

Democracy and Internet Control: Theory and Evidence from Transparency Reports

Shreyas Meher, Pengfei Zhang*

December 17, 2023

Abstract

Internet control has long been considered a feature of authoritarian regimes alone. Drawing data from Google and Twitter transparency reports, we observe that democratic countries remove an equal amount of content as their authoritarian counterparts. The distinction between the two regimes lies not in the quantity but in the method of content removal. Democracy refrains from government takedown and instead delegates the removal right to the users. This paper conjectures that politicians' reputation concern is the key to understanding this phenomenon. To that end, we develop a political agency model that explains the stylized facts and derives testable hypotheses. Using the timing of elections as a natural experiment, we provide supporting evidence that the takedown requests from democratic governments decreased significantly as the election approached. This reputation effect is not observed in authoritarian regimes or other types of requests.

Keywords: Internet Control, Censorship, Pandering, Natural Experiment

*School of Economic, Political and Policy Sciences, The University of Texas at Dallas, pengfei.zhang@utdallas.edu

1 Introduction

Previous literature views internet filtering either as a dichotomy (e.g., censorship) or as a scale (e.g., freedom of speech). However, democratic countries also remove content from the internet, and government behavior on internet filtering is *multi-dimensional*:

- different conception of the internet
- different types of content they filter
- different bundles of rights for the internet users
- different focuses of their content restriction strategies

This paper provides a taxonomy of the different roles governments around the world have played in internet filtering. We explore the cross-country measurements of internet filtering, classify the types of government behaviors based on a data-driven approach, and identify the key trade-off for national internet-filtering strategies. Prior literature has approached Internet censorship with a binary lens—countries either censor the Internet or they do not—or by using a scalar measurement of freedom of expression. However, this dichotomy is overly simplistic and fails to capture the complexity inherent in government interactions with the Internet. Democratic nations, often perceived as bastions of free speech, also engage in various forms of Internet content removal, highlighting that state behavior in the realm of digital information control is multidimensional and nuanced. This paper challenges traditional models by offering a more granular taxonomy of government roles in Internet filtering, reflecting the intricate layers of control exercised by states over the digital landscape.

Why do we strive for a ‘free’ internet in the first place? With the growth of digital technology permeating every facet of modern life, this question becomes increasingly relevant. The pursuit of a ‘free’ internet is rooted in several key principles and ideals that are fundamental to democratic societies and have been shaped by various political economy theories. The concept of a ‘free’ internet is closely aligned with the principle of freedom of information, a cornerstone of democratic theory. This principle asserts

that access to information is a fundamental human right, crucial for the functioning of a democratic society. It enables citizens to make informed decisions, holds governments accountable, and fosters an environment where ideas can be exchanged freely.

Political economy theories, especially those influenced by Habermas's concept of the public sphere, posit that a 'free' internet is crucial for a healthy democratic discourse. The internet, in its ideal form, is viewed as a digital public sphere where citizens can engage in debate and discussion, shaping public opinion and contributing to the democratic process. This aligns with the principles of deliberative democracy, where the quality of democratic governance is linked to the opportunity for all voices to be heard and considered. However, the pursuit of a 'free' internet is not without challenges. Issues such as misinformation, digital divides, and the concentration of power in a few tech giants pose significant threats. These challenges have led to debates about the role of regulation and governance in the digital space. One of the central debates in the political economy of the internet revolves around balancing freedom and the need for regulation. While too much control can stifle freedom of expression and innovation, the absence of regulation can lead to the spread of harmful content and the exploitation of personal data. Finding the right balance is a key focus of contemporary political economy discussions.

Governments have increasingly focused on sophisticated 'third-generation' content restriction strategies that go beyond mere blocking and filtering. These include shaping online discourse through orchestrated information campaigns, deploying surveillance tools to deter undesirable online behavior, and leveraging regulatory frameworks to enforce content moderation on platforms. This research explores the multifaceted nature of government Internet filtering practices across various countries, utilizing cross-country measurements of filtering and employing a data-driven approach to classify types of government behaviors. Our analysis aims to uncover the underlying trade-offs that nations grapple with when devising their national Internet-filtering strategies. What emerges is a spectrum of governmental stances on digital information control, with varied implications for global Internet governance.

The transformative power of the Internet has fundamentally altered access to information, catalyzed political movements, and reshaped social interactions on a global

scale (Zarras, 2016; Hellmeier, 2016). Yet, despite its inherently decentralized architecture, which transcends traditional geographical boundaries (Mueller, 2010), the Internet has become the focal point of regulatory oversight and control by nation-states (Akdeniz and Altiparmak, 2008; Goldsmith, 2007; Mueller, 2010). This paper scrutinizes the evolution of internet filtering, advancing from the early notions of its impracticability (Bambauer, 2009) to the sophisticated, multifaceted strategies employed by both authoritarian and democratic regimes alike, blurring the lines between censorship and governance (Bambauer, 2009; Breindl and Kuellmer, 2013; Pearce et al., 2017).

Censorship online is driven by a complex interplay of cultural, historical, religious, and political factors (Akdeniz and Altiparmak, 2008), where governments, rooted in their respective ideological frameworks, curtail access to a variety of content, ranging from political dissent to intellectual property violations (Hellmeier, 2016; Pearce et al., 2017). The role of major social media platforms in shaping political discourse has only intensified the focus on Internet regulation, making them prime targets for state-mandated censorship efforts (Dick et al., 2012; Zarras, 2016). In addressing a gap in existing literature predominantly centered on single-nation studies (Pearce et al., 2017), our cross-country analysis offers a comparative perspective on the mechanisms of Internet censorship within diverse political landscapes.

2 Background and Related Literature

In an era where the digital domain is intrinsically interwoven with the fabric of political discourse, the interplay between electoral cycles and internet censorship presents a compelling yet under-explored dimension of contemporary politics. This literature review aims to delve into the nuanced impact of electoral processes on requests for content removal on online platforms such as Google and Twitter, a phenomenon that lies at the intersection of digital governance, electoral politics, and public policy. Despite the critical importance of this area, especially in democratic societies, it remains a relatively less trodden path in academic research compared to other aspects of internet regulation and political communication. Requests for content removal from governments or political

entities to major online platforms raise fundamental questions about freedom of expression, state control, and the role of digital intermediaries in shaping political narratives. However, literature specifically addressing the impact of electoral cycles on such internet censorship practices is not as common as one might expect. While there is a wealth of research on broader themes like online misinformation, digital propaganda, and the general influence of social media on politics, studies that directly link electoral timelines with state-led attempts to control or remove online content are relatively sparse. This gap in the literature highlights the need for a more focused inquiry into how impending elections might motivate political entities to shape or sanitize the online information landscape.

Beyond outright content blocking, governments often engage in more subtle forms of censorship through surveillance and data control, aligning with Michel Foucault's concept of 'governmentality'. [Foucault \(2007\)](#) suggests that modern states exert control through the management and regulation of populations, including their activities and discourse online. Surveillance tools enable governments to monitor internet usage and communications, thereby discouraging dissent and controlling public discourse. This form of control is not always overt but can be deeply ingrained in the architecture of the internet within a country. For instance, the requirement for data localization, where data about a country's citizens must be stored within the country, allows governments to access and potentially manipulate this information easily. Furthermore, the use of advanced technologies like deep packet inspection allows for more sophisticated surveillance, enabling governments to track and analyze internet traffic at a granular level. Regarding internet blocking, the application of Foucault's theories of power and surveillance is evident. Governments often justify internet censorship and blocking under the guise of maintaining social order, national security, or moral standards, aligning with Foucault's concept of biopower – the control of populations through regulatory mechanisms. This form of control is not only about restricting access to certain information but also about shaping the acceptable discourse within a society. In some regimes, internet blocking is also a tool for obscuring the functioning of power itself, resonating with Foucault's idea of power being omnipresent and insidious. By controlling the flow of information, governments can limit citizens' ability to understand and critique the mechanisms of power,

thus maintaining control with less overt repression.

Our paper interacts and builds upon the literature on electoral accountability, particularly in the realm of internet censorship and government media capture. We build upon and extend the insights provided by pivotal studies such as those by [King et al. \(2013\)](#), [King et al. \(2014\)](#), [Maskin and Tirole \(2004\)](#), [Hess and Orphanides \(1995\)](#), [Rao \(2021\)](#), and [Williams \(2013\)](#), each contributing unique perspectives on how electoral cycles influence government behavior. King et al. (2013, 2014) provide critical insights into the strategies adopted by the Chinese government to suppress collective actions and control political discourse. Their work highlights the influence of electoral pressures even in authoritarian regimes, revealing how such pressures shape policy decisions. Our research complements these findings by examining similar dynamics in democratic contexts, showing that electoral considerations significantly impact government strategies for internet censorship and content takedown, thus broadening the scope of understanding electoral accountability beyond authoritarian regimes.

The interplay between political accountability and judicial independence, as discussed by Maskin, resonates with our findings in democratic settings. Our paper explores how democratic governments adjust their internet control strategies as elections approach, reflecting a strategic balance between maintaining power and adhering to democratic norms. Hess's analysis of conflict behavior in democracies related to electoral cycles, Rao's exploration of electoral accountability in the context of investor-state arbitration, and Williams's examination of flexible election timing on foreign policy in parliamentary democracies collectively illustrate the broad impact of electoral cycles on various aspects of government behavior. Our research enriches this discourse by highlighting the specific impact of these cycles on internet governance policies. We demonstrate that the approach to internet censorship in democracies is not static but varies significantly with the electoral calendar, adding a nuanced understanding of how democratic leaders navigate the complex terrain of digital governance while being responsive to electoral pressures.

Our paper significantly enriches the literature on censorship and media capture by intertwining the themes of electoral accountability and internet governance across var-

ious political regimes. We engage with and extend the insights offered by [Zittrain et al. \(2017\)](#), [Gaubatz \(1991\)](#), [Akdeniz and Altiparmak \(2008\)](#), [Goldsmith \(2007\)](#), [Hellmeier \(2016\)](#), and [Mueller \(2010\)](#), each contributing to the diverse landscape of internet censorship practices globally. Zittrain's exploration of the global trend toward geopolitical filtering and Gaubatz's study on the behaviors of democratic states across different electoral cycle stages provide a broader understanding of governmental control in the digital realm. Our research complements these findings by demonstrating how the timing of elections in democratic countries influences internet censorship policies, suggesting that such strategies are not merely static state actions but are dynamic and responsive to the electoral environment. The contrasting approaches to internet censorship in Turkey and China, as examined by Akdeniz and Goldsmith & Wu, respectively, highlight the differences in online discourse control between democratic and autocratic regimes. Our paper builds upon these studies by offering a model that elucidates how governments, irrespective of their regime type, manage the complex balance between control and public opinion in the realm of internet governance, especially during election cycles. Hellmeier's systematic examination of internet censorship in autocratic regimes and Mueller's comprehensive view of global internet governance add depth to our understanding of how states justify and implement internet censorship policies. Our research contributes to this discourse by providing a comparative analysis across countries, showcasing the varied strategies and motivations behind censorship practices worldwide.

