# Governments and the Net:

## Defense, Control, and Trust in the Fifth Domain

Dissertation submitted for the degree of
Doctor of Social Sciences (Dr. rer. soc.)

presented by

**Lukas Kawerau**

at


Sektion Politik - Recht - Wirtschaft
Fachbereich Politik & Verwaltungswissenschaft

Konstanz, 2021

| | |
|---|---|
| Date of Defense: | November 26, 2021 |
| First Referee: | Prof. Dr. Nils B. Weidmann |
| Second Referee: | Prof. Dr. Gerald Schneider |
| Third Referee: | Prof. Dr. Anita Gohdes |

Für Mama, ich wünschte du könntest das hier lesen.
Für Josy, ohne dich wäre es nicht gegangen.

## Acknowledgments

This dissertation would not have come to be without many strong hands lending their support.

Foremost, I would like to thank my main supervisor Nils B. Weidmann, the best supervisor one could wish for. Despite many challenges, setbacks, and delays somehow Nils never gave up on me. For close to a decade now Nils has inspired me to do research and to do it well. I am very grateful for his support, his always honest feedback, and his mentorship over the long years we have worked together.

I would also like to thank my other supervisors Gerald Schneider and Anita Gohdes. The wisdom Gerald Schneider imparted in very memorable lectures during the Master's degree was instrumental in keeping the right mindset. Anita provided very valuable feedback, and reading her work provided consistent inspiration. In a similar vein I would like to thank Margaret (Molly) E. Roberts, Alberto Dainotti, and Alistair King. The collaboration with them was foundational for my work in this dissertation and made the analyses possible in the first place.

In addition, I am very grateful for the help and encouragement of my colleagues and friends who provided invaluable feedback on drafts, helped through coffee-fueled brainstorming sessions, and endless exchanges of hilarious memes. Sebastian, Philipp, Eda, Max H., and Sascha – lunches and coffee with you kept me going. Alina, Max, and Steffen – knowing you had my back from the Master through today has helped in ways you'll never know.

I would also be remiss to not mention some other guides and mentors who, sometimes entirely unwittingly through their mere presence online, helped me reach this point. Susanne helped me finish the rabbit ears. Mark and Michael showed me the standard and that the mind is primary. Zed taught me programming, which was literally the first step on this particular journey. Conor built the right tool for how my mind works. To all of you – thank you.

Finally, I would like to thank my family. My parents always believed in me, no matter what was going on. I'll be forever grateful for your love and support, Mama und Papa. I am blessed with the most supportive sister one could wish for, and a mother-in-law always willing to lighten the load. Danke, Laura und Ingeborg. Most importantly, thank you, Josy. I have no words for how grateful I am for your unconditional support, your patience, and your belief in my ability to do this. Thank you, forever.

## Zusammenfassung

In den letzten dreißig Jahren haben sich das Internet generell und das World Wide Web im besonderen von einer Nischentechnologie zum Hauptschauplatz politischen Geschehens entwickelt. Dieser Wandel ist auch an Regierungen nicht spurlos vorbeigegangen. In meiner Dissertation untersuche ich drei Aspekte der Nutzung des Internets durch Regierungen und trage außerdem zur methodischen Weiterentwicklung der Politikwissenschaft bei.

Bisherige Forschung über den Einfluss des Internets auf Politik hat sich vorwiegend darauf konzentriert den politischen Diskurs und die Interaktion zwischen Bürgern und Regierungen zu analysieren. Die technische Realität des Internets bedeutet für Regierungen allerdings auch neue Herausforderungen und Möglichkeiten, die darüber hinaus gehen. In meiner Dissertation beleuchte ich daher Fragen der Cybersicherheit, Zensur und Informationskontrolle sowie des Umgangs mit der inhärent internationalen technischen Struktur des Internets durch Regierungen. Methodisch erweitere ich in diesem Zuge das Repertoire der Politikwissenschaft um Techniken der Internet-Messung aus der Informatik, die empirische Analysen dieser Fragen erst möglich machen.

Im ersten Papier führe ich diese Techniken und ihre Bedeutung für die Politikwissenschaft grundsätzlich ein und wende sie dann auf einen Ländervergleich defensiver Cybersicherheit an. Unter Verwendung von Internet-Messdaten zu Sicherheitsschwachstellen auf Servern, die Websites von Regierungen hosten, konstruiere ich einen neuen, beobachtenden Indikator für defensive Cybersicherheitsfähigkeiten und vergleiche ihn mit einem Indikator, der auf Expertenbefragungen basiert. Meine Analyse zeigt, dass der Beobachtungsindikator plausibel dasselbe Konzept misst wie der auf Expertenbefragungen basierende Indikator und dass Expertenbefragungen durch die Medienberichterstattung über Sicherheitsverletzungen in einer Weise verzerrt sein könnten, wie es bei Beobachtungsindikatoren nicht der Fall ist.

