

The Political Logic of Government Disclosure: Evidence from Information Requests in Mexico

Daniel Berliner*, Benjamin E. Bagozzi†, Brian Palmer-Rubin‡ and Aaron Erlich§

July 23, 2019
Forthcoming, *The Journal of Politics*

*Associate Professor of Political Science and Public Policy, Department of Government, London School of Economics and Political Science, d.berliner@lse.ac.uk

†Associate Professor of Political Science and International Relations, University of Delaware, bagozzib@udel.edu

‡Assistant Professor, Department of Political Science, Marquette University, brian.palmer-rubin@marquette.edu

§Assistant Professor of Political Science, Department of Political Science, McGill University, aaron.erlich@mcgill.ca

Abstract

When citizens ask questions, how does their government answer? Requests for government information confront officials with incentives both for and against disclosure. We argue that officials seek to manage political risks in ways that favor requests from government-aligned regions. We study responsiveness in the context of Mexico's access-to-information law, using publicly available data from several hundred thousand information requests filed with Mexican federal government agencies between 2003 and 2015. Our empirical strategy makes comparisons only among requests sent to similar agencies on similar topics at similar times, while accounting for the complexity, sophistication, and sensitivity of individual requests. We find that requests filed from locales with higher governing-party vote shares receive more favorable responses, across multiple indicators of the nature and timing of responses. Further, we find bias only for requests on publicly relevant topics, supporting a mechanism of mitigating political risks over one of rewarding supporters with greater access to benefits.

Keywords: Transparency, accountability, access to information, Mexico.

Supplementary material for this article is available in the appendix in the online edition. Replication files are available in the JOP Data Archive on Dataverse (<http://thedata.harvard.edu/dvn/dv/jop>). This research was supported by Social Science and Humanities Research Council (SSHRC) Grant #430-2018-1069, Fonds de recherche du Québec - Société et culture (FQRSC) Grant #253243, McGill University, the Arizona State University Program for Transborder Communities, and the University of Delaware General University Research Fund.

Introduction

To whom are agents of accountability themselves accountable? Government officials charged with responding to citizen queries and disclosing official information play a fundamental role in *downward* “vertical accountability,” in which the state answers to citizens for its actions (O’Donnell 1998; Schedler 1999). And yet, as agents of political principals—both their hierarchical superiors and elected politicians—they may face conflicting incentives to also be accountable *upwards*, in ways that inhibit the provision of information to citizens. How do government officials resolve these roles, and with what effects on responsiveness to citizens?

While many scholars have studied how electorally motivated discretion drives government expenditures or public goods provision (Besley and Burgess 2002; Golden and Min 2013; Kramon and Posner 2013; Stokes et al. 2013; Diaz-Cayeros, Estevez and Magaloni 2016; Carlitz 2017), government responsiveness extends beyond spending decisions: governments interface with citizens in millions of everyday interactions, in which agents of the state may vary in their responsiveness to citizen needs. We examine government responsiveness in one case of frequent citizen-government interaction: responses to citizen requests for government information in Mexico. To do so, we use a novel dataset incorporating every information request filed with Mexican federal government agencies between 2003 and 2015.

Much like government *expenditures*, government *information* can be divided into particularistic and public categories. For instance, information useful in applying for government benefits or employment serves the requester’s private interest, while information that reveals government corruption or inefficiency serves the public’s interest in good governance. While this categorization mirrors the public-private good distinction made in the distributive politics literature, the provision of government information produces a different set of incentives for politicians than does distributive spending, in particular because publicly relevant information can present a political risk to the incumbent administration.

We thus focus attention on the distinct incentives at work in the specific domain of government information. Officials charged with responding to citizens serve two distinct roles that are central to democratic governance and accountability. First, they are gatekeepers for information about government activities—from budgets to security to evidence of corruption—relevant for

citizens' retrospective evaluations of government performance, a central component of democratic accountability (Przeworski, Stokes and Manin 1999). Second, they provide information useful for guiding citizens' everyday dealings with government agencies, such as applying for benefits, acquiring titles or permits, or seeking government employment. As the timely access to such information confers material advantages, it can become a currency of clientelism if granted discretionarily according to partisanship (Stokes et al. 2013, 7). Following from these two roles, discretionary non-disclosure of public information can either challenge *political accountability* or can hinder the *quality of governance* by violating the principle of impartiality (Rothstein and Teorell 2008) in the provision of public goods and services. It is the goal of this paper to assess these two forms of discretionary disclosure.

For the agents of accountability at the heart of this process, motivations for responding positively include commitment to the goals of transparency as a tool for downward accountability, as well as avoiding formal or informal sanctions through state institutions of horizontal accountability. They may also be motivated by partisan incentives to provide information that helps political allies access state benefits or navigate regulatory processes. However, they also have many reasons to be less responsive, particularly concerning the risk of unwanted disclosures resulting in scandal or negative media attention that might harm the reputation of their agency's leadership or the political goals of the ruling party — both forms of upwards accountability.

How do officials manage these competing incentives regarding disclosure of government information? We argue that officials filter and discriminate based on the perceived political risk of requests, informed by the partisan alignment of communities where requests originate. However, counter to the distributive politics literature, which expects discretionary targeting of private goods, we argue that officials prioritize areas of core support in the provision of information that is most *publicly* relevant. As the latter type of information presents a greater risk of generating negative attention, politicians and their agents prefer to withhold this information from political opponents where possible. We thus suggest that government responsiveness in the domain of information, rather than spending or material goods, is better understood through such a logic of managing risks rather than a logic of targeting benefits.

We test these expectations by comparing the quality of responses across requests originating

from municipalities with varying levels of support for the governing party—but filed through the same centralized and ostensibly anonymous system. We address heterogeneity in requests originating in areas of high and low governing-party support by considering variation only among requests sent to the same government agency, on the same topic, and in the same year, and by incorporating text-based information on the complexity, sensitivity, legalism, and sophistication of individual requests. To control for the topics underlying requests, we use unsupervised topic models to classify requests into twenty topics, which we categorize as either primarily publicly relevant (e.g. budgets and spending, security issues, and environmental disputes) or primarily related to private goals (e.g. accessing services and benefits, or seeking jobs or contracts).

We find evidence that requests filed from areas of stronger government support receive more favorable responses: faster, less likely to claim the information does not exist, and less likely to raise financial or physical barriers. These results are stronger among publicly relevant topics than among topics with a more private focus, consistent with a mechanism of officials managing political risk, and inconsistent with a mechanism of rewarding allies with preferential access to material benefits. Further analyses reveal evidence of within-municipality variation in responsiveness over time, exploiting changes in political alignment after the Institutional Revolutionary Party (PRI) regained presidential power in 2012.

Using qualitative evidence, we surmise that bias in responsiveness to information requests is primarily due to officials inferring the political risk posed by a given request based on the traits of the municipality from which it originates — information that is often readily available to the bureaucrat. However, in some cases officials may be able to more precisely infer the identity of the requester and thus discriminate against political opponents or notable watchdog groups. Observationally, this more fine-grained discretion would yield the same results that we observe: worse responses to information requests that come from areas where political opponents of the ruling party predominate.

Interviews with bureaucrats at several liaison units charged with handling information requests in government agencies illustrate that the types of discretion we elucidate are procedurally feasible, and clarify the conflicting incentives presented by requests for government information. While Mexico’s access-to-information (ATI) system was designed to “level the playing field” by

anonymizing requests, liaison unit personnel report several informal avenues by which requester traits can potentially become known and influence responsiveness. Further, officials described both incentives favoring disclosure — such as avoiding sanctions for noncompliance or commitment to the goals of the transparency law — as well as considerations of the political risk that disclosure may pose to their supervisors and political principals, particularly for sensitive topics.

While we expect similar political dynamics to be at work in countries other than Mexico, this context offers several advantages for our analysis. On a practical level, the online request system offers richly detailed data over a long period of time. More crucially, as a federal democracy with stable multi-party competition, Mexico presents significant regional variation in political dynamics. Moreover, while Mexico has been at the vanguard among transitional democracies in the design and implementation of ATI institutions, persistent corruption and bureaucratic discretion introduce both motive and opportunity for biased responsiveness to information requests. However, we expect that our findings would extend to other contexts where bureaucrats have discretion in disclosing government information, yet face political risks in doing so.

Government Responsiveness and Information Disclosure

We combine insights from two literatures on government responsiveness. First, studies of *policy* responsiveness address the extent to which government policies respond to voter preferences (Golden and Min 2013). Such studies focus on partisan incentives that shape government spending and public goods provision, typically across subnational units (Besley and Burgess 2002; Magaloni 2006; Cleary 2007). A central debate in this literature regards how politicians allocate benefits between areas of ruling- and opposition-party strength (Diaz-Cayeros, Estevez and Magaloni 2016; Gans-Morse, Mazzuca and Nichter 2014; Stokes et al. 2013).

