since our main independent variable will be state-level corruption, we further drop the 151 requests from the District of Columbia government, yielding a sample of 25,853 requests.

Not all requests are resolved in this binary manner. The full set of designations include: withdrawn by the submitter (3,414), awaiting acknowledgment from the agency (2,181), submitted and awaiting response (52), being processed by the agency (2,085), awaiting appeal to agency (191), fix required by submitter (4,107), in litigation (96), payment required by submitter (1,419), request was only partially completed (811), and no responsive documents provided by the agency (12,638). Most of these categories are ambiguous in their interpretation. While about half of the "awaiting acknowledgment" requests were relatively recent – from 2020-2022 – 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 volume 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

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

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).[11]

MuckRock itself, as reflected in its color-coding of outcomes, views Partially Completed requests as (green-coded) successes and No Responsive Documents as (red-coded) failure, and defines a request as closed if it is Completed, Partially Completed, Rejected, or No Responsive Documents; all other intermediate categories are (yellow-coded) ambiguous and potentially still in-progress. We thus perform a robustness check of our main results in which we use all closed files, classifying "no relevant documents" as rejections and partial completions as complete.[12]

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 in terms of 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.[13] 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 is a range of rationales, however. Wyoming, for example, re-

[11]See https://www.muckrock.com/news/archives/2018/sep/06/foia-faq-nrd-wtf/; last accessed October 20, 2023.

[12]If we include yellow-coded requests as 'intermediate' (e.g., as 0.5) our main results are unchanged.

[13]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.

sponded 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 above, 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 within the same quarter.[14] 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.[15] For our main sample of only completed and rejected requests, there are 15,112 singleton requests, leaving 10,892 observations for the matched groupings analysis. Since we will further be interested in groupings for which there is cross-state variation, when we limit the groupings to those with at least two requests in different states, our sample is reduced by an additional 4,865 observations, to 6,027.

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 – 500 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 October 2022), and 82 groupings with

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

[15]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 pair is smaller) within the matched sample. We verify that identical titles filed by the same person and around the same date are FOIA requests with 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 within the same time period.