Im zweiten Papier, gemeinsam verfasst mit Nils B. Weidmann und Alberto Dainotti, untersuchen wir ob und wie Autokratien zwischen Taktiken der Onlinezensur wählen. Dazu analysieren wir zwei Zensurtaktiken, Website-Blockierung und Denial-of-Service-Angriffe. Für unsere empirische Analyse stützen wir uns dazu auf Internet-Messdaten, um eine neue Messung beizusteuern, mit der sich Denial-of-Service-Angriffe besser auf mögliche Ziele abbilden lassen. Die Ergebnisse unserer Analyse liefern erste Hinweise darauf, dass Autokraten je nach aktueller Situation Taktiken aus ihrem Zensurrepertoire auswählen. In Wochen mit Protesten ist die Beobachtung des Vorhandenseins von Webseiten-Sperrungen mit *weniger* DoS-Angriffen gegen oppositionelle Webseiten verbunden, während sie in Wochen ohne Proteste mit *mehr* DoS-Angriffen korreliert. Dies bestätigt unsere theoretische Erwartung, dass Autokraten zwischen taktischer Verstärkung und taktischer Substitution wählen, wenn sie entscheiden, wie sie die Taktiken aus ihrem Repertoire einsetzen.

Im dritten Papier hinterfrage ich die Beobachtung, dass viele Regierungen ihre offiziellen Webseiten und digitalen Dienste durch Unternehmen online bringen die im Ausland angesiedelt sind und sich damit außerhalb der Kontrolle der Regierungen befinden. Diese Beobachtung widerspricht Annahmen aus der theoretischen Bedeutung von Lieferkettensicherheit und nationaler Datensouveränität. Gegeben der Entscheidung offizielle Regierungsseiten durch ausländische Unternehmen online zu stellen, frage ich welche Faktoren die Wahl beeinflussen könnten in welchem Land eine Regierungen eigene Webseiten ansiedelt. Ich untersuche diese Frage empirisch indem ich Internet-Messdaten zu den Hosting-Providern von Regierungsseiten verwende und inter-staatliches Vertrauen

durch gemeinsame Bündnismitgliedschaft und relativen Demokratiestatus modelliere. Die Ergebnisse dieser Analyse zeigen, dass Regierungen ihre offiziellen Webseiten eher in Ländern ansiedeln, denen sie vertrauen.

Zusammengefasst unterstreicht die vorliegende Dissertation die Notwendigkeit, die Nutzung des Internets durch Regierungen nicht nur im Bezug auf politische Inhalte zu analysieren, sondern auch die tieferen technischen Aspekte dieser Nutzung zu beleuchten. Weiterhin zeige ich wie Politikwissenschaftler ihre methodisches Repertoire um Techniken der Internet-Messung erweitern können um solche Analysen durchzuführen.

## Abstract

Over the last thirty years, the Internet generally and the World Wide Web in particular have gone from niche technology to the dominant venue for political interaction. This development has also affected governments. In my dissertation I investigate three aspects of how governments use the Internet and additionally contribute to the methodological development in political science.

Previous research on the impact of the Internet on politics has focused primarily on analyzing political discourse and interaction between citizens and governments. However, the technical reality of the Internet also presents new challenges and opportunities to governments that go beyond this. In my dissertation, I therefore examine issues of cybersecurity, censorship, and information control, as well as how governments deal with the inherently international technical structure of the Internet. Methodologically, I expand the repertoire of political science in this course with techniques of Internet measurement from computer science, which make empirical analyses of these questions possible in the first place.

In the first paper, I introduce these techniques and their relevance to political science, and then apply them to a cross-country comparison of defensive cybersecurity. Using Internet measurement data of security vulnerabilities found on servers that host government websites, I construct a new, observational indicator for defensive cybersecurity capability and compare it to an indicator based on expert interviews. My analysis shows that the observational indicator plausibly measures the same concept as the indicator based on expert surveys and that expert surveys might be biased by media coverage of security breaches in a way observational indicators are not.