When looking at media capture, [Gehlbach and Sonin \(2014\)](#) presents a theoretical framework for analyzing government control of the media. They highlight how media bias and the incentive for government control depend on the government's internalization of the impact of bias on news consumption. [Schedler \(2010\)](#) complements Gehlbach and Sonin's analysis by discussing how authoritarian governments use manipulation, including media control, to maintain power. This manipulation extends to various institutional arenas, including the judiciary, elections, and the media [Esarey and Xiao \(2011\)](#) study on China's digital communication highlights the state's efforts to control media through commercialization and regulation. Despite the growth of digital media, the Chinese state increased its influence over public opinion as mass media's reach expanded. This is augmented by [Simonov and Rao \(2022\)](#) where they look at government-controlled

online news in Russia. They find that by investing in controlled outlets' quality and non-sensitive content, governments can subtly manipulate public opinion. [Lorentzen \(2014\)](#) further adds to this discussion - providing insights into China's strategic media control, balancing the need for some media freedom against the risk of facilitating widespread discontent. In the context of these themes, our paper focusing on cross-country internet censorship differences, taking into account regime type, adds significant value to the ongoing discourse. It highlights the nuanced ways in which different regimes implement censorship and control media, with a specific focus on the role of the internet as a medium.

[Prat and Strömberg \(2013\)](#) discuss media capture and censorship, emphasizing the role of media pluralism and commercial interests in preventing government control. They highlight the complex interplay between government incentives, media strength, and censorship costs, and how these factors influence policy outcomes and political accountability. [Shadmehr and Bernhardt \(2015\)](#) add to this discussion, examining state censorship in authoritarian regimes, analyzing the strategic decisions of rulers regarding media censorship to control information and prevent revolutions. Their study reveals how advancements in communication technology and media strength can paradoxically benefit authoritarian rulers by forcing them to censor less, thereby maintaining a semblance of media freedom. [Egorov et al. \(2009\)](#) examine the paradox of media freedom in resource-poor dictatorships. The authors discuss how the lack of resources can lead dictators to rely more on public support, consequently allowing greater media freedom as a strategic move to maintain power. This paper adds an economic dimension to the discussion of media control, showing how economic factors can influence a regime's approach to media censorship.

3 Measurement of Internet Control

OpenNet Initiative (ONI) Data The OpenNet Initiative (ONI) conducts country-level audits of government filtering of websites. To identify and document Internet filtering, ONI uses two checklists of websites: a global list and a local list. The global list is comprised

of internationally relevant websites with provocative or objectionable content in English. The local lists are designed individually for each country to document unique filtering and blocking behavior. Local researchers within each country query these pre-defined lists of URLs using a data collection software client. The list of URLs is accessed simultaneously over HTTP both in the country suspected of Internet filtering and a control country known for no governmental filtering. Each country is given a score on a five-point scale for the following four dimensions: political score, social score, conflict/security score, tools score.

1. Political - This category focuses on filtering done by the government on websites or content that relates to political freedom. In particular ONI looks at whether views that are opposing to the ruling party are filtered or not. Other content within this category relate to human rights, rights for political and ethnic minorities along with religious movements.
2. Social - This category looks at filtering on content related to socially sensitive topics such as gambling, pornography, drugs among others. Any topic that is socially sensitive and/or offensive is included within this type of filtering in the ONI data.
3. Conflict/Security - This category looks at content that relates to armed conflicts, disputes near borders along with terrorist groups and separatist movements within the country.
4. Internet Tools - This category scores countries on the filtering done on various internet tools such as email messaging, Voice-over Internet Protocol (VoIP), search engines. This category also looks at whether Virtual Private Networks (VPNs) along with other circumvention methods are available to citizens of a country.

The OpenNet Initiative organization not only provides the dataset used for numerous statistical analysis but also publishes country profiles that offer specific instances and a qualitative perspective on internet filtering efforts undertaken by various nations within the data. ONI's tests do not offer real-time tracking of Internet filtering; they provide snapshots of accessibility at specific points in time. They are also not exhaustive, and

they complement other projects such as Herdict, Google Transparency Report, Chilling Effects, and more. The absence of a country from ONI’s results does not indicate a lack of filtering, and reports of censorship in a country without empirical testing data from ONI should not be dismissed. ONI’s evaluations are based on national filtering regimes and do not account for filtering on private or institutional networks.

Variable/Year	2012	2014	2016	2018	2020
Average FOTN	53.425532	54.47692	53.36923	52.75385	51.69231
Average ONI	6.333333	-	-	-	-
Google Items	352.175000	361.75781	1179.96479	2834.71053	1147.73140
Google Requests	34.166667	53.47656	158.55634	282.60526	97.08678
Twitter Accounts	9.875000	94.26190	404.73494	640.62791	1993.48624
Twitter Requests	3.000000	29.23810	133.97590	266.64773	727.42342
Number of Countries	74	74	74	74	74

Table 1: Summary Statistics

Freedom On the Net (FOTN) Data Freedom on the Net is Freedom House’s annual survey of internet freedom around the world. It measures the different ways in which governments around the world restrict users’ rights online. Each country is assessed on 21 questions and is assigned a numerical score. The questions are divided into the following three categories: Obstacles to access, content limitations, and violation of user rights. We transform the scores into an anti-freedom index by subtracting the original score from the maximum score in each category. The FOTN is a much more extensive than the ONI index, with data ranging from 2009 to 2022. The FOTN also contains data for 70 countries over this time period on average.


Google Transparency Report Another dataset that we use is the Google Transparency Report, where Google publishes detailed information on requests by countries or third parties to remove content on their various platforms. We use this data based on the availability and research has shown that Google Transparency Reports, more so than other similar reports by Microsoft, Apple, Reddit, etc. follows the Santa Clara Principles (SCP) on Transparency and Accountability in Content Moderation ([Urman and Makhortykh, 2023](#)).

There are two main types of requesters: Government Officials and Court Orders. From 2019, for some countries, Google also records requests from special agencies, such as Police, the Information and Communication Authority, the Consumer Protection Authority, the Data Protection Authority. We count requests from Government Officials as Government Requests so as to remove requests that are not sent for censorship reasons such as copyright infringement out of our analysis.


We utilize three key variables from the Google Transparency Reports in our analysis: the Number of Requests, the Quantity of Items Requested for Removal, and the Type of Requester. Requester types in these reports are manifold, but for the purpose of simplification and clearer understanding, we've categorized them into two broad entities: Governmental Organizations and Court Orders. The primary motive behind this categorization is to gain insights into how distinct clusters of nations employ different mechanisms for Internet content filtration. Moreover, the requests directed towards Google influence a plethora of its services including but not limited to YouTube, Web Search, and Blogger. This factor is of substantial significance, aligning with our conceptual framework that pertains to the degree of directness, or its absence, in Internet filtering across the spectrum of our designated clusters.

Figures 1(a) & 1(b) plot the total number of requests across the time period and items requested to be removed of the Google Transparency Report, where the data is combined for both waves of the report and goes from 2009 to 2022. We can see that there is a noticeable uptick in both requests and items requested to be removed after 2015 and the trend remains constant. Through these plots we see that there has been a noticeable increase in usage of the reporting mechanisms of such intermediaries to filter and censor the internet. In future sections, we will break this down by the two regimes that we define using the Regimes of War (RoW) data.

Twitter Transparency Report Similar to the Google Transparency Reports, Twitter also disseminates its transparency reports as part of a common practice among social media platforms and Internet intermediaries. Our research approach treats Twitter's data with a comparative lens to that of Google's, with an emphasis on discerning variances in the efficacy of removal rates across various clusters. The principal aim of these trans-



(a) Total Number of Requests across time



(b) Total Items Requested to be removed across time

Figure 1: Google Transparency Report - Request Trends & Removals Over Time

parency reports is to foster openness and accountability, as well as to furnish an array of stakeholders - including users, policymakers, and other interested parties - with insights into Twitter’s governance, policies, and operational responses.

Specifically, we concentrate on Removal Requests within Twitter’s reports. This encapsulates requests pertaining to both account as well as content deletion. The Removal Requests offer a perspective into the volume of content removal pleas that Twitter fields from governments, organizations, and individual entities. Moreover, it highlights how Twitter handles these entreaties, which may encompass demands to eliminate illegal content or content infringing upon Twitter’s terms of service. An additional component of the report provides statistical information regarding the number of accounts and tweets that have been impacted by such removal actions. Through this study, we aim to elucidate the dynamics of content regulation on the widely-used social media platform, thereby enriching the scholarly understanding of digital communication norms, platform governance, and the challenges of upholding digital rights in an era characterized by widespread content creation and sharing.

Figures 2(a) & 2(b) plot the total number of combined requests and accounts specified over the time period of the data, which ranges from 2009 to 2022. The combined request variable is a combination of Court Orders and Legal Demands that are sent to Twitter for removal. The accounts specified variable does the same for the accounts that are requested for Twitter to flag or remove. As we see from both the plots, there is a steady upward trajectory for both of these variables - indicating the greater usage of states and individuals in using intermediary mechanisms and tools to filter and censor

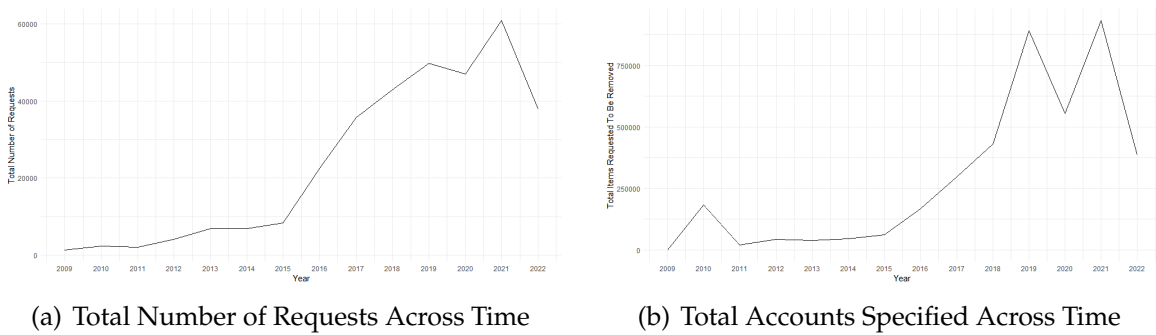


Figure 2: Twitter Removal Requests & Accounts Specified

the internet in various different ways. The slight downturn in requests during 2020 is certainly thought-provoking. While it's tempting to attribute this entirely to the global pandemic, with countries reeling under the health crisis and potentially deprioritizing digital surveillance momentarily, there are likely more layers to this narrative. The socio-political landscape of 2020 was tumultuous, with global movements, economic downturns, and a significant portion of global workforce transitioning to remote work. These shifts might have caused governments to temporarily reevaluate or redistribute their monitoring resources. It also opens up the possibility that, with more users turning to online platforms for information during the pandemic, misinformation might have become more diffused, making it challenging for governments to pinpoint and report.