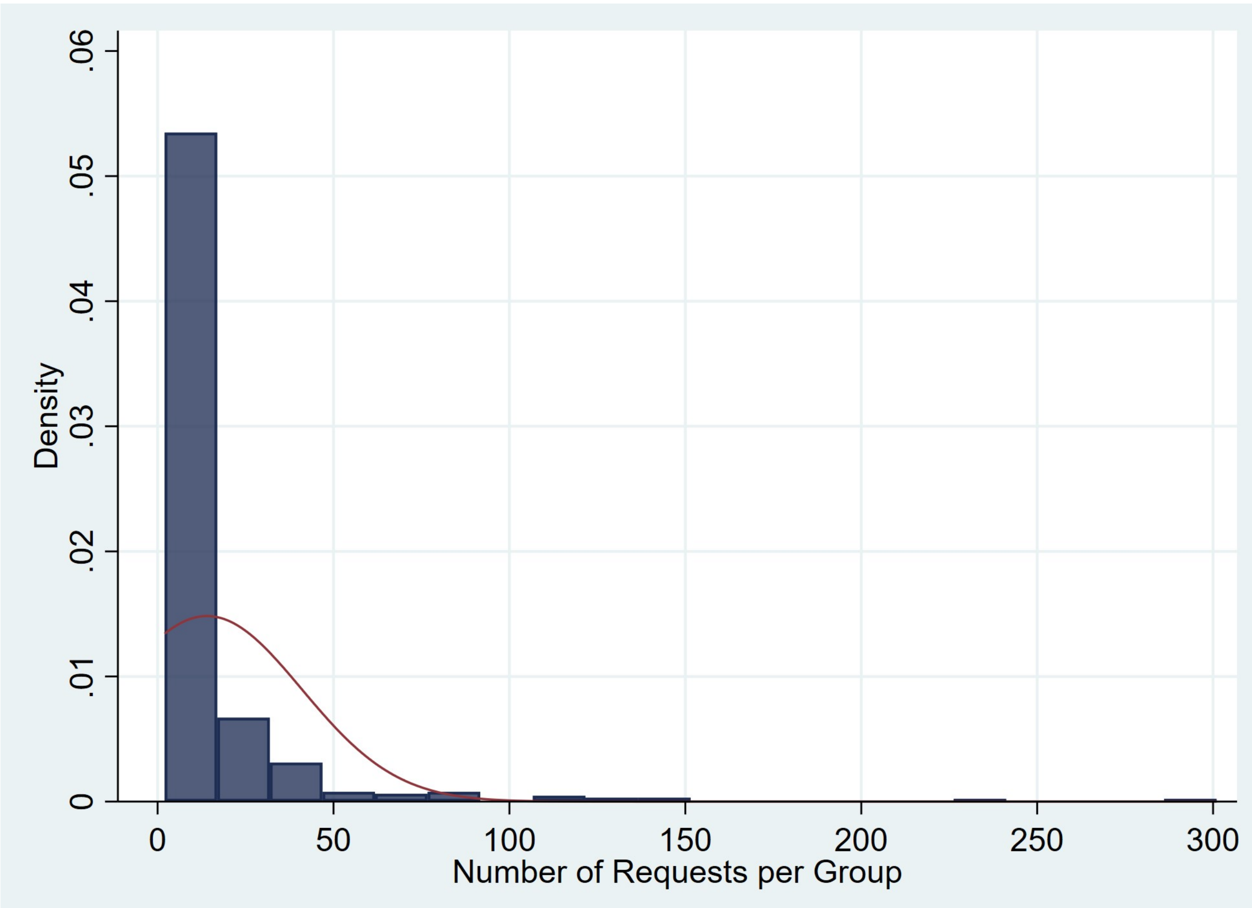

over 30 resolved requests. We make two observations on these small- and large-group requests. First, we note that we generate very similar results when we omit relatively small groupings (e.g., fewer than five requests per group). Second, and perhaps more substantively, larger groupings may introduce 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 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 again generate very similar findings to those in our main results.

For our state-level analyses, the main independent variable is state-level corruption. We utilize the 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 timespan 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. The Reporters Committee for Freedom of the Press compiles and publishes this data and we follow Cordis and Warren (2014) to obtain the FOIA law scores.

2.2 Data for City-level election analysis

In our second set of analyses, we focus on FOIA responses around municipal elections, and so the starting sample is city-level filings. Our focus is on the 100 largest cities in the U.S., as the election dates are readily available via Ballotpedia (a non-profit online encyclopedia of U.S. politics) and the relatively high frequency of FOIA requests in these locations allows for observations from the same city to serve as "treated" election and "control" non-election observations. In the five cases in which there were multiple mayoral elections within a year (due for example to resignations) we focus on the first election; we similarly focus on the initial election in races in which there are runoffs.

To construct the election-level municipal dataset for the 25,387 filings in these 100 cities, 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 mayoral election in the city where the FOIA was filed. Within each grouping, we define an indicator variable, *Election*, that denotes whether there is an election in the next 180 days. We will consider alternate windows ranging from 120 to 360 days 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 median ratio of requests in cities with upcoming elections to those in a benchmark set of control cities over the same time period is exactly one.) Rejection is straightforward to define, as earlier in 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 city that is 90 days in the future and a second in which there is an election 150 days in the future, as well as several requests in cities 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 straightforward deadline for the non-election (control) observations. If we were to define the control deadline based on 90 or 150 days, it would create a mechanical correlation between deadline and treatment status. Since the median number of elections in a grouping is two, this issue arises for a majority of cases, and especially for the larger groups of requests.

Our approach is to simply apply a uniform deadline across all observations within a grouping. In our main analysis, we define a request as $Unfilled$ if it is unresolved within the shortest election window in the matched grouping. In the example above, we would thus assign $Unfilled = 1$ for all requests that are not resolved within 90 days. As a robustness check, we also look at the longest pre-election window to define $Unfilled$ (in the above example, 150 days). We combine rejection and failure to measure a timely response with the variable $Failure$, which denotes requests for which $Unfilled = 1$ or $Rejected = 1$. We will also look at $Unfilled$ and $Rejected$ separately as outcomes. The sample for this set of analyses is comprised of 3,086 observations in 278 matched groupings.