In the second paper, co-authored with Nils B. Weidmann and Alberto Dainotti, we examine whether and how autocracies choose between online censorship tactics. We analyze two censorship tactics, website blocking and denial-of-service attacks. For our empirical analysis, we rely on Internet measurement data to contribute a new measurement to better map Denial-of-Service attacks to possible targets. The results of our analysis provide first evidence that autocrats select tactics from their censorship repertoire depending on the current situation. In weeks with protest, observing the presence of website blocking is associated with *fewer* DoS attacks against opposition websites, while in weeks without protest it is correlated with *more* DoS attacks. This confirms our theoretical expectation that autocrats choose between tactical reinforcement and tactical substitution when deciding how to employ the tactics in their repertoire of techniques.

In the third paper, I investigate the observation that many governments bring their official websites and digital services online through companies that are based abroad and are thus outside the control of governments. This observation contradicts assumptions from the theoretical accounts of the importance of supply chain security and national data sovereignty. Given the decision to bring official government websites online through foreign companies, I ask what factors might influence the choice in which country a government hosts its own websites. I investigate this question empirically by using Internet measurement data on government website hosting providers and modeling inter-state trust through common alliance membership and relative democratic status. The results of this analysis show that governments are more likely to locate their official websites in countries they trust.

In summary, this dissertation underscores the need to analyze the use of the Internet by governments not only in terms of political content, but also to shed light on the deeper

technical aspects of this use. Furthermore, I show how political scientists can extend their methodological repertoire with Internet measurement techniques to conduct such analyses.

# Contents

# List of Figures

# List of Tables

# 1 Introduction

The introduction of the World Wide Web thirty years ago has fundamentally changed politics. Like other major innovations in information and communication technologies (ICTs), from the invention of the printing press over telegraphs to radio and television, the Web has revolutionised how citizens, civil society organisations, and governments interact and communicate with each other. While early theorists argued that the Internet would be "a global social space [...] naturally independent" of governments (Barlow, 1996), this has not come to pass. On the contrary, governments play an increasingly important role online, and the growing importance of the Internet has brought a variety of new challenges as well as opportunities for governments.

It is not surprising, then, that the Internet revolution has spawned a number of new subfields in political science and introduced new questions to existing fields. In international relations, new questions around digital data sovereignty (Irion, 2012), Internet governance (Dutton and DeNardis, 2013; Mueller, Schmidt and Kuerbis, 2013) and the potential of cyberwar (Harknett and Smeets, 2020; Rid, 2012) have received regular attention. Cybersecurity as a research topic straddles a gap between the international relations literature and the new field of e-Governance, investigating the risks of increasing government reliance on the Internet to perform administrative functions (Paquette, Jaeger and Wilson, 2010; Morrison, 2013) and how secure government websites are (Thompson, Mullins and Chongsutakawewong, 2020; Norris et al., 2019; Caruson, MacManus and McPhee, 2012). At the nexus of autocracy research, repression, and social movements scholarship has converged into a new literature on authoritarian interference with the Internet and and "digital toolkits" of repression (Keremoglu and Weidmann, 2020;

Hellmeier, 2016).

However, despite much fruitful work across fields, political science has not yet exhausted the supply of research questions concerning the nexus of governments and the Internet. We still know little about how countries compare in terms of cybersecurity, and the choices they make in bringing official websites and services online. Studies of online censorship tactics remain confined to single tactics (Keremoglu and Weidmann, 2020), and the literature on international relations in particular largely focuses on theoretical considerations of cyberwar and national sovereignty in cyberspace. Where studies in these fields do engage in empirical work, it is often confined to case studies that tell us little about the larger picture or focused on "application layer data" (Keremoglu and Weidmann, 2020) such as website or social media content. One reason for this is likely to be found in missing methods to consistently consider the technical underpinnings of how the Internet actually works and collecting empirical data that can be used for comparative analysis.

Taken together, this presents two problems in political science research: missing insight into how political behavior plays out beyond the immediately visible parts of the Internet (the network layer) and data gathering methods to close this gap. To address these problems, the goal of this dissertation is to explore three facets of how governments engage with the new possibilities and challenges of the Internet, and to provide insight into how researchers can use Internet measurement data to investigate questions at the intersection of politics and the World Wide Web. I show new ways to compare government's defensive cybersecurity capabilities, how autocratic governments choose between different tactics in their digital censorship repertoire, and how governments navigate the inherently international landscape of how the Internet works when publishing their own websites. To do this, I expand beyond the dominant ways of gathering data in current research and rely on data gathering methodologies and data sources from the computer sciences. In that, my dissertation does not just expand our understanding of how governments engage with the challenges and opportunities presented by the Internet, but also expands the toolkit available to political scientists to investigate political behavior online.