Other Supplementary Data Sources Other data sources include the Web Index, World Development Indicator by the World Bank, World ICT Indicators by the International Telecommunication Union, etc.

4 Stylized Facts

The initial preprocessing of the data involved a few steps to ensure a coherent and concise dataset for subsequent analysis. For the Google data, we combined the categories in the 'Requester' column into two broad categories: "Court Orders", combining "Court Order Directed at Google" and "Court Order Directed at 3rd Party", and "Government Requests", which consist of request by Government Officials. This allowed us to focus

on these two significant sources of content removal requests. We excluded the category of "Other" in the requester variable due to its indistinct nature, ensuring a clear distinction between legal and government requests. The efficiency of these requests was subsequently calculated for both "Court Orders" and "Government Requests". The next phase of our analysis involved assessing the temporal changes in request efficiency. We plotted the average "% Efficiency of Requests" against the "Period Ending" variable, separately for "Court Orders" and "Government Requests". This gave us insights into how the efficiency of these request types varied over time, within each regulatory cluster. The Twitter data was used for a similar analysis. However, it provided more granular metrics, such as "Combined Accounts Specified" and "Combined Requests", which were used to construct corresponding time series plots. This enabled a more detailed examination of trends in content regulation requests on Twitter.

Analyzing the content removal trends on prominent online platforms like Google and Twitter provides an intriguing lens to gauge how governments, across different political regimes, seek to manage or control online narratives. This analysis provides valuable insights into the changing dynamics of online content control, as depicted in the various figures and tables presented.

4.1 Democracies remove an equal amount of content as Autocracies

Regime	Requests	Items Removed	Avg. Items
Autocracy	240,175	2,390,667	9.95
Democracy	87,894	1,621,191	18.4
Overall	328,336	4,014,206	12.2
Number of Countries	157	157	157

Table 2: Summary Statistics of the Google Transparency Report

The shift from democracies to autocracies in sending content removal requests to Google, as depicted in Figure 3(a), isn't just a statistical trend — it narrates the evolving digital landscape. In the early 2010s, the internet was rapidly becoming a primary medium for expression, especially in developed democracies where the penetration and

acceptance of digital platforms were high. Democracies, with their inherent emphasis on freedom of expression, might have sent removal requests mainly to regulate potentially harmful content or to maintain the sanctity of the digital space.

However, the narrative began to shift as autocratic governments recognized the profound influence of the internet on public opinion and political stability. As suggested by the graph, this possibly led to the surge in their content removal requests post the mid-2010s. These governments might perceive unregulated digital content as a threat, prompting rigorous online monitoring and censorship to maintain a preferred narrative.

The analysis of items requested to be removed in Figure 3(b) adds another layer of nuance. While democracies sent fewer requests than autocracies in the latter years, they targeted more items per request, possibly indicating a more specific or focused approach to content regulation. This changing dynamic potentially reflects the increasing urge among autocratic regimes to exert greater control over digital content as the internet expands its reach.

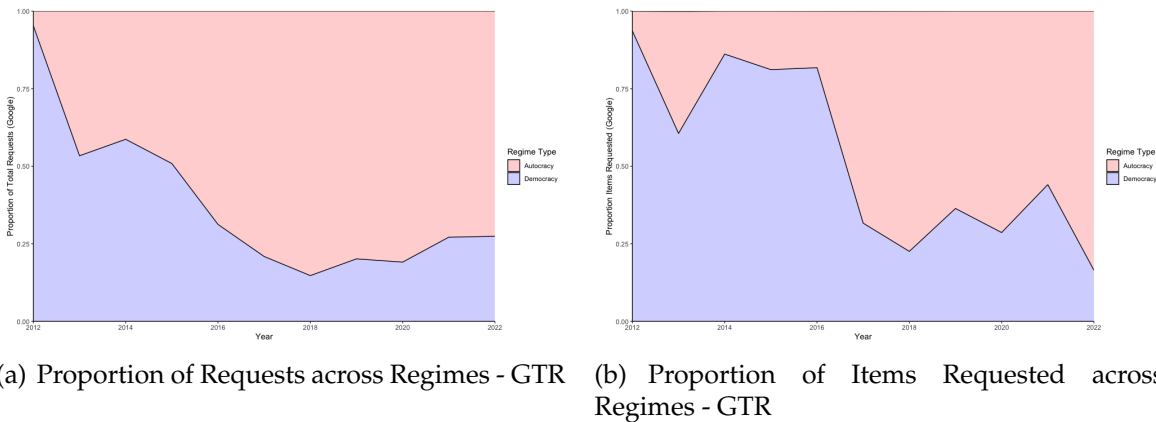


Figure 3: Google Transparency Proportions

Regime	Requests	Accounts	Compliance Rate
Autocracy	141,600	341,461	38.5%
Democracy	130,997	554,594	38.7%
Overall	272,597	896,055	38.6%
Number of Countries	91	91	91

Table 3: Summary Statistics of the Twitter Transparency Report

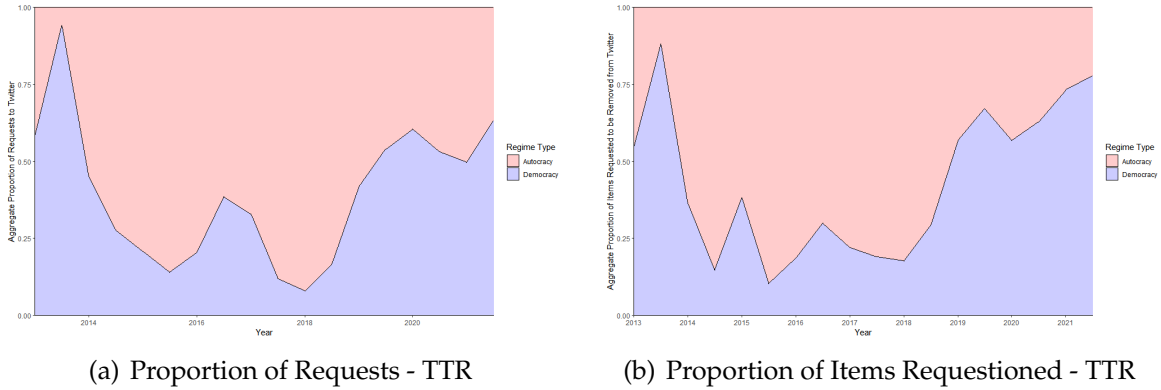


Figure 4: Twitter Transparency Proportions

The proportion graphs, as depicted in Figures 4(a) and 4(b), not only serve as a statistical representation but also as a mirror reflecting the evolving priorities of democracies worldwide. The surge post-2016 is not merely a statistical anomaly but perhaps a direct consequence of the upheavals and challenges that democracies faced during this period. Given the rise of international incidents, coupled with growing concerns over foreign interference in electoral processes, the imperative for democracies to secure their digital frontiers has never been more pressing. These trends might also be indicative of democracies acknowledging the malleable nature of public opinion in the digital age, with the realization that narratives, whether factual or fabricated, can shape electoral outcomes, policy directions, and even diplomatic relations.

Furthermore, the emphasis on regulating content, especially that which borders on hate speech, cyberbullying, or extremism, might be democracies' response to the societal demand for safer online spaces. As digital platforms transform from mere communication tools to significant components of citizens' daily lives, there's an increasing call for them to be accountable and conducive to mental and societal well-being. The digital age, while bringing the world closer, has also spotlighted the darker aspects of human interaction, necessitating a more significant push by governments to curb malicious online behaviors that threaten societal harmony. As Twitter evolves from a microblogging site to an indispensable conduit for news dissemination, its heightened relevance has undoubtedly caught the attention of global powers. Recognizing Twitter as a pulse-check for societal sentiments, governments are perhaps trying to ensure that the information flowing through this channel is in sync with their national interests. This shift in percep-

tion of Twitter, from a mere social platform to a digital powerhouse, has reshaped how states strategize their digital diplomacy and public relations efforts. No longer can governments afford to overlook the influence of such platforms; they need to be proactive, adaptive, and strategic in their engagement.

4.2 Autocracies Censor, but Democracies Delegate to the Users

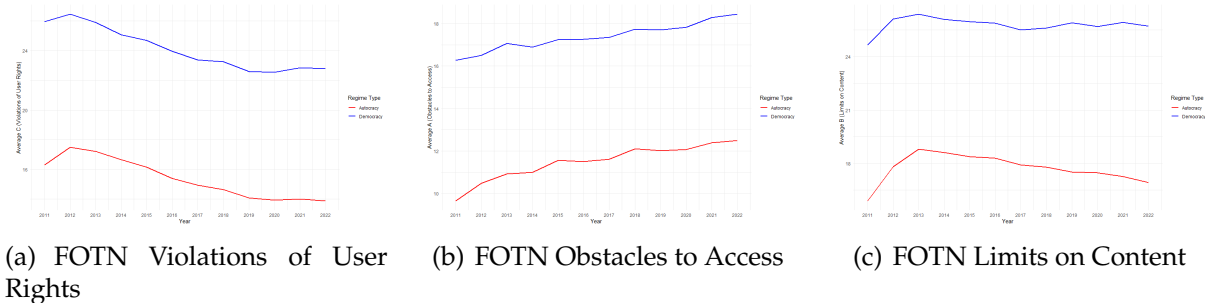



Figure 5: FOTN Index - Trends Over Time

The graph in Figure 5(a) provides a compelling illustration of the divergence between democracies and autocracies concerning the Freedom on the Net (FOTN) Violations of User Rights over the span of a decade. It is evident that democracies, represented by the blue line, have generally maintained a higher standard of user rights, experiencing a slight decline from 2011 to 2022. Conversely, autocracies, depicted by the red line, present a pronounced decline over the years, reflecting an increased violation of user rights. This disparity underscores the fundamental differences in the approach and value systems between these two governance models when it comes to online user rights.