A second group analyzes *individual* responsiveness, the extent to which individual citizen needs are met in everyday interactions with the state. Recent studies have identified bias regarding citizens' inferred race, ethnicity, gender, and social status (Butler and Broockman 2011; McClendon 2015; White, Nathan and Faller 2015; Distelhorst and Hou 2014). Studies of responsiveness to information requests, the present focus, have identified differences based on the threat of collective action in China (Chen, Pan and Xu 2015) and the design of government agen-

cies in the United States (Wood and Lewis 2017). And studies of ATI in Mexico specifically have identified limited favoritism of citizens with implied political connections (Lagunes and Pocasangre 2019)¹ and biases against citizens seeking information about the drug war (Almanzar, Aspinwall and Crow 2017). Yet we lack understanding of how broader partisan objectives shape bureaucratic responsiveness in interactions with individual citizens.

ATI institutions enable citizens to request information from public officials, and obligate officials to respond, subject to certain legal exemptions. In practice, officials operate in conditions of limited capacity (in terms of staff, finances, and records-management systems), substantial discretion, and inconsistent oversight, making ATI institutions a promising venue to study responsiveness. Further, they generate data about responsiveness across two quite different areas of citizen activity: private goals for the benefit of the requester, and public goals related to governance and accountability (Michener and Worthy 2015). Information requests in the former category include, for example, questions about how to access government benefits or programs, conduct bureaucratic procedures, or seek government contracts or jobs. Public goals, on the other hand, include investigating corruption, patronage, or other wrongdoing, evaluating decision-making or policy effectiveness, and mobilizing for political action. Similar to studies that find different targeting logics for different types of distributive goods (Albertus 2012; Diaz-Cayeros, Estevez and Magaloni 2016; Kramon and Posner 2013; Palmer-Rubin 2016), we postulate that public and private varieties of information generate different incentive structures for officials.

As shown in Table 1, certain incentives favor responsiveness across both public and private information. Disclosure of requested information provides citizens a positive impression of government performance. Officials may feel committed to complying with legal mandates or the ethos of transparency, or may wish to avoid sanction for nondisclosure, either within their agencies or from oversight bodies. Officials also have reasons to be unresponsive to all types of information requests. Aside from standard criteria for withholding information (confidentiality, national security, commercial secrets, etc.), bureaucrats may be discretionarily unresponsive out of an interest in minimizing workloads, particularly in understaffed agencies.

¹A similar study in Brazil (Michener and Rodrigues 2015) found no such biases.

Table 1: Incentives For and Against Information Disclosure

	Incentives for Disclosure	Incentives against Disclosure
Information for Public Goals	Demonstrating performance; Personal commitment; Avoiding sanctions	Limiting workload; Complying with legal exemptions; <i>Managing political risks</i>
Information for Private Goals	Demonstrating performance; Personal commitment; Avoiding sanctions; <i>Targeting benefits</i>	Limiting workload; Complying with legal exemptions

However, other potential incentives, denoted in italics, are specific to the type of information. A core consideration in disclosing public information is *political risk*: the potential of the requested information to damage the reputation of the agency, its leadership, or the governing party. Political elites from the ruling party clearly wish to avoid scandals, investigations, or information revelations in the media that might harm themselves or their allies. Agency leaders also seek to protect their own future political careers. These considerations weigh against disclosure for topics of information that pose greater political risks. In contrast, for information related to private goals, disclosure might follow a logic of *targeting benefits*, as does distributive spending, strategically disclosing information that helps the requester get a job, win a contract, or receive government benefits, in order to reward supporters and punish opposition.

Officials responding to information requests consider both their own interests (including individual career goals and organizational reputation) and the interests of their political principals. The prevalence of political appointees and patronage in bureaucratic agencies produces dependence on the governing political party’s interests. We thus expect that officials responding to information requests operate, at least in part, as agents of the party controlling the presidency, and act in accordance with their political interest.

Both public and private goals can be furthered by favoring co-partisans. Officials can discriminate directly in cases of known requesters, or use the political alignment of the location from which a request originates to infer the likelihood that the requester is an “ally” or “adversary.” Either case means that officials respond in ways that disproportionately favor government supporters and government-supporting areas. Thus, our *Partisan Discretion Hypothesis* predicts that, across all types of requests, those from government-supporting areas receive more favorable

responses than those from opposition-supporting areas.

We further pose two subsidiary and competing hypotheses, regarding whether such partisan discretion follows a logic of targeting benefits or one of mitigating political risk. A naive deployment of findings from the distributive politics literature would predict a strategy based on targeting benefits. As such, if officials manipulate disclosure decisions as part of a clientelistic strategy to target voters with particularistic benefits (Diaz-Cayeros, Estevez and Magaloni 2016; Gans-Morse, Mazzuca and Nichter 2014; Stokes et al. 2013), we would expect to see stronger effects for topics of information requests that contribute to *private*, rather than public goals. The *Targeting Benefits Hypothesis* predicts that partisan effects are stronger for requests on topics of information about private needs or benefits than about information with public potential.

In light of the above discussion about the acute political risk posed by disclosure of information with public relevance, however, we have reason to expect the opposite. Given the high costs associated with scandal weighed against the marginal payoff for helping allies access benefits—as well as incentives to appear impartial and maintain high response rates—we hypothesize that officials may exercise more discretion in disclosing information that poses political risks than that which generates private benefits. Therefore, the *Mitigating Risk Hypothesis* predicts that partisan effects are stronger for requests on topics of information with public potential than about private needs or benefits.

Access to Information in Mexico

Mexico's *Ley Federal de Transparencia y Acceso a la Información Pública Gubernamental* (LFTAIPG) passed in June 2002, taking effect one year later. An ATI law was a campaign promise of President Vicente Fox's Partido Acción Nacional (PAN) in the watershed 2000 presidential election. Although Fox initially proposed weak draft legislation, sustained civil society, media, and opposition pressure contributed to the ultimate passage of a law with cross-party support, widely hailed as one of the world's strongest (Pinto 2009; Michener 2011; Berliner and Erlich 2015). The law's implementation and wide usage is seen as a model among developing countries (Bookman and Guerrero Amparán 2009; Michener 2011).

Notably, the ATI law created an independent information commission, the *Instituto Federal*

de Acceso a la Información (IFAI),² tasked with promoting the new law, monitoring compliance, and resolving appeals. IFAI also created a centralized online information portal through which government agencies manage requests. Citizens file requests and receive responses almost exclusively (97% of all requests) through this portal, known as INFOMEX until 2016. When citizens do file handwritten requests (only 3% of all cases), INFOMEX agency officials manage these and enter the relevant information into the portal's database.

When a request is filed, it arrives from the INFOMEX system first to the *unidad de enlace*, or transparency liaison unit, within the government agency to which the request was directed. The unit's staff reviews the request and directs it to the relevant administrative units within the agency, unless simple or frequent enough that it can be handled directly. Administrative units then return preliminary responses to the liaison unit, with further communication if necessary before responses are returned via INFOMEX. Most agencies manage these internal processes using combinations of emails and spreadsheets, with only ten percent reporting automated management systems even as late as 2016 (Ríos Cázares, Castañeda and García 2017). Each agency also has an internal information committee that reviews and approves some categories of denial or deadline extensions. The head of the liaison unit sits on the information committee, along with a representative of the agency head and an internal control officer.³ This creates a formal venue for discussions around how to respond; although informal discussions can also take place.

While several observers claim that INFOMEX maintains requester anonymity (Sobel et al. 2006; Guerrero-Amparan and Toledo 2009; Mizrahi and Mendiburu 2012), this is not generally the case. Indeed, Lagunes and Pocasangre (2019) observe that "While Mexico's FOIA system is supposed to be identity-blind, we rarely see instances in which the identity of the individual is blacked out from requests." And Fox and Haight (2007, p. 40) note that "At first, the IFAI tried to prevent agencies from accessing personal data in SISI [INFOMEX's earlier name] accounts, however agencies insisted on knowing who was requesting information. The IFAI accepted this position, in the absence of a legal basis for defending the anonymity of requestors."

²In 2015, its name changed to INAI.

³Several procedural details were changed by a 2016 reform, but this paper refers to the situation during the period under study.