Table 1: Summary Statistics

	Mean	Median	Std. Dev.	Number of Obs.
Panel A: FOIA Responses				
All Agencies				
Filed Lawsuit or Appealed	0.010	0	0.099	81,180
Payment Required	0.019	0	0.137	81,180
Rejected if completed	0.209	0	0.407	38,299
Filing Per Journalist	24.29	3	172.45	3,342
Federal Agencies				
Filed Lawsuit or Appealed	0.019	0	0.135	28,182
Payment Required	0.005	0	0.071	28,182
Rejected if completed	0.296	0	0.456	12,295
Filing Per Person	16.29	2	117.84	1,729
State Agencies				
Filed Lawsuit or Appealed	0.005	0	0.071	16,317
Payment Required	0.031	0	0.174	16,317
Rejected if completed	0.201	0	0.401	8,090
Filing Per Person	11.99	2	57.79	1,361
City Agencies				
Filed Lawsuit or Appealed	0.006	0	0.075	36,681
Payment Required	0.025	0	0.156	36,681
Rejected if completed	0.153	0	0.360	17,914
Filing Per Person	18.01	2	104.24	2,037
Panel B: Corruption Measures				

continues on the next page

Table 1: Summary Statistics (cont.)

	Mean	Median	Std. Dev.	Number of Obs.
State Level				
Corruption Score	0.275	0.240	0.132	50
Panel C: Other Variables				
State FOIA Law Score	6.035	6	2.286	50
State Population	4,802,412	3,239,950	5,131,990	50
State Income	12,888	12,791	1,825	50

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 this particular sample of requests. We begin in Panel A of Table 1 with some summary statistics for the FOIA data (Panels B and C provide summary statistics for our corruption measure and for the control variables, respectively). For the full sample, among requests that received unambiguous resolution as either rejected or fully fulfilled, 21% were rejected. 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 46% – simply because of the many requests that receive a "no relevant documents" response. There is considerable heterogeneity across levels of government. The rejection rate for federal agency FOIAs is nearly double that of municipal agencies.

In Appendix Table A1, we show the rejection rates (conditional on being completed or rejected) for the 15 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 40 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.

2.3 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 examine 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 assess 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 5,263 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.

And while this affects the extent to which we may generalize about our findings, it is less clearly an issue for the internal validity of our findings, which are based on matched-groupings of identical requests. Thus, within the corpus of documents we focus on, we are making comparisons requests that are the same but for 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 on request rejections which, as discussed above, often are presented by responding agencies as reflecting confidentiality concerns. This naturally raises the question of whether states that have lower rejection rates are "too open" relative to some optimal benchmark.[16] The question of optimal transparency is beyond the scope of our analysis; however, the fact that 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 revealed transparency

We begin by presenting 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

[16] 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 first set of cross-jurisdictional analyses.

(2014), before proceeding to a parallel analysis based on our set of matched groupings.

In our first analysis, we collapse the data to the state-year level, 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 end in either full completion or rejection in state s in year t. 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 (Arkansas, Deleware, Kentucky, Tennessee, and Virginia), which in our inspection of individual FOIA requests we have found to be a very common reason for rejection.[17] 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 0.35 (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 is negatively correlated with rejection ($p < 0.10$); however, we will observe shortly that this 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 0.12, so that a standard deviation increase in corruption rate is associated with a 23 percent increase in the average rate at which FOIA requests are rejected, relative to the baseline average rejection rate of 19%.

In Figure 2 we illustrate the relationship between *Corruption Rate* and *Rejection Rate*.

[17] We have confirmed that our results are not sensitive to the treatment of Alabama, which has ambiguous rules on residency.

Table 2: State Level Corruption and FOIA Response

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), whereas the outcome variable, *Rejection Rate*, is the average FOIA rejection rate by the given state. The mean of the dependent variable, *Rejection Rate*, is 0.19, whereas the standard deviation of *Corruption Rate* is 0.12. Standard errors are clustered at the state level and reported in parantheses. * 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	0.383**	0.373**	0.361**	0.276**	0.361**	0.319**
	(0.150)	(0.152)	(0.143)	(0.123)	(0.135)	(0.125)
Log(Income)			-0.131	-0.0730		-0.0392
			(0.162)	(0.136)		(0.141)
Log(Population)			-0.0158	-0.00646		-0.0167
			(0.0184)	(0.0156)		(0.0153)
Average FOIA Score				-0.0198**		
				(0.00928)		
***1**(Residents Only)*				0.177***		0.178**
				(0.0603)		(0.0672)
Fixed Effects						
Year		X	X	X		
N	450	450	450	450	50	50
R^2	0.044	0.085	0.098	0.181	0.117	0.290

Figure 2: State Level Corruption and FOIA Response

This figure plots the relationship between the state level corruption rate and average FOIA rejection rate. We take the average of all state-year level corruption rates to aggregate it to the state level. Similarly, we take the average of state-year level FOIA rejection rates to obtain a state level rejection rate measure. The 95% confidence interval (CI) is also plotted.