Delving into the Obstacles to Access, as visualized in Figure 5(b), we observe a similar trend. Democracies have remained relatively stable over the years, with minor fluctuations, ensuring that access to the internet remains largely unhindered. On the contrary, autocratic regimes show an increasing trend in obstacles to access from 2011 to around 2016, with a subsequent plateau. This trend possibly indicates tightening controls and restrictive measures implemented by autocratic governments to regulate and monitor internet access among their populations. Figure 5(c) delves into the limits imposed on online content. Once again, democracies have maintained a more liberal stance, with minimal

variations in content restrictions over the observed period. In stark contrast, autocracies present a more pronounced decline, indicating heightened censorship and limitation of online content. This data, in tandem with the previous metrics, underscores the restrictive nature of autocratic regimes, further limiting freedom of expression and access to information.

Taken collectively, the trends observed in the FOTN Index over time provide a comprehensive perspective on the digital liberties and restrictions prevalent in different governance systems. The data showcases the inherent virtues of democratic governance in preserving online freedoms, while simultaneously highlighting the suppressive tendencies of autocratic regimes. These findings serve as a testament to the importance of fostering and preserving internet freedom as a pivotal aspect of modern civil liberties.




(a) Number of Requests - Gov. Officials vs Court Orders (b) Number of Requests - Gov. Officials vs Court Orders


Figure 6: Government Requests vs Court Orders

The compliance rate graphs deepen our understanding of the relationship between these platforms and governments. Despite the larger number of requests from autocracies, their compliance rate for government requests isn't substantially higher (as observed in 7(b)). This might indicate Google's hesitancy or stringent review process when dealing with requests from regimes where censorship might be used as a political tool.

When these removal requests are broken down by their origin—government requests versus court orders—a stark contrast is evident from Figures 6(a) and 6(b). The pronounced reliance of autocracies on direct government mandates speaks volumes about their governance mechanisms. The relative bypassing of judicial routes can suggest ei-



(a) Compliance Rate - Court Orders (Google)



(b) Compliance Rate - Government Requests (Google)

Figure 7: Compliance Rate - Google

ther an overarching government control, wherein judiciary’s independent checks might be limited, or a more streamlined, albeit centralized, digital governance model in these countries. Democracies, with their more pronounced separation of powers, seem to employ the judiciary actively, ensuring that content removal doesn’t infringe upon fundamental rights.

The compliance rate graphs serve not just as statistical indicators but as mirrors to the intricate dynamics that exist between tech behemoths and sovereign states. On the surface, one might anticipate that Google, being a commercial entity, would be more compliant to larger numbers of requests, especially from autocratic regimes where non-compliance could result in significant repercussions, such as service bans or financial penalties. However, the data, as illustrated in 7(b), suggests otherwise. Despite receiving a deluge of requests from autocratic regimes, the non-disproportionate compliance rate underscores Google’s commitment to its foundational principles, which emphasize open access to information. This isn’t to say that Google adopts an adversarial stance against such governments, but rather highlights the company’s complex position. On one hand, Google must respect local laws and sensibilities to operate in a country; on the other, it has its own set of values and a global user base that demands adherence to principles of open information and privacy.

Moreover, the observed compliance rates might also be a product of the nature of requests. Autocracies, in their urgency to control narratives, might send requests that

are broad, generalized, or not well-substantiated. Google’s review mechanisms, which prioritize user privacy and freedom of information, would likely subject such requests to rigorous scrutiny, leading to a lower compliance rate. Additionally, Google’s stance can also be seen as part of a broader narrative where tech companies are increasingly being recognized not just as commercial entities but as significant stakeholders in the global discourse on freedom, privacy, and rights in the digital age. By not disproportionately complying with autocratic requests, Google indirectly emphasizes the importance of balancing governance with digital rights.

In the broader ecosystem, this dynamic between Google and autocratic regimes becomes particularly salient. Tech companies today wield significant influence, often comparable to nation-states, and their policies can have ripple effects across political, social, and cultural domains. The compliance rates, hence, aren’t just transactional statistics but indicators of a more profound dialogue on the ethics, responsibilities, and boundaries of digital governance in the 21st century.

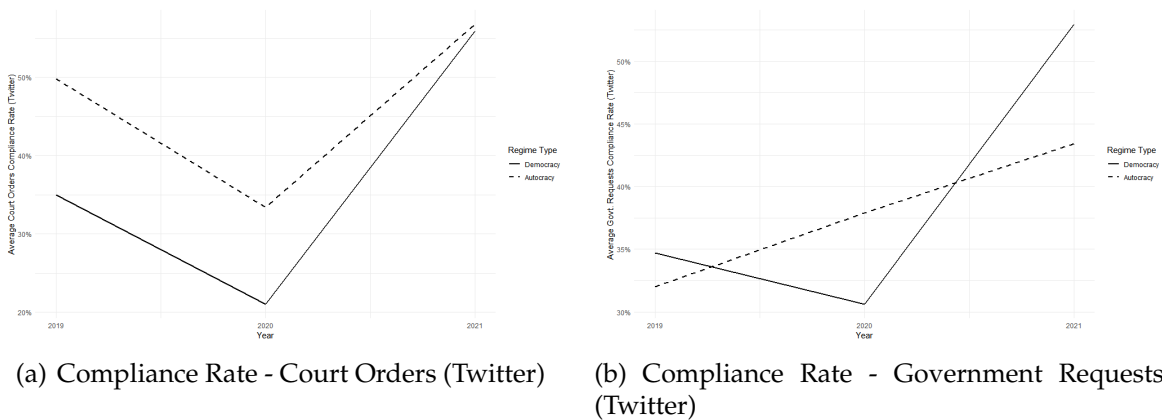


Figure 8: Compliance Rate - Twitter

Drawing parallels to the Google data, the Twitter compliance rate graphs 8(a) & 8(b) offer additional insights into the platform’s interactions with governmental bodies. The figures reveal the nuanced dance of influence and resistance between a global social media platform and the varying government regimes it engages with.

In autocratic regimes, Twitter’s lower compliance rate for government requests, akin to Google’s, might signify a cautious approach toward demands that may potentially sti-

file freedom of speech or contribute to political censorship. Despite potentially facing harsh penalties for non-compliance, such as platform bans or legal challenges, Twitter seems to exhibit a commitment to upholding a certain standard of open communication. The pattern displayed suggests a rigorous internal review process that filters out requests that do not align with international norms or the company's policies on user rights and data privacy. On the other hand, democracies show a fluctuating yet distinct compliance rate trend that suggests a more systematic legal engagement. Democracies' higher reliance on court orders, as opposed to direct government requests, may reflect a stronger rule of law where due process is followed, and removal requests are typically funneled through judicial review. This legalistic approach, while potentially more cumbersome for platforms like Twitter, underscores a structural balance between government oversight and the preservation of civil liberties, at least in the context of digital content management.

5 Model

This section develops a simple political agency model to explain our empirical findings. The model explains why democracy delegates the removal to users: reputation concerns.

¹

5.1 Model

There are two politicians (denoted by P), an incumbent and a challenger, and a representative citizen (denoted by C). The incumbent holds office in the first term, but at the beginning of the second term faces an election against the challenger to determine who holds office in the second term. In each term, the politician in office has to choose whether and how to filter the internet, denoted by $x \in \{0, 1\}$. The politician can decide which player has the right to filter the internet: he can either choose whether to remove the content himself $i = P$, or he may also assign the decision rights to the citizen $i = C$.

¹For a classical reference on political agency models, see [Besley \(2006\)](#).

Therefore in this model, content can be removed in two ways: either through government censorship $i = P, x = 1$, or through user moderation $i = C, x = 1$.

The payoffs to the citizen and the politician depend on the state of the internet $\omega \in \{0, 1\}$ indicating whether the internet content is harmful $\omega = 1$ or harmless $\omega = 0$. Let λ be the probability of the prior belief that the internet is harmful. The citizen's payoff is as follows:

$$u_C = \begin{cases} v - h\omega & \text{if } x = 0 \\ 0 & \text{otherwise,} \end{cases} \quad (1)$$

where v is the value of the internet, and h is the magnitude of the harm. There is a signal s that perfectly reveals the state of the internet. The incumbent and the citizen have an equal ex-ante probability of observing the realization of the signal. The player who observes the signal will update her belief regarding ω by Bayes' rule, while the player who does not will keep the prior belief. Which player observes the signal is common knowledge. We denote $s = P$ as the signal observed by the politician and $s = C$ as the signal observed by the citizen.

There are two types of politicians - unbiased or biased - the type is denoted by $\delta \in \{0, 1\}$. An unbiased politician $\delta = 0$ shares the citizens' policy preferences. The biased politician $\delta = 1$ has a private benefit from strict internet control δx . Both types of politician gets a non-policy related return r from holding the office. Both types get a payoff of 0 when not in office. Putting things together, the politician's payoff is as follows:

$$u_P = \begin{cases} r + u_C + \delta x & \text{if the politician holds the office} \\ 0 & \text{otherwise.} \end{cases} \quad (2)$$

Let π be the probability that a randomly picked politician from the pool of candidates is unbiased. We assume that only the politician observes his type δ , whereas the citizen does not. The citizen will, however, observe the incumbent's policy choice and update her belief about the incumbent's type rationally.

The timing of the game is described in figure 9. There are five time periods denoted




Figure 9: Timeline of the Political Agency Model

v	the value of the internet
h	the internet harm
$\omega \in \{0, 1\}$	the state of the internet
$\lambda \in (0, 1)$	the prior probability that the internet content is harmful
$\delta \in \{0, 1\}$	the type of the politician
$\pi \in (0, 1)$	the prior probability that the politician is unbiased
$x \in \{0, 1\}$	whether to remove the internet content

Table 4: List of Mathematical Notations of the Model

by $t \in \{0, 1, 2, 3, 4\}$. At $t = 0$, nature reveals the state of the internet ω to either the incumbent or the citizen. At $t = 1$, the incumbent chooses the decision maker $i \in \{P, C\}$. At $t = 2$, the chosen decision maker i chooses whether to remove the content or not $x \in \{0, 1\}$. At $t = 3$, the citizen chooses whether to vote for the incumbent or the challenger. At $t = 4$, the elected politician chooses his preferred internet policy, and payoffs are realized.