Individuals must create an account to use INFOMEX, and provide not only their name and contact information but also their physical address. This requester information can be seen by liaison unit staff who first receive requests and distribute them within their agency. Some agencies use specialized software to keep requester information out of internal communications, but most do not. A 2016 survey of liaison units found that 99 percent reported communicating about requests over email and 86 percent over telephone (Ríos Cázares, Castañeda and García 2017, p. 61). Even where specialized software exists, however, anonymity is not guaranteed. In one large, cabinet-level ministry, liaison unit staff claimed that their internal system *did* allow other officials to see full requester information.⁴

This complicated situation regarding requester information leaves several avenues for discrimination, either by specific requester identity (which would only sometimes be informative, and which we cannot observe in the publicly-available data) or by location (which is both systematically informative and which we can observe). First, officials in the liaison unit—who can formally see requester information—might alter the speed or nature of initial decisions based on the requester. Second, the agency’s information committee might make different decisions regarding extensions or claims of classification, confidentiality, or nonexistence of the requested information. Third, other officials may see requester information despite mandates to the contrary, due to the agency’s internal procedures or informal pressure. Several interviewees suggested that other officials indeed attempt to find out information about the requester (and potentially incorporate such information into their decisions), and are sometimes successful.

There is a clear incentive for selective disclosure, as information requests often lead to public scandals or negative media coverage. For instance, newspaper stories frequently reference the use of information requests in reporting and investigation. These include information on potential misuse of funds,⁵ public employees lacking qualifications,⁶ irregularities in procurement or

⁴Interviews were conducted with officials in the liaison units of seven different government agencies in March 2017, under Arizona State University IRB number 00005773.

⁵*Reforma*, Feb. 28, 2005, “Acusan anomalías en gasto de GDF.”; *Reforma*, March 20, 2015, “Financian en Pemex tarjetazos a directivos Conserva plana mayor beneficios pese a crisis.”

⁶*Reforma*, April 11, 2014, “Detecta SAT a 12 funcionarios ‘cachirules.’ ”

handling of supplies,⁷ and even to investigate potential arbitrary detention by police forces and the discovery of mass graves.⁸

In managing the political risks of disclosure, officials face conflicting incentives. Interviewed liaison unit officials emphasized both their commitment to the law and to serving citizens, and their fear of “getting in trouble” with IFAI, or of attracting negative attention to the agency for poor compliance records—all incentives in favor of disclosure. But officials also discussed political considerations. One official said that, for some sensitive requests (often relating to patronage spending), higher-level officials simply order the denial of the request, leaving the liaison unit to come up with a defensible justification. Another said that agency leadership sometimes argue directly with the information committee in favor of denying specific requests. The head of yet another liaison unit directly invoked the reputation of the head of the agency, saying “we have to take care of the secretary . . . we have to protect him.”

A 2007 survey (CIDE 2007) of 1,221 high level staff at 150 federal entities further demonstrates these mixed incentives. The survey found 73% agreeing that “the federal government should give access to all the information that it generates” (p. 37) and 76% aware of sanctions for noncompliance, 29% of whom cited suspension as a likely penalty (p. 26). Yet at the same time, 53% reported that they were justified in withholding requested information when they suspected the information could be used to attack the agency (p. 43), and respondents were broadly aware of a variety of strategies to withhold information, such as declaring it nonexistent or the responsibility of another agency, or requiring the requester to collect a physical copy (p. 44).

While government agencies and their officials do not directly face electoral incentives, they cannot remain fully independent of them either. Top agency leadership are appointees accountable directly to the President, and often seek to protect their own reputations in pursuit of future political advancement. Liaison unit heads are frequently replaced, although in some cases they remain in place through presidential transitions (even in 2012). Lower-level staff within each

⁷*Reforma*, Jan. 22, 2008, “Indagan pérdidas por medicina caduca.”; *El Universal*, January 29, 2013, “Piden a PF informar precio de flota vehicular.”

⁸*Reforma*, Feb. 6, 2015, “Hallan 100 cuerpos y nadie investiga!”; *Reforma*, Nov. 4, 2012, “Aumenta arraigo, no las condenas.”

unit, particularly those with substantial experience or legal expertise, are more likely to remain in place through a transition. Nearly all relevant staff to this process are *empleados de confianza*—technically patronage appointees who may be fired at will, despite usually pursuing longer-term careers in the public sector, and thus face incentives to protect their principals’ reputations. One interviewed official even noted liaison unit directors had been fired in the past for disclosing information that displeased agency leadership.

A 2016 survey of liaison units (Ríos Cázares, Castañeda and García 2017, p. 58-61) found that 88% of unit directors (and 76% of responsible staff) were *empleados de confianza* (with the bulk of the remainder being even less secure temporary employees), averaging 91 (93) months in the agency, and 33 (57) months in the response unit, respectively. Nearly all had at least undergraduate degrees, while roughly 40% of directors had either masters or doctorate degrees. The average response unit had 4-5 staff, including the director. A survey of officials at the ATI system’s founding (Gill and Hughes 2005, p. 132) found that “nearly one-half of respondents thought government officials would only partially comply with the new access law and said potential sanction by superiors, or the fear of such sanction, would prompt non-compliance.” Furthermore, given a policy requiring that unit directors report directly to the agency head, these ATI authorities are usually high-level agency officials, invested in the well-being of the agency and its leadership (Ríos Cázares, Castañeda and García 2017, p. 52).

Sample, Preprocessing, and Text Analysis

We focus on Mexico’s publicly available administrative data containing all ATI requests filed with federal government agencies during the period from June 2003 (when the law came into effect) through August 2015, as made available by the Mexican government via INFOMEX.⁹ These records contain detailed information on each request, including the date and time filed, the agency to which the request is directed, the full request text, the date of response, the nature of the response, and links to attached files associated with the request or response. Another field differentiates electronic requests (97% of the total) from manual requests filed on paper

⁹This excludes state-level authorities and constitutionally autonomous bodies, as these fall under different legal regimes.

and entered electronically by officials. While the publicly available data protect the identities of requesters by omitting names, user IDs, or answers to voluntary demographic questions, these data do include geographic information on a requester’s state and municipality, selected from ‘drop-down’ menus of pre-defined options, thereby preventing misspellings.

While INFOMEX cannot prevent individuals from entering false information, misrepresentations of requesters’ locations appears to be relatively rare, as there is a strong relationship between municipality size and request volume. Across municipalities, the logged total number of requests filed and logged population are correlated at 0.76. If a large share of requesters were selecting false locations, we would not observe such a strong relationship. Further, any measurement error in this regard simply sets a higher bar for us to identify meaningful results.

The text of each request corresponds to the requester’s own open-ended description of the information they seek. While the majority of requesters described the nature of their requests within one of two primary text fields, roughly 13% of cases included either some or all of the request as a file attachment. We implemented an automated routine to download these based on links contained in the INFOMEX data, convert to machine-readable text, and incorporate into our corpus. We also removed all official requests for confidential requester personal information.¹⁰ Finally, as a small number of requests were extremely long, often where an attached file contained an example report or a large spreadsheet, we truncated any request over 1,000 words, affecting only 0.02% of requests.

We applied a set of preprocessing techniques to convert our final corpus of requests into a collection of individual “documents” compatible with the unsupervised topic models used below, as well as to fix common typos, standardize the use of accents, and remove any non-graphical characters. We then processed each request text to remove all punctuation, numbers, names of months, and stopwords. The removal of these character-sets is standard preprocessing for our topic model techniques (e.g. Roberts et al. 2014) and helps to eliminate extremely common, but highly uninformative, character strings. We also removed common phrases used by requesters to indicate that the information request is not provided in the text field, but rather is attached.

¹⁰While these use INFOMEX, they are governed by different legal requirements, and additionally do not make public their request texts.

Next, we omitted sparse terms that do not occur in at least 0.1% of the documents in our corpus, which is a common means of eliminating extremely rare words that offer little insight into the shared topics across one's text corpus (Quinn et al. 2010; Grimmer 2012). Finally, all remaining words were converted to lower case, stemmed, and re-structured into a document-term-corpus before analysis. Altogether, these preprocessing steps created a corpus with 1,023,267 processed request documents and 6,696 total unique words.

To uncover the different topics of interest to citizens making information requests, we use Latent Dirichlet Allocation (LDA; Blei, Ng and Jordan 2003). LDA is an unsupervised topic model that identifies latent dimensions of texts. LDA estimates such topics via a hierarchical Bayesian model that treats documents as a finite mixture of underlying topics, and topics as a mixture over a latent set of topic probabilities (Blei, Ng and Jordan 2003, 993). In this manner, each identified topic can be characterized by a set of words that LDA determines as having the highest probability of association with that topic. These topic-word associations can then be used to (i) interpret each topic for its substantive meaning and/or (ii) to classify one's documents in terms of the topic(s) most likely to have generated each document, based upon its observed distribution of words.

With LDA, the analyst must explicitly choose the number of topics, k , to be estimated. In Appendix B, we evaluate different approaches to this choice, and identify a $k = 20$ as most appropriate for this application based on both qualitative evaluation and quantitative model fit statistics. We then estimate an LDA model on our full corpus of 1,023,267 preprocessed request texts. For this LDA model, in addition to setting $k = 20$, we assign the model's α parameter to favor the assignment of a single topic to each document,¹¹ and use Gibbs sampling for estimation. As such, this LDA model identifies the 20 topics that best characterize the request texts across the entire 2003-2015 period. We then classify each request according to its single most "dominant" associated topic (although we also utilize each request's full set of twenty topic probabilities in our analyses).