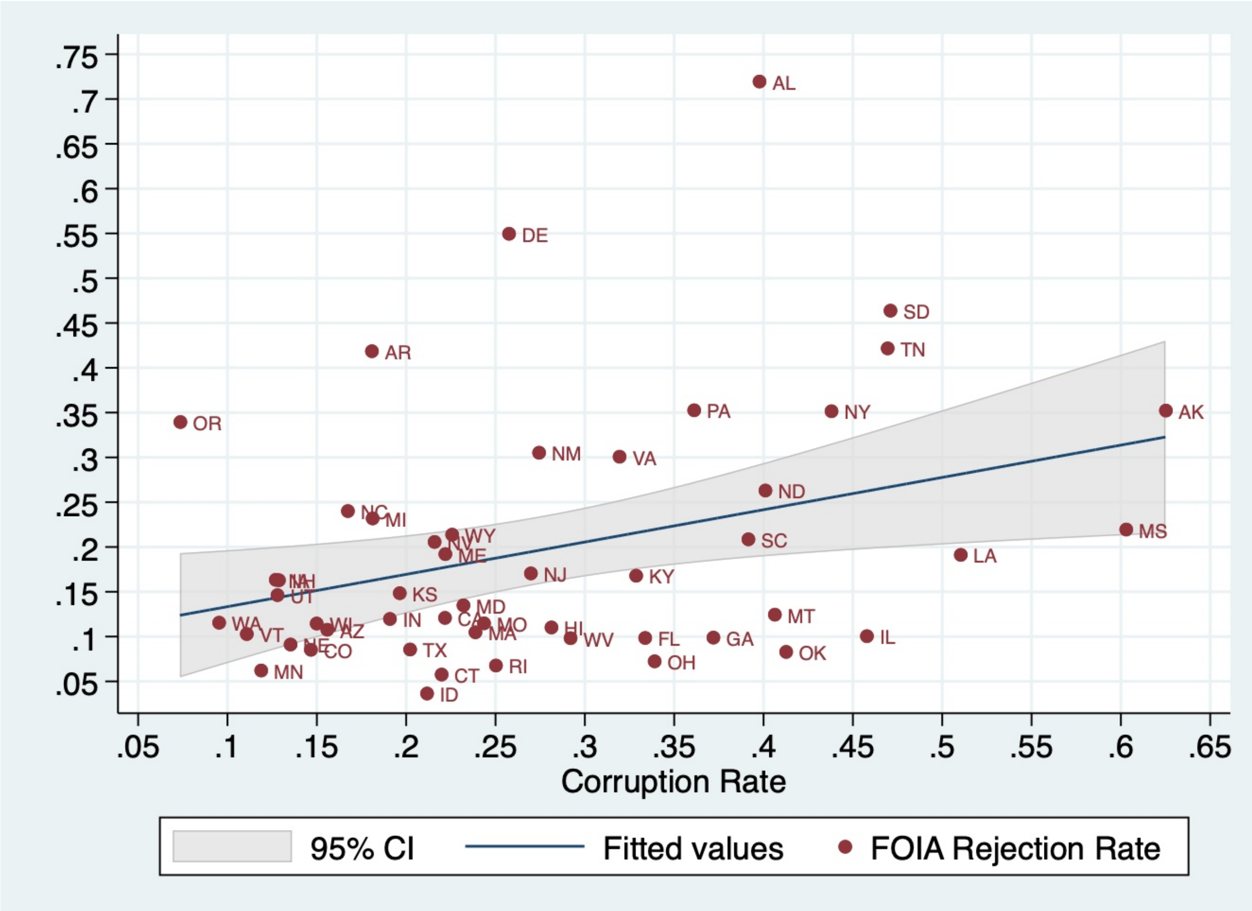

No individual outlier is driving the result, and the relationship is approximately linear. In Appendix Table A2, we use our alternative measure of rejection rate, in which we include partial completions as completed requests, and "no document" responses as effectively the same as rejection. We obtain similar results to those in our main analysis.

Finally, in Appendix Table A3 we repeat our analysis based on equation (1) but using the natural log of average response time. We observe no significant relationship between the time to fill a request and any state-level attribute.[18]

[18]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.

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 within the same calendar quarter. We amend the above specification to account for the different data structure:

$$Rejection\ Rate_{gst} = \alpha + \beta \times Corruption\ Rate_s + X_{st} + \gamma_g + \epsilon_{gst} \tag{2}$$

In all specifications we include 1,215 group fixed effects, and use two-way clustering – again by state and also by grouping of matched requests. 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.