In this game, the politician's pure strategy is a mapping $(i(\delta, \omega), x(i, \delta, \omega))$ that picks the decision maker and the removal choice if he is the decision maker for every realization of the pair (δ, ω) . The citizen's pure strategy is a tuple $(x, a)(i)$ that picks the removal choice and the voting decision for every history of the politician's behavior. Let π_C be the citizen's posterior belief on the probability that the incumbent politician is unbiased. The solution concept is perfect Bayesian equilibrium.

We normalize the bias term of the biased politician to be 1. We assume that $v \in (0, 1)$ such that bias overrides consumer welfare, $v < \lambda h$ such that the expected harm exceeds the value of the content, and $r > \lambda v$ such that the private return from holding office is not too small.

5.2 Analysis

In a dictatorship, the politician's re-election is unaffected by the citizen's voting. Whoever holds the office will simply choose his preferred internet policy, as there are no gains from doing anything else. For an unbiased politician $\delta = 0$, his behavior is determined by whether he has more information. If he observes the state ω , he will match the policy with the state. If the content is harmful $\omega = 1$, he would remove it $x = 1$; if instead the content is harmless $\omega = 0$, he would keep it $x = 0$. If he knows the citizen observes the state ω , he will choose $i = C$ and delegate the decision completely to the citizen. The citizen will then match the removal decision with the state $x = \omega$. This is the delegation principle of public regulation when the affected party is closer to the fact.

A biased politician $\delta = 1$, however, will choose internet control $x = 1$ regardless of the state x . Putting things together, when the politician is unconstrained by the electoral accountability, his strategy is

$$\begin{cases} i = P, x = \omega & \text{if } \delta = 0, s = P \\ i = C, x = \omega & \text{if } \delta = 0, s = C \\ i = P, x = 1 & \text{if } \delta = 1. \end{cases} \quad (3)$$

This implies a welfare loss for the citizen. Notice that efficiency requires $x^e = \omega$. The above policy is inefficient when $\delta = 1$ and $\omega = 0$. This is the censorship risk: *the biased politician removes the content when he knows the content is no harm to the citizen*. This happens because the politician faces no accountability constraints. The welfare loss of internet control in an authoritarian regime is the censorship risk. A censorship regime blocks the internet and sends huge amounts of removal requests to platforms.

In a democratic regime, the politician's re-election is determined by the voting. The first question might be: is the policy (3) still incentive compatible? It is not. To see this,

suppose $s = P$, and the citizen observes $x = 1$, her posterior belief π_C would be such that

$$\begin{aligned}\pi_C &= \frac{\Pr(x = 1|\delta = 0)\pi}{\Pr(x = 1|\delta = 0)\pi + \Pr(x = 1|\delta = 1)\pi} \\ &= \frac{\pi/2}{\pi/2 + (1 - \pi)} \\ &= \frac{\pi}{2 - \pi} < \pi\end{aligned}\tag{4}$$

The citizen believes the politician is more likely to be biased than an average politician from the challenger pool, and thus will not re-elect the incumbent. The biased politician will deviate from policy (3) because losing office is too high a cost.

In a democracy, when the citizen has more information on the content $s = C$, a pooling equilibrium might arise in which both types of politicians choose to delegate the internet policy to the citizen $i = C$. Since $i = C$ for both $\delta = 0$ and $\delta = 1$, the citizen learns nothing about the incumbent's type from the equilibrium policy choice and thus has no reason to change her prior belief. If the incumbent were to choose $i = P$, what would be the off-path belief of the citizen? Since a biased politician prefers $i = P$, it is reasonable to assume that the citizen would think the incumbent is more likely to be biased if she observes $i = P$. As a result, the citizen's posterior belief would be such that

$$\pi_C = \begin{cases} \pi & \text{if } i = C \\ \tilde{\pi} & \text{if } i = P, \end{cases}\tag{5}$$

where $\tilde{\pi} < \pi$. The citizen will re-elect the incumbent if she observes $i = C$, but will elect the challenger if she observes $i = P$.

For delegation to be an equilibrium, it remains to check whether the best responses of the citizen and both types of incumbents are consistent with the above belief profile. The citizen will match the removal decision with the state $x = \omega$ since she is perfectly informed. An unbiased politician faces no trade-off since he can match the state and still get re-elected - his payoff aligns with the citizen's payoff. If $\omega = 1$, the biased politician has no incentive to deviate since he prefers removal and the citizen will remove the content. Suppose instead $\omega = 0$ that the content is harmless, the biased politician would have a

trade-off between policy preference and re-election. If he chooses to remove the content himself $i = P, x = 1$, his payoff in the first term is $r + \delta$, but he will not get re-elected and will obtain a payoff of 0 in the second term. If he chooses to delegate the decision to the citizen $i = C$, his payoff in the first term is $r + v$, but he will get re-elected and obtain a payoff of $r + \delta$ in the second term. The biased incumbent would weigh up his first term rent from choosing $x = 1$ and being voted out of the office (which is $r + \delta$) against his future payoff of being re-elected (which is $2r + v + \delta$). The biased politician is willing to go against his preference and pretend to be a good type because reputation is important for re-election.

When the politician has more information on the content $s = P$, the same pooling equilibrium exists in which both types of politicians delegate to the citizen $i = C$ regardless of the state. Since $i = C$ for both $\delta = 0$ and $\delta = 1$, the citizen learns nothing about the incumbent's type from the equilibrium policy choice and thus has no reason to change his prior belief. Any incumbent picking $i = C$ will be re-elected by the citizen, and off-path belief implies that the incumbent will be voted out if he chooses other policies. The expected utility of the citizen if she decides to keep the content $v - \lambda h$. By the assumption $h > v/\lambda$, the citizen will remove the content if she is assigned to be the decision maker $i = C$. Anticipating this, the biased politician $\delta = 1$ would indeed choose $i = C$, because he prefers removal, and delegation also gets him re-elected. The same reason applies to the unbiased politician when $\omega = 1$. If he knows the content is harmful and the citizen would remove it anyway, he will assign the decision rights to the citizen. Suppose instead $\omega = 0$ that the content is harmless, the unbiased politician would have a trade-off between inefficiency and re-election. If he chooses to implement the efficient internet policy himself $i = P, x = 0$, his payoff in the first term is $r + v$, but he will not get re-elected and will obtain a payoff of 0 in the second term. If he chooses to delegate the decision to the citizen $i = C$, his payoff in the first term is just r , but he will get re-elected and obtain a payoff of $r + (1 - \lambda)v$ in the second term. By the assumption $r > \lambda v$, the present value of choosing $i = C$ is larger than that of $i = P, x = 0$ (i.e., $2r + (1 - \lambda)v > r + v$). To get re-elected, the unbiased politician finds it optimal to delegate the decision rights to the less informed citizen.

The equilibrium outcome under democracy is summarized as follows:

$$\begin{cases} i = C, x = \omega & \text{if } \delta = 0, 1 \text{ and } s = C \\ i = C, x = 1 & \text{if } \delta = 0, 1 \text{ and } s = P. \end{cases} \quad (6)$$

Regardless of which party has more information, the equilibrium outcome is the same: the politician delegates the right to the citizen and never directly remove the content. But the welfare properties of the delegation equilibrium are different depending on who observes the signal. If $s = C$, the equilibrium outcome is *efficient*. When the citizen has more information about the content, delegation is the welfare-maximizing policy. Re-election incentives motivate the incumbent to build reputation. Reputation-building imposes political discipline on the biased politicians such that he has to mimic the behavior of the unbiased. This happens because of electoral accountability: voters can remove the politician out of office if the politician has a bad reputation.

Reputation building can sometimes be counter-productive. If instead $s = P$, the equilibrium outcome is *inefficient*. The welfare loss of the uninformed delegation is over-removal: when $\omega = 0$ and $\delta = 0$, even though both the unbiased politician and the citizen would prefer to keep the good content $x = 0$, the content is removed $x = 1$. Both types of politicians choose to delegate the removal decision to the citizen $i = C$ irrespective of whether it is the best option for them or the citizen. This is the pandering risk: *the unbiased politician delegates the moderation decision to the citizens even when he knows the citizens are making sub-optimal removals*. This happens because the incumbent wants to preserve good reputation with the citizen and he ends up delegating too much power to the citizen.

If we compare the filtering outcome of autocracy and democracy in the model, we can reproduce the stylized facts presented in Section 4. Under autocracy, a content is removed unless the politician is unbiased and $\omega = 0$. The ex-ante probability of removing a content (or the fraction of content removed) is $1 - (1 - \lambda)\pi$. Under democracy, a content is removed if either the citizen is informed and $\omega = 1$ or the citizen is uninformed. The probability is $\frac{1}{2}(1 + \lambda)$. When $\lambda = \frac{\pi}{\pi+1/2}$, the two probabilities are equal. This leads to the following proposition.

Proposition 1. *There exists λ and π such that democracy is equally likely to remove content*

$x = 1$ as an autocracy.

Under autocracy, delegation happens only if the politician is unbiased and has more information. The probability of $i = C$ is $\frac{\pi}{2}$. Under democracy, delegation happens all the time with probability 1. This leads to the following proposition.

Proposition 2. *Compared to autocracy, democracy is more likely to delegate content removal to the citizen $i = C$.*

Under autocracy, internet control is inefficient because of censorship risk. It happens when the content is harmless but the politician is biased. The expected welfare loss is $(1 - \lambda)(1 - \pi)v$. Under democracy, internet control might be inefficient because of pandering risk. It happens when the content is harmless and the politician has more information. The expected welfare loss is $\frac{1-\lambda}{2}v$. We have the following proposition.

Proposition 3. *For $\pi \leq \frac{1}{2}$, the internet policy of democracy is more efficient than that of autocracy.*

The above proposition on social welfare comparison speaks to the consistently higher ranking of democracy in expert assessment surveys such as FOTN and V-Dem.

The key to understanding the relationship between democracy and internet control is the reputation concern of the incumbent politician. Higher electoral constraint reduces censorship $i = P, x = 1$.

Hypothesis 1. *Among democratic countries, higher accountability constraints lead to fewer government requests.*

Reputation building behavior is closely tied to the within-term election cycles. A stock idea in political economy is that politicians behave differently when elections loom. The agency model provides a natural way of thinking about such issues in general. Arguably the reputation concern is more salient when the election comes closer. In our model, the voting citizen's payoffs are different depending on whether an election is close or not. The expected payoff of voters in period one is ... while in period two it is ... In this

model, post-election government behavior (period two) is more ill-disciplined than that at the beginning of the term (period one) which seems to parallel the conventional wisdom on political cycles. We expect government control on the internet to be shown earlier in an election term while delegation is more likely when an election looms. Restraint on government requests is a signal to the voters.