In Appendix C, we discuss an extensive set of evaluations aimed at interpreting these topics, including examination of the words most representative of each topic, and a close reading of sev-

¹¹We further justify this choice in Appendix B.

eral hundred individual requests sampled from those strongly associated with each topic. These interpretations are developed further in Berliner, Bagozzi and Palmer-Rubin (2018). Based on these evaluations, we split these twenty topics into two sets: thirteen topics with higher public potential, and seven topics oriented towards more private goals. These are listed in Table 2. We describe our LDA estimation strategy, and dominant topic approach, in further detail in Appendices B and C.

By topics with higher public potential we mean those that pose greater political risk to government entities. These include issues such as budgets and spending, crime and violence, environmental impacts, oil and utilities, distributive social programs, public procurement, and a topic capturing requests asking for information about specific public servants. These have clear politically relevant uses for investigation, oversight, media scrutiny, or political organizations. Requests on budgets and spending or on distributive programs might be aimed at identifying misuse of resources or favoritism. Requests about local environmental impact, permits for construction projects, or on decision procedures and compliance evaluations for government contracts might be aimed at uncovering corruption in allocating contracts and permits, or non-compliance with legal rules. Details on the scale of violence and the conduct of police and military in the “Drug War,” or on the qualifications of public officials, both also have clear potential to generate negative media attention or potential scandals. While even in these topics we still find some requests that are not likely filed with public goals, it is nonetheless these categories of requests that are most likely to activate a consideration of potential negative political consequences by officials.

Topics with lower public potential primarily reflect individual citizen needs, commercial interests, and demand for aggregate statistics. These pose less risk of generating negative media or investigative attention. For example, one topic largely comprises requests for information to fulfill individual citizen needs, such as how to access specific government services or how to complete specific procedures. Another primarily comprises requests for aggregate statistics pertaining to specific medical conditions. Relatively legalistic and sophisticated requests for commercial information, likely filed by companies themselves for private purposes, dominate a third topic. And while we identify other topics focusing on public employees as having higher public potential (due to a focus on specific individuals or qualifications that might be used to

identify unqualified patronage hiring), another topic is primarily made up of requests seeking information on aggregate salary scales, benefits, and pensions – likely filed by employees themselves.

In Appendix C, we describe all twenty topics in more detail, provide example requests, and explain our categorization of them as mainly “public” or mainly “private.” This division of topics is not absolute: an individual request for budget information might be mundane and apolitical, while a request pertaining to aggregate statistics might be aimed at evaluating government performance. However, these divisions nonetheless enable the testing of our second hypothesis by capturing the set of requests that are more likely to trigger political risk-management processes on the part of government officials. As a robustness check (Appendix A), we also analyze only the most clearly private topic—Individual Needs—and the most clearly sensitive topic—Military, Police, and Crime—individually.

Table 2: Twenty Topics of Requested Information

Topic Label	Percent of Requests
Environment and Land	8.64
Individual Needs	8.50
Budgets and Spending	7.98
Commercial Information	7.86
Distributive Programs	7.06
Military, Police, and Crime	6.84
Government Employees 1: Salaries and Benefits	6.71
Government Employees 2: Functions and Qualifications	6.26
Energy and Public Utilities	5.65
Internal Procedures and Official Documents	5.19
Medical Statistics	4.67
Government Employees 3: Specific Personnel	4.10
Public Procurement 1: Service Providers	4.09
Banking, Finance, and Taxes	3.83
Education	3.66
Medical Supplies 1: Contracts and Suppliers	2.82
Public Procurement 2: Procedures and Documents	2.32
Public Procurement 3: Anti-Corruption	2.03
Medical Supplies 2: Purchases and Spending	1.24
Medical Supplies 3: Inventories	0.56

Note: Twenty topics in order from largest to smallest after assigning each request to one primary topic, with the percent of all requests. Topics designated as having higher public potential are shaded in gray.

Modeling Government Responsiveness

Dependent Variable

We are interested in government responsiveness to citizen information requests. In the context of information requests, responsiveness is the extent to which citizens receive the information they seek, in useful form, and in a timely fashion, *except* in cases legitimately subject to legal exemptions from disclosure. Given the complexities in defining and measuring these characteristics, we use information on both the timing and types of responses to triangulate across different possible measures of “non-responsiveness.”

First, we measure the time-to-response in working days. In calculating this variable, we exclude weekends as well as official Mexican government holidays. However, given the presence of legal time limits for responses, we also employ two dichotomous measures of whether or not the response was late. While the statutory deadline was twenty working days, officials could request extension of up to forty working days. We use both of these thresholds to construct separate indicators of late responses. Across the whole dataset, 66.4% of requests received responses within 20 working days, while 89.3% of requests received responses within 40 working days. Although a late response may still provide the requested information, many information requests concern time-sensitive matters, particularly when publicly relevant. Further, delays often reflect a sensitive request that required greater intra-agency negotiation, either with the liaison unit or information committee.

We then employ three different measures of the type of response, based on information provided by officials categorizing the nature of their responses. Here we must confront two issues: First, that requests may be denied for legally allowable reasons, and second, that state agents themselves categorize their response types, and may not always do so accurately. Fox, Haight and Palmer-Rubin (2011), for example, find that some responses (from the 2003-2005 period) were categorized as positive with information delivered electronically, but the actual electronic attachments contain written justifications for *not* delivering the requested information. Nonetheless, they still concluded that for responses classified as positive, “information was actually provided in a satisfactory manner approximately 87 percent of the time” (Fox, Haight and Palmer-Rubin 2011, p. 9).

We thus focus first on responses categorized as “Inexistencia de la información solicitada,” claiming that the requested information does not exist, making up 5.9% of all responses. According to Fox, Haight and Palmer-Rubin (2011, p. 11), this is “the category of denial that leaves the greatest opening for agency opacity.” They found that many agencies denied requests under this category without written justification, concluding that “agencies have discovered that this is the least risky way to deny requests for information in cases where they did not want it released or when assembling it would be a large burden,” especially since “in contrast, the burden of proof is on the agency if it claims that information requested is ‘confidential’ or ‘reserved’” (Fox, Haight and Palmer-Rubin 2011, p. 14). Thus, while a denial for other reasons may well be legitimate under the law, we consider the “Inexistencia” responses to be more clearly examples of unresponsive state agents.

Second, we construct another dependent variable for responses categorized as either “Negativa por ser reservada o confidencial” or “Información parcialmente reservada o confidencial,” indicating that the information requested is either fully or partially classified or confidential. Although many of these denials will be justified, there may be political bias on the margins. These responses account for 3.6% of the total.

Finally, we construct a dependent variable to capture situations where officials do disclose the requested information, but in a non-electronic form that creates additional financial or physical barriers to the requester. Whereas the majority of positive responses are categorized as “Entrega de información en medio electrónico,” meaning the information was delivered electronically, some are physically shipped (requiring the requester to first pay a fee), while others require the requester to physically appear at a government office. These place additional burdens on the requester (and further compromise their anonymity) that in some cases may be prohibitive—indeed several thousand requests appear abandoned in the INFOMEX system, without the requester having ever paid the shipping fee. Lagunes and Pocasangre (2019), discussing the burdensome procedure required to make such payments, suggest that “charging a fee may be a strategy to impede access to information or at least discourage citizens from pursuing a request.” We capture these cases with three categories of responses: “Notificación de envío” (shipping notification), “Notificación lugar y fecha de entrega” (place and date to pick up the in-

formation), and “Respuesta del solicitante a la notificación de entrega de información con costo” (awaiting the requester paying the fee). Together, these types of responses account for 4.9% of the total.

Each of these dependent variables captures different dimensions of responsiveness on which we might expect to see political bias in accordance with our hypotheses. As an initial visualization of these measures of responsiveness in relation to political alignment, Figures 2 and 3 plot each measure against state-level vote shares.¹² Figure 2 plots average responsiveness to requests in each state, over the period up until the 2012 presidential election, against the state-level PAN vote share averaged between the 2000 and 2006 presidential elections. Figure 3, in turn, plots average responsiveness over the period after newly elected PRI President Enrique Peña Nieto took office in December 2012, against state-level PRI vote share in the 2012 presidential election. In each case except one, the line-of-best-fit slopes downwards, suggesting that areas of higher governing party support tended to receive better responses. However, our subsequent approach moves beyond the state-level to use the individual request as the unit of analysis.