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 reassuringly close to those based on the preceding cross-sectional analyses, though the point estimates are more noisily estimated. 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 favored within-group analysis. For this subsample, the point estimate increases marginally, to 0.43 (significant at the 5% level). In the final two columns we present results for our preferred specifications that include grouping fixed effects. The point estimate on *Corruption Rate* is 0.37, significant at the 5% level. The final column also includes a measure of FOIA *legal* stringency (*Average FOIA Score*). It is noteworthy that this 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

Table 3: Matched Sample Analysis of Corruption and FOIA Response

This table presents the association between the average corruption rate and the average FOIA responses by a given city during a given date when we match identical FOIA requests that were filed to a given city's department by the same person during the same period. The data is, therefore, at the date t and agency i of city $c(k)$ − agency i of city $d(l)$ pair level. The mean of the dependent variable, *Rejection Rate*, is 0.13, whereas the standard deviations of *Corruption Rate* is 0.11. In columns 4, 5, and 6, we require that the standard deviation of corruption within a pair should be greater than 0, whereas we do not put such a condition in columns 1, 2, and 3. Standard errors are double clustered at the state and pair level.

	Dependent Variable: *Rejection Rate*					
	(1)	(2)	(3)	(4)	(5)	(6)
Corruption Rate	0.288*	0.325*	0.348*	0.429**	0.366**	0.292*
	(0.165)	(0.181)	(0.191)	(0.197)	(0.179)	(0.157)
Log(Income)			-0.203	-0.206	-0.274	-0.154
			(0.132)	(0.169)	(0.167)	(0.122)
Log(Population)			-0.0164	-0.0301**	-0.0239*	-0.0147*
			(0.0103)	(0.0133)	(0.0125)	(0.00858)
Average FOIA Score						-0.00509
						(0.00787)
***1**(Residents Only)*						0.297***
						(0.0544)
Fixed Effects						
Year		X	X	X	X	X
Pair-ID					X	X
Condition						
Pair-level $\sigma(Corruption\ Rate) > 0$				X	X	X
N	10,892	10,892	10,892	6,027	6,027	6,027
R^2	0.008	0.026	0.035	0.044	0.282	0.313

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 to over 40 percent of our overall sample, and is 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 variation that exists within our groups – not simply across-state variation. Motivated by this observation, we now turn to exploring 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 explore responsiveness to FOIA requests in 180 days leading up to city elections. As we explained in the introduction, the relationship between electoral pressures and FOIA responsiveness is decidedly ambiguous. Most obviously, officials may be less apt to respond to FOIA requests if the resulting revelations risk political embarrassment. On the other hand, electoral pressures may lead to greater responsiveness, lest opacity become an election issue in itself. We see it as an empirical question, and one that we evaluate to illustrate the potential benefits of using granular request data that may be disaggregated in any way by time, geography, and even agency.

Our specification focuses on groupings for which at least one request is to an agency in a city where there is an election in the upcoming 180 days, and at least one request in a city where there is *not* an election in the next 180 days as a benchmark. Recall that we will also consider a shorter and also longer pre-election windows (from 120 to 360 days).

As discussed in Section 2, there are two ways of avoiding a pre-election FOIA response: delay ($Unfilled$), or outright rejection (*Rejection*), and 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. Our analyses are based on the following specification:

$$\begin{aligned} Failure_{rct} = \alpha + \beta_1 \times Election_{rct} + \beta_2 \times Corruption\ Rate_{s(c)} \\ + \beta_3 \times Election_{rct} \times Corruption\ Rate_{s(c)} + \gamma_g + \epsilon_{rct} \end{aligned} \tag{3}$$

for request r at time t in city c. As before, we include a set of matched group fixed effects, γ_g, for requests with the same submitter, content, and quarter of submission.[19] In some specifications, we will also include state fixed effects, λ_s, which absorbs the level effect of corruption. We cluster standard errors by matched-group and also by state.

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 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 -0.05, indicating an 5 percentage point *lower* probability of having a "failed" FOIA request if a municipal election is in the relatively near future, compared with the probability of failure of an identical request made in a city without an upcoming election (though the point estimate

[19]We use a matching algorithm to construct the matched groups by matching on filer, title, and date. We manually verify all matched groups by randomly selecting at least five (or all requests if the total number is smaller) FOIA requests within the matched sample. We verify that identical titles filed by the same person and around the same date are FOIA requests that also contain identical contents in the body of the request, apart from differences in agency name.

is significant only at the 10% level).