Hypothesis 2. *Closer to election time, government takedown requests will decrease.*

6 The Effect of Electoral Accountability on Internet Control

The digital age has bestowed upon governments a powerful instrument in the form of information requests to internet platforms, which could potentially influence the public discourse, especially in the context of electoral processes. The timing of these requests in relation to electoral cycles may serve as a strategic tool for governments to control or sway public opinion, thereby impacting electoral outcomes. This study employs a two-way fixed effects regression model to critically analyze the causal relationship between the timing of government requests to internet giants such as Google and Twitter and impending elections.

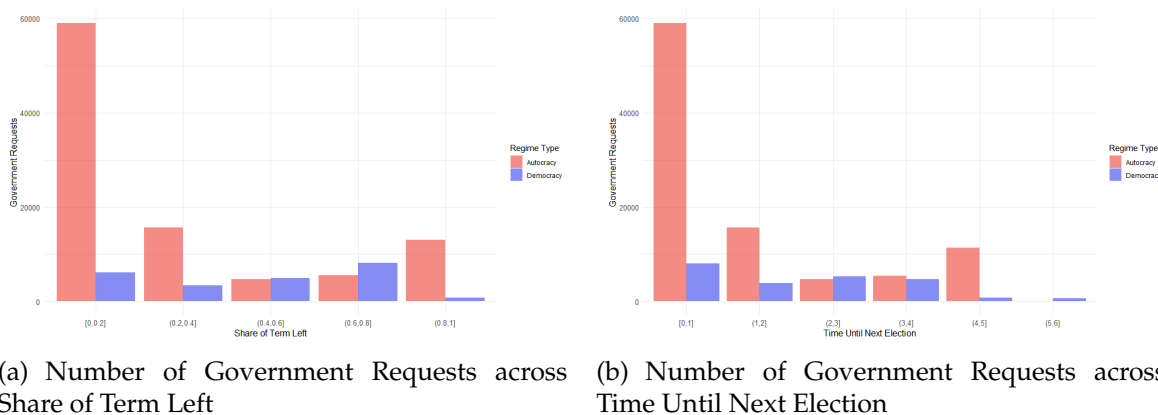


Figure 10: Distribution of Government Requests across Time

Figure 10(a) delineates the patterns of government requests in relation to the share of term left, with autocracies (red) and democracies (blue) depicted side by side. The

data reveals a stark contrast: autocratic regimes exhibit a consistent or even heightened level of requests throughout the term, suggesting that their actions are less influenced by electoral cycles. This contrasts with democratic regimes, where we see that the requests increase as share of term decreases, but there is a reduction during the election year. Such a trend in democracies could reflect an initial period of assertiveness post-election, possibly for setting the governance agenda, which then tapers off as the next election nears, underscoring a potential restraint imposed by impending electoral accountability. Similarly, Figure 10(b) contrasts the frequency of requests by governments against the years until the next election. The visual representation indicates that while democratic regimes show variable levels of requests as elections approach — likely due to the electoral check on power — autocracies demonstrate a lack of correlation with approaching elections, underlining the lesser electoral constraints on their governance behaviors. Autocracies may maintain or increase control measures independent of electoral timing, reflecting the less constrained nature of their governance. This comparison underscores the fundamental differences in how regime types engage with information control in the shadow of electoral mechanisms.

These visual insights necessitate a more rigorous econometric approach to disentangle the observed correlations from causations. To this end, we employ a two-way fixed effects regression model that allows us to control for both unobserved time-invariant country characteristics and common temporal shocks. By including fixed effects, we aim to mitigate the bias in causal inference that may arise from omitted variable concerns, such as inherent differences in government types or global trends in digital governance. To elucidate the potential causal links, we consider the frequency and nature of requests made by governments for user data or content removal in the run-up to elections. The two-way fixed effects model allows us to control for both unobservable time-invariant characteristics of individual countries and common shocks across all countries within specific time periods. The key variable of interest in our analysis is “Election Proximity,” which measures the time left until the next scheduled national election. The operationalization of “Election Proximity” is analogous to the “Share of Term Left” and “Time Until Next Election”. In our context, this variable gauges the temporal distance from an election, positing that the closer the period to the election, the higher the probability that government re-

quests to internet platforms will be strategically timed to impact voter behavior. The premise rests on the hypothesis that incumbent governments manipulate online information flows to shape electoral outcomes in their favor.

The regression equation is specified as follows:

$$\begin{aligned}
 Req_{c,t} = & \alpha + \beta_1 ShareOfTermLeft_{c,t} + \beta_2 TimeUntilNextElection_{c,t} \\
 & + \Theta Controls_{c,t} + \mu_c + \tau_t + \epsilon_{c,t}
 \end{aligned}
 \tag{7}$$

where $Req_{c,t}$ represents the count of either government requests or court orders by government 'c' in time 't', $ShareOfTermLeft_{c,t}$ is the normalized score denoting the share of term left for the current government, and $TimeUntilNextElection_{c,t}$ represents the time until the next election. The vector $Controls_{c,t}$ includes control variables such as economic indicators, political stability measures, and internet penetration rates. The term μ_c represents country-fixed effects, τ_t denotes year-fixed effects, and $\epsilon_{c,t}$ is the error term.

6.1 Operationalizing the Independent Variable

In our empirical analysis, we operationalize two critical independent variables to capture the political context: the Share of Term Left and the Time Until Next Election. These variables are pivotal in exploring how the temporal proximity to elections influences governmental behavior, particularly in the realm of internet governance. The CCP's dataset, specifically the "Characteristics of National Constitutions, Version 4.0," offers detailed insights into the constitutional frameworks governing term limits and electoral cycles in various countries. This information is crucial for our construction of the independent variables 'Share of Term Left' and 'Time Until Next Election'. By integrating the CCP's data on term limits (referred to as 'hosterm' in our study) with our datasets, we can accurately compute the political timelines relevant to each nation's governance structure. This integration allows us to normalize the years in office of political leaders against their respective constitutional term limits.

Share of Term Left The variable ‘Share of Term Left’ is conceptualized to reflect the remaining tenure of a political leader within a given term. This measure is crucial for understanding the potential impact of nearing the end of a political term on a government’s decision-making process, especially concerning strategies for internet control. In operational terms, for each country-year observation, we calculate this variable by normalizing the years in office against the host country’s constitutional term limit. This normalization accounts for scenarios involving consecutive terms, ensuring a cyclical reflection of a leader’s tenure within the constitutional bounds of their office. We then transform this variable to represent the proportion of the term that remains, under the hypothesis that the closer a leader is to the end of their term, the more their actions, including internet control decisions, might be influenced by considerations of political survival and public perception.

Time Until Next Election The ‘Time Until Next Election’ variable quantifies the years remaining until the next scheduled national election. This variable is instrumental in assessing how the imminence of elections may sway governmental actions. It offers a temporal lens to examine the strategic considerations of governments in relation to internet censorship or freedom, hypothesizing that impending elections could lead to shifts in governance strategies to favor public opinion or consolidate power.

6.2 Natural Experiment Setup

The relationship between government requests for internet content removal and electoral cycles offers a unique lens to understand the exogeneity of elections in the context of political maneuvering. Typically, elections are scheduled events, following a predictable cycle that is determined well in advance. This predictability provides a natural experiment setting, where the timing of elections – ordinarily exogenous to a government’s day-to-day decision-making – can be correlated with variations in government requests for content removal. This setup allows us to isolate the impact of impending elections on government behavior, particularly in terms of internet censorship and control.

However, the occurrence of snap elections adds a layer of complexity to this analysis. Unlike regular elections, snap elections can be called at short notice and are often a strategic tool employed by incumbent governments to capitalize on favorable political conditions or to catch opposition parties off-guard. This introduces an element of endogeneity, as the decision to call a snap election may be influenced by the same underlying factors that affect a government’s propensity to control information. Nevertheless, our natural experiment framework remains robust for several reasons. First, snap elections, while strategically timed, are still relatively rare and follow legal and constitutional procedures, limiting the extent to which they can be manipulated purely for information control purposes. Second, the majority of elections in our dataset are regular, scheduled elections, thereby reducing the overall impact of snap elections on our analysis.

In this context, our study navigates through these complexities by focusing predominantly on regular electoral cycles, while acknowledging the nuances introduced by snap elections. By examining patterns of government requests in relation to both scheduled and unscheduled elections, we can better understand the dynamics at play between political cycles and internet censorship. Our approach takes into account the possibility of strategic behavior by governments while still leveraging the predominantly exogenous nature of election timing to draw insights about the relationship between political cycles and government censorship behavior on the internet.

6.3 Regression Results

Table 5 examines the effect of a political leader’s remaining term share on the number of government requests sent to Google. The baseline model in column (I) shows the effect of an election coming close and the effect of the regime type. As share of term left decreases, the number of government requests increases. This is consistent with the behavior of a rational, self-interested politician who tries to remove any adverse news before the election. Compared to autocracy, democracy reduces the government requests substantially. This is consistent with the stylized facts we present in Section 4. Democracy imposes political discipline on the politician’s internet strategy through electoral accountability. This is shown by the significant positive effect of the share of term left on the govern-

ment requests in democracy. As Column (II) shows, compared to autocracy, a one percent decrease in the share of term left in democratic countries led to a decrease of 972 requests. The estimated effect is robust to the inclusion of control variables, the state-fixed effect, and the year-fixed effect. The estimates in column (5) tell us that holding all other things constant, forthcoming re-election in democracy, in contrast to autocracy, led to 1,368 fewer requests. Collectively, the results in Table 5 suggest that there are significant effects of election timing on internet takedowns, and importantly the reputation effect we conjecture is only salient in democratic countries.