Independent Variables

Our primary independent variable captures political alignment at the municipality level. For each information request, this variable takes the value of the governing party’s vote share in the municipality of origin, in the previous presidential election. Here, we first match each request to the municipality from which it was sent and to corresponding municipality-level data on presidential election vote shares. The resulting variable, Government Vote Share, reflects PAN vote share from the elections of Vicente Fox (2000-2006) and Felipe Calderón (2006-2012), and PRI vote share from the election of Enrique Peña Nieto (2012-2018).

¹²We omit Morelos, as a Morelos-based group filed thousands of highly unusual requests (see Appendix C) whose more negative responses heavily influenced the state’s averages. While we do not omit Morelos from our subsequent analyses, our results are similar if we do.

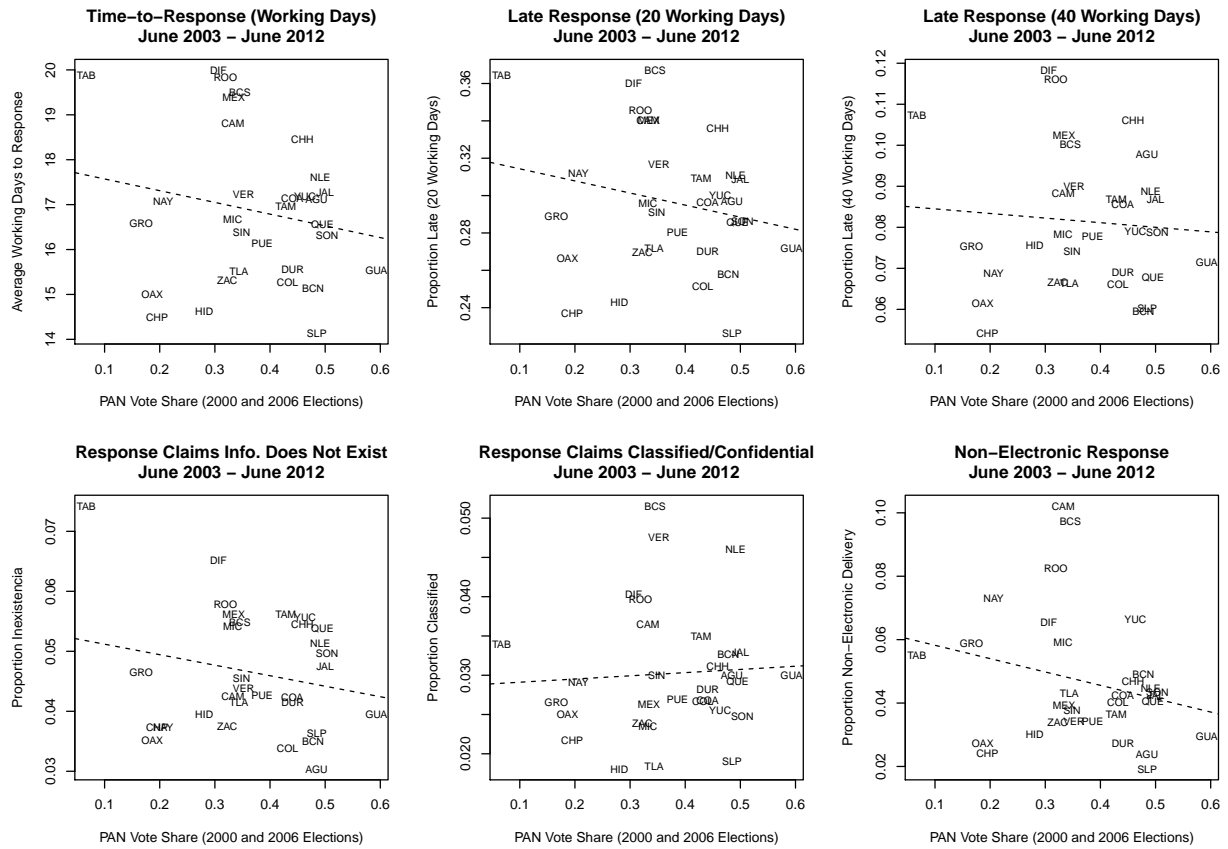


Figure 1: State-level averages for six different measures of responsiveness, plotted against state-level governing party vote shares. Response measures are averages for all requests filed in each state over the period up until the 2012 presidential election. Vote shares are PAN vote share averaged between the 2000 and 2006 presidential elections. Dashed line is a line-of-best-fit from a linear regression of response (state averages) on state vote share. Morelos omitted as an outlier (see footnote 12).




Figure 2: State-level averages for six different measures of responsiveness, plotted against state-level governing party vote shares. Response measures are averages for all requests filed in each state over the period after newly elected PRI president Enrique Peña Nieto took office in December 2012. Vote shares are PRI vote share in the 2012 presidential election. Dashed line is a line-of-best-fit from a linear regression of response (state averages) on state vote share. Morelos omitted as an outlier (see footnote 12).

We control for the logged population of each municipality to avoid confounding by differences in requests from large urban areas and smaller rural locales. We also control for the logged “economic marginality” of each municipality, an index of economic and demographic indicators, to ensure that the effect of governing party political support is not merely proxying for underlying economic differences.

We also include several request-level measures, in order to capture features of requests that may influence the nature of official responses. First, we control for all twenty sets of topic probabilities, reflecting the extent to which each request contains words associated with each of the twenty different potential topics of request texts. We include indicators of whether or not each request included an attached file, and the medium of each request (electronic or physical), although nearly all are electronic. We also measure the (logged) word count of each request, and the (logged) ratio of characters to words. This ratio of characters to words is included as a proxy for “readability,” a measure of the complexity of text.

We also measure two other features of requests that might reflect their complexity or sophistication—legalism and punctuation. Some information requests are highly legalistic, citing provisions of the ATI law or otherwise conveying a level of expertise on the part of the requester that might influence the ways that officials respond. Other requests are clearly filed by individuals with little political sophistication, use very ordinary language, and often omitting punctuation. To capture legalism, we first count the number of times each request includes a series of words or phrases that reflect references to the ATI law or related legal provisions,¹³ and then measure this count as a ratio of the length of each request. We also measure the incidence of punctuation¹⁴ as a ratio of the length of each request.

To measure the investigative risk of individual requests in a more fine-grained manner beyond topics, we constructed a “Corruption Discourse Index” using a dictionary of words drawn from newspaper articles. We first identified over 300,000 articles from two major newspapers

¹³These are: “ley federal de transparencia,” “ley de transparencia,” “artículo,” “fracción,” “sujeto obligado,” and “LFTAIPG”—both with and without accents where relevant and regardless of case.

¹⁴Periods, commas, colons, semicolons, question marks, and inverted question marks.

(*Reforma* and *El Universal*) mentioning any of twenty large federal agencies with unambiguous names and abbreviations. We then developed a set of eighteen search terms pertaining to corruption or scandal, and identified the set of articles containing such terms. We used these to produce a dictionary of 4,568 words that were distinctive to the corruption articles, as compared to the non-corruption articles about the same agencies. Each information request, in turn, was scored with the proportion of words that matched this dictionary. As expected, this measure on its own is strongly associated with worse responses.

We also control for a potentially relevant time-varying factor at the agency level, calculating a measure of “agency workload” to capture the (logged) number requests received over the preceding thirty days. For each agency, this variable takes a different value on each day of the period under study, reflecting the fact that officials facing a larger “queue” of pending requests may respond in different ways. We also include indicators for each month of the year, to capture any seasonal trends in response patterns.

Finally, we include a set of fixed effects for each government agency (thereby making only within-agency comparisons for requests from different regions of the country), each year (to capture across-the-board time trends), and each “dominant” topic category (reflecting the topic with the highest probability of association with each request). Rather than including these fixed effects separately (which we show in Appendix A), we employ a separate fixed effect for every single agency-topic-year combination. That is, our models make comparisons *only* among requests filed with the same agency, on the same topic, and in the same year—but from different parts of the country with varying political alignment. For example, comparisons are made among requests filed with the environmental ministry, on the topic of budgets and spending, in the year 2008—but *not* across different topics of information in requests filed with that same ministry, and *not* across requests for budget and spending information in requests filed with different ministries, and not across different time periods. This distinction is important, as we do not wish to assume that each government agency has the same baseline procedures for responding to different topics of requests, or to assume that changes in internal procedures over time took place according to identical timelines across agencies. These very fine-grained fixed effects automatically capture the effect of any omitted variable that varies across agency-topic-

year combinations, and allow us to more closely approximate the level of comparability that would be obtained in an experimental approach.