Overall, our results suggest that the "electoral accountability" effect dominates. But it is then natural to consider whether accountability *in general* interacts with electoral pressures. We thus consider whether the effect of elections on responsiveness differs as a function of (state-level) corruption, adding the term $Election \times Corruption\ Rate$ in column (2). The coefficient on this interaction is 0.63 (significant at the 5% level), and implies a "crossing point" of $Corruption\ Rate = 0.35$, where the *Election* effect is zero, at approximately the 75th 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 coefficients on both *Election* and $Election \times Corruption\ Rate$ are relatively unaffected by this inclusion.

Table 4: Matched Sample Analysis of Corruption and Failure to Response Around City Mayoral Elections

This table presents the test of whether cities with a mayoral election failure to response, defined as either a rejection or failure to response prior to an election, during the year prior to the city's mayoral election when we match the given FOIA request to other identical FOIA requests that were filed to a different city's department by the same journalist during the same period. The variable $\mathbb{1}$(Election) takes the value of one if the FOIA request was filed with a department in a city that had a mayoral election one year prior to the election date. The mean of the dependent variable, $\mathbb{1}$(Failure), is 0.33, whereas the mean of $\mathbb{1}$(Rejected) and $\mathbb{1}$(No Decision) are 0.11 and 0.25, respectively. Standard errors are double clustered at the pair and state level.

	Dependent Variable								
	$\mathbb{1}$(Failure)			$\mathbb{1}$(Rejected)			$\mathbb{1}$(No Decision)		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
$\mathbb{1}$(Election)	-0.0513*	-0.220***	-0.189***	-0.0258*	-0.120***	-0.0826**	-0.0314	-0.141**	-0.131**
	(0.0278)	(0.0671)	(0.0561)	(0.0144)	(0.0388)	(0.0337)	(0.0252)	(0.0551)	(0.0500)
Corruption Rate	0.0592	-0.0451		0.308*	0.250		-0.199	-0.267*	
	(0.242)	(0.218)		(0.180)	(0.172)		(0.137)	(0.132)	
$\mathbb{1}$(Election) × *Corruption Rate*		0.633**	0.496**		0.353**	0.217*		0.410*	0.352*
		(0.264)	(0.226)		(0.143)	(0.119)		(0.221)	(0.205)
Fixed Effects									
Pair-ID	X	X	X	X	X	X	X	X	X
State			X			X			X
N	3,086	3,086	3,086	3,086	3,086	3,086	3,086	3,086	3,086
R^2	0.216	0.218	0.270	0.257	0.258	0.326	0.213	0.214	0.253

The remaining columns of Table 4 look separately at the two components of $Failure$, i.e., $Rejection$ and $Unfilled$. We observe comparable patterns for both: a marginal improvement from elections on average, but in a manner that is dependent on the broader institutional environment.

In Appendix Tables A5-A9, we show the preceding patterns for both shorter and longer windows of pre-election requests, from 120 to 360 days. While the point estimates naturally vary in magnitude and significance, the patterns we describe above are generally observed over all windows.

4 Conclusion

In this paper we introduce a new tool for measuring government transparency within the U.S. 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.

Our primary interest in this paper is demonstrating the credibility and potential usefulness of our "revealed transparency" measure. We do so by showing an intuitive correlation between state-level FOIA rejection rates and state-level corruption prosecutions, which we take as a rough proxy for government probity, and to illustrate the flexibility of our approach, we examine how government transparency differs as a function of upcoming elections.

We see many potential directions to make use of our ever-expanding dataset, which may take advantage of the variation in our data across time, location, and type of government agency. 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. We may also explore how particular events impact openness – the 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. T. Enevoldsen, R. Pande, and M. Walton (2023). Public information is an incentive for politicians: Experimental evidence from delhi elections. *American Economic Journal: Applied Economics*.

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.

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.

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 (2020). The effect of sustained transparency on electoral accountability. *American Journal of Political Science*.

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

Harden, J. J. and J. H. Kirkland (2021). Does transparency inhibit political compromise? *American Journal of Political Science 65*(2), 493–509.

Honig, D., R. Lall, and B. C. Parks (2022). When does transparency improve institutional performance? evidence from 20,000 projects in 183 countries. *American Journal of Political Science*.

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.

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.

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*.