	<i>Number of Government Requests</i>				
	(1)	(2)	(3)	(4)	(5)
Share of Term Left	-339.137*	-960.898***	-1,127.193***	-1,515.886***	-1,459.937***
	(173.898)	(288.516)	(314.651)	(368.019)	(368.617)
Democracy	-338.783***	-827.482***	-963.908***	-876.315**	-795.679*
	(107.625)	(210.665)	(237.927)	(430.001)	(431.682)
Share of Term Left * Democracy		972.343***	1,156.462***	1,521.103***	1,482.328***
		(360.801)	(391.385)	(435.482)	(437.625)
GDP Per Capita			-76.213	-864.249**	-917.237**
			(93.824)	(396.097)	(437.011)
% of Internet Users			6.028	7.734	5.453
			(4.120)	(7.506)	(9.976)
Urban Population %			1.851	5.039	-14.071
			(4.248)	(68.418)	(71.621)
State Fixed Effect				x	x
Year Fixed Effect					x
Observations	829	829	774	774	774

Note: Standard errors in parentheses. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

Table 5: Number of Government Requests with Share of Term Left

Table 6 provides insights into how the proximity of elections and regime type influence government requests for content removal. In model (5), the coefficient for Time Until Next Election is -281.434, indicating a notable increase in government requests as elections draw nearer. This suggests that governments, anticipating elections, may intensify efforts to control online content. The role of regime type is highlighted in the Democracy variable, where transitioning from an autocracy to a democracy leads to a significant decrease in requests, with model (5) showing a reduction of approximately 848 requests. This substantial decline underscores the restraining influence of democratic governance on government censorship efforts. The interaction term between Time Until Next Elec-

	<i>Number of Government Requests</i>				
	(1)	(2)	(3)	(4)	(5)
Time Until Next Election	-22.921 (20.428)	-70.516*** (27.285)	-85.325*** (29.668)	-291.537*** (65.629)	-281.434*** (65.741)
Democracy		-578.896*** (166.759)	-685.433*** (191.208)	-938.742** (436.546)	-848.326* (438.703)
Time Until Next Election * Democracy		71.013* (42.846)	89.781* (46.625)	278.209*** (77.801)	268.342*** (78.283)
GDP Per Capita			-84.508 (95.307)	-875.858** (395.197)	-920.629** (436.140)
% of Internet Users			6.586 (4.132)	8.200 (7.485)	6.381 (9.947)
Urban Population %			2.133 (4.340)	7.104 (68.285)	-10.486 (71.503)
State Fixed Effect				x	x
Year Fixed Effect					x
Observations	829	829	774	774	774

Note: Standard errors in parentheses. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

Table 6: Government Requests with Time Until Next Election

tion and Democracy in model (5) shows an increase of around 268 requests. This suggests that in democracies, the decrease in government requests as elections approach is more pronounced than in autocracies, reinforcing the notion of democratic accountability as a deterrent to government censorship. The results in model (5) are robust to the inclusion of GDP per capita, the percentage of internet users, urban population percentage, and fixed effects for state and year. Overall, Table 6 reinforces the idea that electoral timing and regime type are crucial determinants of government behavior in internet content management, particularly in democratic contexts where electoral accountability plays a significant role.

The analysis of Table 7 reveals a crucial aspect of the judicial process in internet control, particularly in contrast to direct government requests. The primary variables of interest, 'Share of Term Left' and 'Democracy,' demonstrate non-significant coefficients across all models, most notably in model (5), which incorporates the most controls. This non-significance is not just a statistical observation but a substantive finding that highlights the unique nature of court orders. Unlike government requests, which are often influenced by political factors and electoral cycles, court orders appear to be less sus-

	<i>Number of Court Orders</i>				
	(1)	(2)	(3)	(4)	(5)
Share of Term Left	-31.649 (24.406)	-42.337 (41.472)	-26.650 (33.199)	-8.979 (22.101)	-10.573 (21.235)
Democracy	-0.378 (15.247)	-8.508 (29.711)	42.088* (24.475)	-29.182 (26.084)	-28.123 (25.040)
Share of Term Left * Democracy		16.361 (51.310)	2.866 (40.865)	10.108 (26.134)	19.077 (25.144)
GDP Per Capita			-10.712 (9.751)	-44.718* (23.482)	-78.857*** (24.962)
% of Internet Users			0.066 (0.446)	2.574*** (0.456)	1.249** (0.586)
Urban Population %			1.538*** (0.442)	-0.048 (4.071)	-4.274 (4.156)
State Fixed Effect				x	x
Year Fixed Effect					x
Observations	848	848	791	791	791

Note: Standard errors in parentheses. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

Table 7: Court Orders with Share of Term Left

ceptible to such considerations. The lack of a significant relationship between the key variables and the number of court orders suggests that judicial decisions related to internet control are more insulated from immediate political pressures. This aligns with the theoretical understanding that courts, unlike politicians, do not operate under the same temporal constraints and reputation concerns associated with electoral cycles. Furthermore, the significant coefficients observed in control variables, such as GDP Per Capita and the percentage of Internet Users, underscore that socioeconomic factors might have a more pronounced role in the context of court orders. This could imply that judicial decisions on internet control are more reflective of the broader economic and technological environment rather than immediate political motivations. In essence, Table 7 substantiates the hypothesis that court orders, as a mechanism of controlling internet content, operate on a different paradigm than government requests. They are less entwined with the fluctuations of political cycles and more grounded in broader, systemic factors, making them a distinct entity in the landscape of internet governance.

Table 8 extends our understanding of the dynamics between electoral cycles and internet control, specifically through the lens of court orders. The analysis here aligns with

	<i>Number of Court Orders</i>				
	(1)	(2)	(3)	(4)	(5)
Time Until Next Election	-7.889*** (3.003)	-7.687* (4.052)	-3.679 (3.263)	-2.157 (3.945)	-2.043 (3.789)
Democracy	-8.041 (15.558)	-6.730 (23.537)	47.311** (19.761)	-31.397 (26.450)	-29.844 (25.403)
Time Until Next Election * Democracy		-0.449 (6.040)	-3.521 (4.842)	2.517 (4.679)	3.840 (4.505)
GDP Per Capita			-8.031 (9.833)	-44.758* (23.473)	-78.875*** (24.960)
% of Internet Users			0.049 (0.444)	2.572*** (0.455)	1.254** (0.585)
Urban Population %			1.442*** (0.450)	-0.019 (4.071)	-4.243 (4.156)
State Fixed Effect				x	x
Year Fixed Effect					x
Observations	848	848	791	791	791

Note: Standard errors in parentheses. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

Table 8: Court Orders with Time Until Next Election

the patterns observed in Table 7, underscoring the non-significant influence of election timing on the issuance of court orders. In models (1) and (2), while we observe a statistically significant coefficient for 'Time Until Next Election,' the effect size is minimal, indicating a negligible impact on the number of court orders. This is in stark contrast to the findings for government requests, where proximity to elections showed a substantial effect. This difference highlights a key aspect of judicial decision-making in internet governance: court orders are less subject to the ebb and flow of political cycles and electoral pressures, reflecting a more constant, non-partisan approach. The interaction term 'Time Until Next Election * Democracy' also yields non-significant results across all models, reinforcing the notion that, unlike direct government actions, judicial decisions on internet content removal are less influenced by the political climate or regime type. This further illustrates the distinct roles and motivations governing the actions of political entities and judicial bodies in the realm of internet control. In conclusion, Table 8 consolidates our earlier proposition that court orders represent a different facet of internet governance, one that is less malleable to political and electoral influences. This distinction is crucial for understanding the multifaceted nature of internet control mechanisms across different governance structures.

6.4 Discussion of Results

Our empirical analysis, built upon the foundation of our political agency model, uncovers intricate relationships between political regimes, electoral cycles, and internet control strategies. To elucidate these connections, we draw upon the model’s key parameters, including the probability π denoting the likelihood of encountering an unbiased politician, and the binary variables δ and ω representing the politician’s type and the state of internet content, respectively. These parameters serve as a theoretical lens through which we interpret our empirical findings, providing valuable insights into the dynamics of internet governance.

In democratic contexts, where the probability π of encountering an unbiased politician increases, our results reveal a corresponding reduction in government requests for internet content removal, particularly as elections draw near (see Table 5). This empirical observation aligns with the model’s depiction of democratic leaders, who, when under heightened public scrutiny ($\omega = 0$) and seeking re-election, tend to refrain from overt censorship ($x = 0$). This behavior mirrors the model’s assertion that democratic governments, while considering the electoral implications and potential reputation costs (r), tend to adopt less intrusive internet policies.

Furthermore, the model posits that democracies tend to delegate content removal decisions to users, especially in the presence of high electoral accountability (see Table 5). Our data substantiates this hypothesis, indicating a discernible shift from state-led control to a more user-centric approach ($i = C$) as democratic governments approach elections. This shift resonates with the model’s portrayal of politicians carefully weighing the advantages of policy efficiency against the risks of electoral backlash in democratic settings.

Moreover, our empirical findings bolster the model’s proposition that democratic internet policies may exhibit greater efficiency compared to autocracies, particularly when politicians deliberate over the trade-offs between policy efficiency and potential electoral consequences (see Table 5). The observed decrease in government requests for content removal as elections approach suggests a pivot toward more efficient and less intrusive

internet policies in democracies, potentially mitigating the perils of over-censorship often associated with authoritarian regimes ($\delta = 1$).

In contrast, the data regarding user decisions (delegation in our model) reveal a different narrative, characterized by a lack of significant correlation with electoral cycles and regime type (see Table 7). This finding supports the model's theoretical distinction between the motivations and mechanisms of political entities and user-based decisions (delegation) in the realm of internet governance. Unlike direct government actions, which are influenced by political and electoral dynamics ($i = P$), user-based decisions concerning internet content removal ($i = C$) appear to be guided by more stable, non-partisan considerations.

In summary, our study offers an in-depth exploration of internet governance mechanisms across diverse political structures. We have substantiated our analysis with references to specific table numbers, linking theoretical insights from our model to empirical observations. This approach sheds light on the intricate interplay between political regimes, electoral accountability, and strategies of internet control. It provides profound insights into the multifaceted nature of internet governance and the pivotal role played by democratic norms and electoral processes in shaping these mechanisms.

7 Conclusion

Our study offers a comprehensive view of how political regimes, particularly democracies, interact with the digital sphere, especially in the context of electoral cycles. The findings challenge the conventional narrative that democracies are inherently liberal in terms of internet governance. Instead, we uncover a more complex reality where democratic governments, much like their authoritarian counterparts, engage in strategic timing of internet control activities, albeit with differing methodologies and implications.