Sample and Modeling Approaches

Rather than using the entire sample of over one million requests, we carefully select the set of requests for which our hypotheses are most appropriate. First, we omit the roughly half of all requests originating from the Federal District, where Mexico City is located. Given the preponderance of media, political, legal, and civil society organizations filing requests in the capital, we would be concerned that these requests are less comparable with those in the rest of the country, would make the municipality of origin a less reliable proxy for the likely political alignment of the requester, and would potentially lead to misleading results given the high support for the PRD in the capital city. However, results remain highly similar when we instead retain Federal District observations (Appendix A). Second, we omit agencies that only rarely receive requests, as their responses are likely to be more idiosyncratic. We include only agencies that received, on average, more than one request per week. While this retains only 189 of 304 possible agencies, it removes only 36,074 requests, a small fraction of the total. Results remain highly similar when we retain these (Appendix A). Third, to avoid any bias from “right-censoring,” we remove any requests filed after May 1, 2015 or otherwise unresponded to by our August 2015 cutoff. Finally, we omit all requests filed in the period between July 1, 2012—the date of the 2012 presidential election—and the end of 2012 when the new PRI President Enrique Peña Nieto took office.¹⁵ For purposes of evaluating the effects of political alignment, this period is likely to be idiosyncratic, and it is not clear which party’s supporters we would expect potential political bias to favor. The resulting sample includes 456,747 individual requests.

We employ two main modeling approaches. The first is to model responses as a function of the previously mentioned covariates, making cross-municipality comparisons between requests sent from areas of varying political alignment. By including fixed effects for each agency-topic-year combination, we account for potential unobserved confounders on those dimensions. Our

¹⁵Although the new president took office December 1, 2012, we exclude requests through December 31, 2012, to ensure that indicators for 2012 do not extend across two separate administrations.

second approach instead leverages changes in political alignment over time, adding fixed effects for municipalities and removing fixed effects for each year. This approach takes into account *only* the over-time variation in political alignment between municipalities and the central government, primarily arising from the change in the governing party after 2012.

For both these approaches we employ linear probability models, as employing categorical models with such a large number of fixed effects poses problems for both computation and interpretation.¹⁶

Results

Table 3 presents our main set of results, including the previously discussed control variables and fixed effects for every agency-topic-year combination. For five of the six dependent variables, the coefficient for government vote share is negative and statistically significant at $p < .05$. Only for classified/confidential responses is the coefficient not significant. Notably, the inclusion of fixed effects means that this analysis already accounts for any unobserved variation at the level of agencies, topics, years, or any combinations thereof, as well as controlling for other municipality-level factors, request-level complexity, sensitivity, and sophistication, and the time-varying workload of each agency. These results offer strong support in favor of the *Partisan Discretion Hypothesis*. Ministry personnel do in fact appear to exercise discretion in responding to information requests, favoring request from areas of core support. But which political logic drives this discretionary behavior—favoring allies with access to benefits, or keeping potentially damaging information out of the hands of opponents?

To adjudicate between the *Targeting Benefits Hypothesis* and the *Mitigating Risk Hypothesis*, we examine whether political effects of government vote share appear differentially for topics of requests with higher and lower public potential. Table 4 presents results of models restricting the sample to only topics we categorized as having higher public potential (this and all subsequent tables display only the main coefficients of interest, although all controls from Table 3 are still included). All negative coefficients are greater in magnitude. Table 5 restricts the sample to only

¹⁶In Appendix A, we show logit results for robustness. In all models, we employ two-way clustered standard errors by both municipality and agency.

topics categorized as having lower public potential. In this case, only one of the coefficients is statistically significant.¹⁷ These results suggest that greater political targeting of responses takes place for topics of information that pose potential political risks, and not for politically mundane or purely private topics of information. This finding is thus more consistent with the *Mitigating Risk Hypothesis* than the *Targeting Benefits Hypothesis*. In Appendix A, we also show results employing an interaction term between Government Vote Share and the indicator of higher public potential. The coefficient on the interaction term is negative in all cases, and statistically significant for four of the six models.

In Appendix A, we present additional robustness checks for our main findings from Table 4, including the following. First, we include in the sample requests from the Federal District, which do not substantially change our results. Second, we measure the main independent variable in several alternate ways, including using state-level rather than municipal-level vote shares (in case officials incorporate information about political alignment at a more geographically aggregated level), or measuring the governing party's vote margin over the next-largest party rather than vote share, or using vote shares from mayoral elections in each municipality (though these are more limited in data availability) rather than in presidential elections. Results from these models remain largely in line with our primary results. Third, we control for measures that capture the extent to which a locale represents a "swing" district, as a potential alternate hypothesis to ours might be that officials target better responses to such areas, in attempt to provide valuable services or better demonstrate government performance. We alternately use the Effective Number of Parties in presidential vote shares of each municipality, or the closeness of the vote margin between the nationally-governing party and the largest other party. In either case, the coefficient for the "swing" variable is close to zero, and not statistically significant, in all six models.

¹⁷A coefficient equality test shows that the coefficients for government vote share in Tables 4 and 5 are significantly different from each other for Models 1, 2, and 6.

Dependent Variable	Model 1 LogDays	Model 2 Late20	Model 3 Late40	Model 4 Inexist.	Model 5 Classif.	Model 6 NonElec.
Gov. Vote Share	-0.124** (0.038)	-0.030* (0.013)	-0.022** (0.008)	-0.017** (0.006)	0.004 (0.004)	-0.012* (0.005)
Log Population	0.007 (0.006)	-0.001 (0.002)	-0.005 (0.003)	-0.002 (0.002)	-0.000 (0.001)	0.002* (0.001)
Log Econ. Marginality	-0.049** (0.015)	-0.016** (0.005)	-0.016** (0.004)	-0.005* (0.002)	-0.002 (0.001)	-0.004 (0.003)
Request Length	0.189** (0.019)	0.065** (0.007)	0.032** (0.003)	0.010** (0.001)	0.007** (0.002)	0.010** (0.002)
Request Readability	0.106** (0.032)	0.045** (0.014)	0.030** (0.007)	-0.006 (0.005)	0.015* (0.007)	0.033** (0.010)
Request with Attachment	-0.169** (0.018)	-0.043** (0.007)	-0.017** (0.006)	-0.007 (0.005)	-0.011** (0.004)	-0.003 (0.005)
Request Medium	-0.038 (0.054)	-0.071** (0.021)	-0.023** (0.008)	-0.013* (0.005)	0.011 (0.009)	0.070 (0.038)
Request Legalism	1.060 (0.656)	0.521 (0.330)	0.219 (0.184)	0.369 (0.259)	0.183 (0.179)	-0.060 (0.211)
Request Punctuation	0.365** (0.089)	0.060 (0.033)	-0.009 (0.026)	-0.008 (0.016)	0.001 (0.016)	0.028 (0.016)
Request Corruption Discourse	0.154* (0.069)	0.097** (0.026)	0.055** (0.010)	0.014 (0.013)	0.048** (0.012)	0.002 (0.009)
Log Agency Workload	0.057** (0.012)	0.030** (0.005)	0.014** (0.005)	0.008* (0.004)	-0.003 (0.003)	-0.008** (0.002)
DV Mean	2.561	0.332	0.100	0.056	0.032	0.042
Num. obs.	456747	456747	456747	456747	456747	456747
Adj. R ²	0.323	0.240	0.161	0.123	0.092	0.121

** $p < 0.01$, * $p < 0.05$

Table 3: Full models of government responsiveness to information requests. Models additionally include twenty topic probabilities, monthly indicators, and fixed effects for every agency. Two-way clustered standard errors on municipality and agency.

	Model 1	Model 2	Model 3	Model 4	Model 5	Model 6
Dependent Variable	LogDays	Late20	Late40	Inexist.	Classif.	NonElec.
Gov. Vote Share	-0.162** (0.040)	-0.049** (0.014)	-0.030** (0.009)	-0.019** (0.007)	-0.000 (0.005)	-0.021** (0.004)
DV Mean	2.593	0.346	0.109	0.062	0.033	0.038
Num. obs.	305115	305115	305115	305115	305115	305115
Adj. R ²	0.334	0.251	0.173	0.121	0.096	0.098

** $p < 0.01$, * $p < 0.05$

Table 4: Models of government responsiveness to information requests, with sample restricted to topics with higher public potential. Models additionally include the same control variables as in Table 3, twenty topic probabilities, monthly indicators, and fixed effects for every agency-topic-year combination. Two-way clustered standard errors on municipality and agency.

	Model 1	Model 2	Model 3	Model 4	Model 5	Model 6
Dependent Variable	LogDays	Late20	Late40	Inexist.	Classif.	NonElec.
Gov. Vote Share	-0.052 (0.053)	0.005 (0.019)	-0.009 (0.012)	-0.014* (0.006)	0.012 (0.007)	0.006 (0.011)
DV Mean	2.500	0.303	0.083	0.045	0.030	0.050
Num. obs.	151632	151632	151632	151632	151632	151632
Adj. R ²	0.300	0.213	0.129	0.127	0.088	0.158

** $p < 0.01$, * $p < 0.05$

Table 5: Models of government responsiveness to information requests, with sample restricted to topics with lower public potential. Models additionally include the same control variables as in Table 3, twenty topic probabilities, monthly indicators, and fixed effects for every agency-topic-year combination. Two-way clustered standard errors on municipality and agency.

Fourth, we use alternative topic models, setting the number of topics to 19 or 21, instead of 20, to address potential concerns over topic instability (Wilkerson and Casas 2017). Fifth, we employ separate fixed effects for each agency, topic, and year (rather than for every combination thereof). Sixth, we instead use more fine-grained fixed effects that make comparisons only among requests with identical agency, topic, year *and* similar ‘level of difficulty.’ Seventh, we take into account potential relationships among the different outcome variables by controlling for measures of response type in our models of response time, and vice versa. Finally, we include additional request text-based control variables, assessing the prevalence of misspellings and irregular capitalization. In all of these checks, the results remain substantially similar.

We next turn to a different approach, employing municipality fixed effects to account for any unobserved municipal-level characteristics, while also identifying results based solely on over-time changes within each municipality. That is, the coefficient on government vote share in these models is primarily informed by changes at the 2012 election, before which its value reflects the PAN vote share and after which its value reflects the PRI vote share in each municipality. If our first hypothesis is correct, then requests from a largely pro-PAN municipality should receive more favorable responses before 2012, and less favorable responses afterwards, and vice versa for a largely pro-PRI municipality. Table 6 presents the results of models including municipality fixed effects, in addition to the other control variables in the models from Table 3, as well as fixed effects for every agency-topic combination. We do not incorporate fixed effects for year, in order to capture within-municipality variation over time in political alignment. In four of the six models, the coefficient on government vote share is statistically significant, in all four cases at a $p < .01$ level. We further test Hypothesis 2 in the municipality-fixed effects approach (Appendix A), and find that the effects of political alignment are stronger among topics of information with higher public potential, than those with lower public potential.

To assess the substantive sizes of these effects, we compare the responses to hypothetical requests from municipalities with governing party vote shares of 0.25 and 0.75, based on the main results in Table 3. A request from the more government-opposed municipality would be 4.9% more likely to receive a response later than 20 working days, 13.5% more likely to receive a response later than 40 working days, 19.8% more likely to receive a response claiming that

	Model 1	Model 2	Model 3	Model 4	Model 5	Model 6
Gov. Vote Share	-0.291** (0.058)	-0.143** (0.030)	-0.065** (0.020)	-0.035** (0.011)	0.001 (0.007)	-0.014 (0.010)
Dependent Variable	LogDays	Late20	Late40	Inexist.	Classif.	NonElec.
DV Mean	2.561	0.332	0.100	0.056	0.032	0.042
Num. obs.	456747	456747	456747	456747	456747	456747
Adj. R ²	0.263	0.194	0.130	0.086	0.067	0.099

** $p < 0.01$, * $p < 0.05$

Table 6: Models of government responsiveness to information requests, including municipality fixed effects. Models additionally include the same control variables as in Table 3, twenty topic probabilities, monthly indicators, and fixed effects for every agency-topic combination. Two-way clustered standard errors on municipality and agency.

the information does not exist, and 12.8% more likely to receive a non-electronic response.¹⁸ Further, these effects are larger for the sample of topics with higher public potential.

We emphasize that these represent average effects, across a wide variety of agencies, time periods, and topics of information. Many requests will be routine, and pose no chance of activating a consideration of political risk on the part of officials. Only for some subset of requests would we expect any chance of officials considering the identity of the requester or the political alignment of their location in the first place. We are thus capturing differences averaged across a set of requests that falls within the scope conditions of our argument, and a set of requests that falls outside those scope conditions. While we differentiate between two sets of requests based on the likely public potential of their topics, this can only capture a partial level of detail.

While we do not wish to engage in “cherry-picking” subsets of our data without a priori theoretical reasons, we do offer one illustrative example of a subset of information requests with results indicating much stronger effects of local political alignment. In Appendix A, we present results for a sample of only requests filed with *Petróleos Mexicanos*, or PEMEX, the state oil company (and again omitting requests filed from the Federal District). All coefficients are much larger in magnitude than in the main results from Table 2. Comparing hypothetical requests

¹⁸Percentage increases on the baseline probability of each outcome. We skip classified responses as we find no significant effects in line with our expectations for this outcome-type.

from municipalities with varying vote shares as above, our results suggest that a request from the more government-opposed municipality would be 19% more likely to receive a response later than 20 working days, 31.1% more likely to receive a response later than 40 working days, 44.6% more likely to receive a response claiming that the information does not exist, and 78.6% more likely to receive a non-electronic response. This illustration suggests that the extent of political discrimination is much stronger than our primary results suggest for some agencies, types of information, or specific time periods. On the other hand, many other types of requests are routine and likely to be handled in a purely impartial manner, and some entire agencies may be fully professionalized and impartial such that they allow no discrimination whatsoever.

Conclusion

We study responsiveness in the provision of government information, using publicly available data containing every information request filed with Mexican federal government agencies from 2003 to 2015. Similar to studies of distributive policies, we analyze how responsiveness is shaped by discretionary partisan bias. Yet while something of a consensus in the distributive politics literature has emerged that politicians target core supporters with private benefits – useful in clientelistic exchange – more than public goods, we find the opposite. Officials favor areas of core support with information that lends itself to generating *public* accountability and demonstrate little to no partisan bias in the provision of information useful in generating *private* benefits for the requester.

Underlying these findings are multiple countervailing incentives confronting the agents of accountability who are charged with responding to information requests. Although ostensibly institutions of downwards accountability, as well as horizontally accountable to the oversight body IFAI, officials are also accountable upwards to their agency heads and politicians with electoral incentives. These multiple forms of accountability create different motivations for and against disclosure in different circumstances. Favoring disclosure, they may be motivated by an interest in sustaining a positive image of their ministries, career incentives based on performance, sanctions for noncompliance, or personal commitment to transparency. The main incentives disfavoring disclosure are the political risk associated with releasing evidence of corruption or

poor governance to the general public, along with the needs to comply with official exemptions, and cope with limited capacity. While agencies certainly could exercise partisan discretion in delivering information useful for reaping private benefits, they instead exercise caution in the provision of information that has public potential, given the high cost associated with scandals.

By incorporating information on the topics, complexity, sensitivity, and sophistication of requests, and making comparisons only among requests filed with the same government agency, on the same topic, and in the same year, we address concerns over underlying differences in the requests themselves. We find evidence that municipalities with higher vote shares for the President's party receive more favorable responses, across multiple measures of responsiveness. However, this finding holds only for requests on topics of information with higher public potential, and not for more private topics of information.

Our findings demonstrate that, even in a relatively routine mode of citizen-government interaction featuring decentralized decisions made by large numbers of officials, ruling-party interests and the need to manage political risks can generate dynamics of responsiveness favoring areas of core government support. These findings also suggest that, despite formal rules intended to preserve requester anonymity and procedural fairness in Mexico, informal political dynamics shape bureaucratic processes and allow political discrimination to take place. Yet, on the other hand, we find no evidence of a clientelistic logic in the discretionary disclosure of information to confer private benefits. This suggests some success in insulating accountability mechanisms from patronage politics, yet their replacement with partisan bias of a different form.

Our results, we believe, are applicable well beyond the case of Mexico. We would expect to find similar bias in responsiveness to citizen information requests in any system featuring the means (functioning ATI institutions), motive (electoral or institutional sanctions for politicians), and opportunity (bureaucratic discretion and compromised anonymity) for such discrimination. This includes a wide variety of democratic and non-democratic regimes. To be sure, geographic targeting on the basis of electoral returns is not the only strategy available to officials to discriminate in disclosure; they may also base decisions on other cues that signal the provenance of the request, such as the sophistication of the language, signals of insider knowledge, or even the precise identity of the requester—as in Roberts' (2005) study of Canada. However, given

the politically damaging nature of corruption revelations in even the most mature democracies, we expect the filtering of potentially damaging information to be a primary consideration in discretionary disclosure.

Reforms aimed at minimizing such abuses could include measures to insulate bureaucrats charged with handling information requests from partisan motives, strengthen oversight agencies to detect and sanction noncompliance, and better procedural protections for citizen anonymity. However, such policies are typically much easier to propose than to meaningfully establish.

More broadly, the problem of discretionary responsiveness in institutions of accountability is a particularly daunting challenge for all democracies, transitional democracies in particular. Institutions created or strengthened through “second-generation” democratization—such as anti-corruption agencies, state auditors, and judiciaries—have proven quite resistant to attempts to augment their autonomy and aggressiveness (O’Donnell 1998). Our findings extend this concern to transparency institutions, demonstrating how an institution designed to empower citizens to hold politicians accountable can be utilized by state actors as a tool of selective responsiveness.