Our model, as detailed in Section 5, provides a theoretical framework that enhances our understanding of the empirical findings regarding internet control and electoral cycles. In democracies, the behavior of politicians, as influenced by their reputation con-

cerns, plays a pivotal role in shaping internet policy decisions. The model suggests that both types of politicians, biased ($\delta = 1$) and unbiased ($\delta = 0$), opt to delegate content removal decisions to citizens ($i = C$) when they possess greater information on the content ($s = P$). This strategy ensures re-election as it aligns with voter preferences, avoiding the risks associated with direct censorship. The model's implications are consistent with the findings from our empirical analysis. In democratic regimes, the significant decrease in government requests for content removal as elections near, as shown in Table 5, can be attributed to politicians' desire to maintain a positive reputation with voters. By delegating the decision-making power to citizens, politicians in democracies can avoid the direct responsibility for content removal, thus navigating the delicate balance between maintaining public order and upholding the values of internet freedom.

Furthermore, the model explains why court orders do not exhibit the same pattern as direct government requests (Table 7 and Table 8). Court orders, being less directly associated with the incumbent politicians, do not significantly impact their reputation in the same way. This distinction is crucial as it underscores the multifaceted nature of internet governance in democracies, where various branches of government play different roles. In autocratic regimes, however, the model predicts a higher likelihood of direct internet control ($i = P, x = 1$) due to lower reputational concerns and electoral constraints. This is consistent with the observation that autocracies exhibit a more constant pattern of internet control, irrespective of electoral cycles.

Our research has significant implications for the understanding of digital governance in the modern world. It calls for a nuanced approach to analyzing internet freedom and control, moving beyond the simplistic dichotomy of democratic liberalism versus authoritarian censorship. The findings suggest that democratic governments are not immune to the temptations of controlling online discourse, especially when faced with the pressures of impending elections. This insight is crucial for policymakers, internet platforms, and civil society in advocating for more transparent and accountable internet governance practices. Looking ahead, this study opens several avenues for future research. One key area is examining how different electoral systems within democracies might influence government behavior in digital governance. Another important aspect is exploring the role of international norms and agreements in shaping or constraining

government actions in internet control. Additionally, understanding the impact of these governmental strategies on public opinion, democratic participation, and overall trust in government institutions could provide valuable insights into the broader consequences of digital governance strategies.

In conclusion, our study provides a critical contribution to the discourse on internet governance, highlighting the intricate ways in which electoral cycles and political structures shape online content control strategies. As the digital realm becomes increasingly central to political, social, and economic life, understanding these dynamics is essential for fostering an internet environment that is both free and respectful of democratic principles.

References

- Akdeniz, Yaman and Kerem Altiparmak**, *Internet: restricted access: a critical assessment of Internet content regulation and censorship in Turkey*, Imaj Kitabevi & Imaj Yayinevi, 2008.
- Akgül, Mustafa and Melih Kırılıdoğ**, "Internet censorship in Turkey," *Internet Policy Review*, 2015, 4 (2), 1–22.
- Badawy, Adam, Emilio Ferrara, and Kristina Lerman**, "Analyzing the digital traces of political manipulation: The 2016 Russian interference Twitter campaign," in "2018 IEEE/ACM international conference on advances in social networks analysis and mining (ASONAM)" IEEE 2018, pp. 258–265.
- Bambauer, Derek E**, "Filtering in Oz: Australia's foray into Internet censorship," *U. Pa. J. Int'l L.*, 2009, 31, 493.
- Barquin, Sonia, Guillaume de Gantès, HV Vinayak, and Duhita Shrikhande**, "Digital banking in Indonesia: Building loyalty and generating growth," *McKinsey & Company*, February, 2019, 6.
- Besley, Timothy**, *Principled agents?: The political economy of good government*, Oxford University Press on Demand, 2006.
- Breindl, Yana and Bjoern Kuellmer**, "Internet content regulation in France and Germany: Regulatory paths, actor constellations, and policies," *Journal of Information Technology & Politics*, 2013, 10 (4), 369–388.
- Budnitsky, Stanislav and Lianrui Jia**, "Branding Internet sovereignty: Digital media and the Chinese–Russian cyberalliance," *European Journal of Cultural Studies*, 2018, 21 (5), 594–613.
- Chung, Jongpil**, "Comparing online activities in China and South Korea: The internet and the political regime," *Asian Survey*, 2008, 48 (5), 727–751.
- Dick, Archie L, Lilian I Oyieke, and Theo JD Bothma**, "Are established democracies less vulnerable to Internet censorship than authoritarian regimes? The social media test," *FAIFE spotlight*, 2012.

- Egorov, Georgy, Sergei Guriev, and Konstantin Sonin**, "Why resource-poor dictators allow freer media: A theory and evidence from panel data," *American political science Review*, 2009, 103 (4), 645–668.
- Esarey, Ashley and Qiang Xiao**, "Digital communication and political change in China," *International Journal of Communication*, 2011, 5, 22.
- Faris, Robert and Nart Villeneuve**, "Measuring global Internet filtering," *Access denied: The practice and policy of global Internet filtering*, 2008, 5.
- Fish, Eric**, "Is Internet censorship compatible with democracy? Legal restrictions of on-line speech in South Korea," *Asia-Pacific Journal on Human Rights and the Law*, 2009, 10 (2), 43–96.
- Foucault, Michel**, *Security, territory, population: lectures at the Collège de France, 1977-78*, Springer, 2007.
- Frydman, Benoît and Isabelle Rorive**, "Regulating Internet content through intermediaries in Europe and the USA," *Zeitschrift für Rechtssoziologie*, 2002, 23 (1), 41–60.
- Gaubatz, Kurt Taylor**, "Election cycles and war," *Journal of Conflict Resolution*, 1991, 35 (2), 212–244.
- Gehlbach, Scott and Konstantin Sonin**, "Government control of the media," *Journal of public Economics*, 2014, 118, 163–171.
- Goldsmith, Jack**, "Who controls the Internet? Illusions of a borderless world," *Strategic Direction*, 2007, 23 (11).
- Hellmeier, Sebastian**, "The dictator's digital toolkit: Explaining variation in internet filtering in authoritarian regimes," *Politics & Policy*, 2016, 44 (6), 1158–1191.
- Hess, Gregory D and Athanasios Orphanides**, "War politics: An economic, rational-voter framework," *The American Economic Review*, 1995, pp. 828–846.
- Hobbs, William R and Margaret E Roberts**, "How sudden censorship can increase access to information," *American Political Science Review*, 2018, 112 (3), 621–636.

- King, Gary, Jennifer Pan, and Margaret E Roberts**, "How censorship in China allows government criticism but silences collective expression," *American political science Review*, 2013, 107 (2), 326–343.
- , – , and – , "Reverse-engineering censorship in China: Randomized experimentation and participant observation," *Science*, 2014, 345 (6199), 1251722.
- Kolozaridi, Polina and Dmitry Muravyov**, "Contextualizing sovereignty: A critical review of competing explanations of the Internet governance in the (so-called) Russian case," *First Monday*, 2021.
- Lian, Bradley and John R Oneal**, "Presidents, the use of military force, and public opinion," *Journal of Conflict Resolution*, 1993, 37 (2), 277–300.
- Lorentzen, Peter**, "China's strategic censorship," *American Journal of political science*, 2014, 58 (2), 402–414.
- Maskin, Eric and Jean Tirole**, "The politician and the judge: Accountability in government," *American Economic Review*, 2004, 94 (4), 1034–1054.
- Mchangama, Jacob and Joelle Fiss**, "The digital Berlin Wall: How Germany (accidentally) created a prototype for global online censorship," *Copenhagen: Justitia and Authors*, 2019.
- Mueller, Milton L**, *Networks and states: The global politics of Internet governance*, MIT press, 2010.
- Pearce, Paul, Roya Ensafi, Frank Li, Nick Feamster, and Vern Paxson**, "Augur: Internet-wide detection of connectivity disruptions," in "2017 IEEE Symposium on Security and Privacy (SP)" IEEE 2017, pp. 427–443.
- Petrov, Nikolay, Maria Lipman, and Henry E Hale**, "Three dilemmas of hybrid regime governance: Russia from Putin to Putin," *Post-Soviet Affairs*, 2014, 30 (1), 1–26.
- Prat, Andrea and David Strömberg**, "The political economy of mass media," *Advances in economics and econometrics*, 2013, 2, 135.

- Qiang, Xiao**, "Liberation technology: the battle for the Chinese internet," *Journal of Democracy*, 2011, 22 (2), 47–61.
- Rao, Weijia**, "Domestic Politics and Settlement in Investor-State Arbitration," *The Journal of Legal Studies*, 2021, 50 (1), 145–185.
- Roberts, Margaret E.**, *Censored: Distraction and Diversion Inside China's Great Firewall*, Princeton University Press, 2018.
- Schedler, Andreas**, "Authoritarianism's last line of defense," *J. Democracy*, 2010, 21, 69.
- Shadmehr, Mehdi and Dan Bernhardt**, "State censorship," *American Economic Journal: Microeconomics*, 2015, 7 (2), 280–307.
- Simonov, Andrey and Justin Rao**, "Demand for online news under government control: Evidence from russia," *Journal of Political Economy*, 2022, 130 (2), 259–309.
- Sivetc, Liudmila**, "Controlling free expression "by infrastructure" in the Russian Internet: The consequences of RuNet sovereignization," *First Monday*, 2021.
- Tréguer, Félix**, "From deep state illegality to law of the land: The case of internet surveillance in France," in "7th Biennial Surveillance & Society Conference (SSN 2016):" Power, performance and trust"" 2016.
- Urman, Aleksandra and Mykola Makhortykh**, "How transparent are transparency reports? Comparative analysis of transparency reporting across online platforms," *Telecommunications Policy*, 2023, p. 102477.
- Warf, Barney**, "Geographies of global Internet censorship," *GeoJournal*, 2011, 76, 1–23.
- Williams, Laron K**, "Flexible election timing and international conflict," *International Studies Quarterly*, 2013, 57 (3), 449–461.
- Wulf, Volker, Dave Randall, Konstantin Aal, and Markus Rohde**, "The Personal is the Political: Internet Filtering and Counter Appropriation in the Islamic Republic of Iran," *Computer Supported Cooperative Work (CSCW)*, 2022, 31 (2), 373–409.
- Yesil, Bilge and Efe Kerem Sozeri**, "Online surveillance in Turkey: Legislation, technology and citizen involvement," *Surveillance & Society*, 2017, 15 (3/4), 543–549.

Zarras, Apostolis, “Leveraging Internet services to evade censorship,” in “International Conference on Information Security” Springer 2016, pp. 253–270.

Zittrain, Jonathan L, Robert Faris, Helmi Noman, Justin Clark, Casey Tilton, and Ryan Morrison-Westphal, “The shifting landscape of global internet censorship,” *Berkman Klein Center Research Publication*, 2017, (2017-4), 17–38.