References

- Albertus, Michael. 2012. “Vote buying with multiple distributive goods.” *Comparative Political Studies* 46(9):1082–1111.
- Almanzar, Tanya, Mark Aspinwall and David Crow. 2017. “Freedom of information in times of crisis: The case of Mexico’s war on drugs.” *Governance* Forthcoming.
- Berliner, Daniel and Aaron Erlich. 2015. “Competing for Transparency: Political Competition and Institutional Reform in Mexican States.” *American Political Science Review* 109:110–128.
- Berliner, Daniel, Benjamin E Bagozzi and Brian Palmer-Rubin. 2018. “What information do citizens want? Evidence from one million information requests in Mexico.” *World Development* 109:222–235.
- Besley, Timothy and Robin Burgess. 2002. “The Political Economy of Government Responsiveness: Theory and Evidence from India.” *Quarterly Journal of Economics* 117(4):1415–1451.
- Blei, David M., Andrew Ng and Michael Jordan. 2003. “Latent Dirichlet allocation.” *Journal of Machine Learning Research* 3:993–1022.

- Bookman, Zachary and Juan-Pablo Guerrero Amparán. 2009. "Assessing the Implementation of Mexico's Freedom of Information Act." *Mexican Law Review* 1(2):25–49.
- Butler, Daniel and David Broockman. 2011. "Do Politicians Racially Discriminate Against Constituents?" *American Journal of Political Science* 55(3):463–477.
- Carlitz, Ruth D. 2017. "Money Flows, Water Trickles: Understanding Patterns of Decentralized Water Provision in Tanzania." *World Development* 93:16–30.
- Chen, Jidong, Jennifer Pan and Yiqing Xu. 2015. "Sources of Authoritarian Responsiveness: A Field Experiment in China." *American Journal of Political Science* .
- Cleary, Matthew R. 2007. "Electoral Competition, Participation, and Government Responsiveness in Mexico." *American Journal of Political Science* 51(2):283–299.
- Diaz-Cayeros, Alberto, Federico Estevez and Beatriz Magaloni. 2016. *The political logic of poverty relief: Electoral strategies and social policy in Mexico*. Cambridge University Press.
- Distelhorst, Greg and Yue Hou. 2014. "Ingroup Bias in Official Behavior: A National Field Experiment in China." *Quarterly Journal of Political Science* 9(2):203–230.
- Fox, Jonathan and Libby Haight. 2007. Transparency Reforms: Theory and Practice. In *Mexico's Right-to-Know Reforms: Civil Society Perspectives*, ed. Jonathan Fox, Libby Haight, Helena Hofbauer and Tania Sanchez. Mexico City: FUNDAR/Woodrow Wilson International Center.
- Fox, Jonathan, Libby Haight and Brian Palmer-Rubin. 2011. "Proporcionar transparencia Hasta qué punto responde el gobierno mexicano a las solicitudes de información pública?" *Gestión y Política Pública* 20(1):3–61.
- Gans-Morse, Jordan, Sebastian Mazzuca and Simeon Nichter. 2014. "Varieties of clientelism: Machine politics during elections." *American Journal of Political Science* 58(2):415–432.
- Gill, Juliet and Sallie Hughes. 2005. "Bureaucratic Compliance with Mexico's New Access to Information Law." *Critical Studies in Media Communication* 22(2):121–137.
- Golden, Miriam and Brian Min. 2013. "Distributive Politics Around the World." *Annual Review of Political Science* 16(1):73–99.
- Grimmer, Justin. 2012. "A Bayesian Hierarchical Topic Model for Political Texts: Measuring Expressed Agendas in Senate Press Releases." *Political Analysis* 18(1):1–35.
- Guerrero-Amparan, JP and MS Toledo. 2009. "The Right to Information for Marginalized

- Groups: The Experience of Proyecto Comunidades in Mexico: 2005-2007.”.
- Kramon, Eric and Daniel N Posner. 2013. “Who benefits from distributive politics? How the outcome one studies affects the answer one gets.” *Perspectives on Politics* 11(02):461–474.
- Lagunes, Paul and Oscar Pocasangre. 2019. “Dynamic transparency: An audit of Mexico’s Freedom of Information Act.” *Public Administration* 97(1):162–176.
- Magaloni, Beatriz. 2006. *Voting for autocracy: Hegemonic party survival and its demise in Mexico*. Cambridge University Press Cambridge.
- McClendon, Gwyneth. 2015. “Race and Responsiveness: A Field Experiment with South African Politicians.” *The Journal of Experimental Political Science* .
- Michener, Greg. 2011. “FOI Laws Around the World.” *Journal of Democracy* 22(2):145–159.
- Michener, Greg and Ben Worthy. 2015. “The Information-Gathering Matrix: A Framework for Conceptualizing the Use of Freedom of Information Laws.” *Administration & Society* .
- Michener, Greg and Karina Rodrigues. 2015. Assessing Discrimination in Transparency and Freedom of Information Regimes. Global Conference on Transparency Research, Lugano: .
- Mizrahi, Yemile and Marcos Mendiburu. 2012. “Implementing Right to Information: A Case Study of Mexico.”. The World Bank.
- O’Donnell, Guillermo. 1998. “Horizontal Accountability in New Democracies.” *Journal of Democracy* 9(3):112–126.
- Palmer-Rubin, Brian. 2016. “Interest organizations and distributive politics: Small-business subsidies in Mexico.” *World Development* 84:97–117.
- Pinto, Juliet G. 2009. “Transparency Policy Initiatives in Latin America: Understanding Policy Outcomes from an Institutional Perspective.” *Communication Law and Policy* 14(1):41–71.
- Przeworski, Adam, Susan C. Stokes and Bernard Manin. 1999. *Democracy, Accountability, and Representation*. Cambridge University Press.
- Quinn, Kevin M., Burt L. Monroe, Michael Colaresi, Michael H. Crespin and Dragomir R Radev. 2010. “How to Analyze Political Attention with Minimal Assumptions and Costs.” *American Journal of Political Science* 54(1):209–228.
- Ríos Cázares, Alejandra, Carmen Castañeda and Adriana García. 2017. “Transparencia. Diagnóstico Institucional 2016.”.

- Roberts, Alasdair S. 2005. "Spin control and freedom of information: Lessons for the United Kingdom from Canada." *Public Administration* 83(1):1–23.
- Roberts, Margaret E., Brandon M. Stewart, Dustin Tingley, Christopher Lucas, Jetson Leder-Luis, Shana Gadarian, Bethany Albertson and David Rand. 2014. "Structural Topic Models for Open Ended Responses." *American Journal of Political Science* 58(4):172–178.
- Rothstein, Bo and Jan Teorell. 2008. "What Is Quality of Government? A Theory of Impartial Government Institutions." *Governance* 21(2):165–190.
- Schedler, Andreas. 1999. "Conceptualizing accountability." *The self-restraining state: Power and accountability in new democracies* .
- Sobel, David, Bethany Davis Noll, Benjamin Fernandez Bogado, TCG Group and Monroe Price. 2006. "The Federal Institute for Access to Information and a Culture of Transparency."
- Stokes, Susan C, Thad Dunning, Marcelo Nazareno and Valeria Brusco. 2013. *Brokers, Voters, and Clientelism: the puzzle of distributive politics*. Cambridge University Press.
- White, A., N. Nathan and J. Faller. 2015. "What do I need to vote? Bureaucratic discretion and discrimination by local election officials." *American Political Science Review* 109(1):129–142.
- Wilkerson, John and Andreu Casas. 2017. "Large-scale Computerized Text Analysis in Political Science: Opportunities and Challenges." *Annual Review of Political Science* Forthcoming.
- Wood, Abby K. and David E. Lewis. 2017. "Agency Performance Challenges and Agency Politicization." *Journal of Public Administration, Research, and Theory* 27(4):581–595,.

Biographical Statements

Daniel Berliner is an Associate Professor of Political Science and Public Policy at the London School of Economics and Political Science, London, United Kingdom, WC2A 2AD.

Benjamin E. Bagozzi is an Associate Professor of Political Science & International Relations at the University of Delaware, Newark, DE, 19716.

Brian Palmer-Rubin is an Assistant Professor of Political Science at Marquette University, Milwaukee, WI 53233.

Aaron Erlich is an Assistant Professor of Political Science at McGill University, Montreal, Quebec, Canada H3A 2T7.

Acknowledgements

Many colleagues provided invaluable feedback on this project, including James Alt, Andreu Casas, Daniel de Kadt, Bryce Dietrich, Todd Eisenstadt, Daniel Kinderman, Emily Marshall, Greg Michener, Alejandra Ríos Cázares, Caroline Tolbert, and Joachim Wehner. Manuel Gutiérrez provided outstanding research assistance. Earlier versions of this paper were presented at the American Political Science Association in 2015, Midwest Political Science Association in 2017, and in seminars at CIDE, Dickinson College, LSE, University of Delaware, University of Iowa, University of Washington, and Vanderbilt University.