## 1.1 Governments and the Internet

As the Internet and the World Wide Web in particular has become the dominant venue for political interaction, governments have been forced to engage with this new technology in different ways. In this section, I discuss research from three different fields related to how governments use and engage with the Internet. First, I discuss previous work on the nexus of governments and cybersecurity. Second, I introduce related literature on how governments use ICTs to control information. Third, I provide an overview of research concerning international relations and the Internet. Finally, I summarize the

gaps identified in these literatures.

### 1.1.1 Governments and Cybersecurity

The main benefit of the Internet as a whole has been the dramatic improvement in access to and sharing of information between individuals, businesses and institutions. Motivated by a desire to make administrative processes more efficient and less costly, as well as to expand access to services for citizens (Ni and Bretschneider, 2007), governments have increasingly moved parts of their communication infrastructure and citizen-facing services online. However, this shift to digital infrastructure has meant an increase in vulnerabilities for governments from multiple directions.

By moving much of their data and communication to services connected to the Internet, governments now face digital threats from political activists, criminals, and rival governments. Activists see government websites as targets for their political messaging – taking government websites down or "defacing" them, i.e. replacing a website with a political message, can send powerful messages to the citizens accessing these websites. Examples include defacements of e-Government portals in the Philippines or Myanmar to protest against the government (Neil, 2021; Cimpanu, 2021). The data stored on government servers connected to the Internet is also of interest to activists (Lee, 2020), as well as for criminals who can sell personal identifiable data on the black market (Norris et al., 2019). Breaches of important government sites like the CIA's system for communicating with covert sources (Dorfman and McLaughlin, 2018) or the Office for Personnel Management (McLaughlin and Dorfman, 2019), or the hack of the German Bundestag in 2015 by Russia (Von Der Burchard, 2020), show that governments also have to consider rival nations as threats to their digital systems (Selby, 2017).

The increasing reliance of governments on the Internet has received theoretical attention in the literature from two vantage points. The first, smaller, literature deals with general risks associated with moving government process and data online, while the second, larger, literature mainly situates cybersecurity in the context of the threat of cyberwar. Because of this deep link to international relations, I examine this literature more closely in subsection 1.1.3. Noting that "federal government efforts to quickly adopt and adapt to a new technology have not always been successful" (Paquette, Jaeger and Wilson, 2010, p. 251), Paquette, Jaeger and Wilson (2010) conduct a comprehensive review of the risks associated with government use of cloud computing. Categorizing risks to government cybersecurity along four dimensions (Access, Availability, Infrastructure, and Integrity) Paquette, Jaeger and Wilson (2010) highlight risks in the form of contractor and internal compliance, access problems in the event of network interruptions and need for IT governance mechanisms that are able to "identify, assess, and mitigate the risks" (Paquette, Jaeger and Wilson, 2010, p.251), concluding that "the tendency to implement cloud infrastructure and worry about the consequences later will lead to unpredictable

and undesirable consequences to the nation's information" (Paquette, Jaeger and Wilson, 2010, p.252). Going beyond the risks outlined by Paquette, Jaeger and Wilson (2010), Morrison (2013) notes that governments also face threats across the entire supply chain of technology providers. One cited example is Stuxnet, a malicious computer worm introduced into Iranian IT systems via the supply chain (Morrison, 2013, p.755). More recently, concerns over supply chain risks for government cybersecurity have culminated in a ban of Huawei, a Chinese IT equipment provider, in the United States and efforts to convince US allies to enact similar bans (Whalen, 2021).

Previous work has also begun to study the practical implementation of cybersecurity across different levels of government. Norris et al. (2019) show through a nation-wide survey of local governments in the United States that local governments constantly face cyberattacks, yet still lack in adequate implementation of cybersecurity practices, highlighting a problem identified by Paquette, Jaeger and Wilson (2010). Work by Caruson, MacManus and McPhee (2012), which surveys Florida county officials, suggests this might be a consequence of a lack of knowledge by local public administrators. Going beyond survey data, some limited work has also used observational data on cybersecurity practices. Zhao and Zhao (2010) use web content analysis, information security auditing, and computer network security mapping to investigate the security of US state e-government websites, finding that most websites exposed information that could be used for hacking. In a similar study, Thompson, Mullins and Chongsutakawewong (2020) audited and compared security practices of e-government websites in Australia and Thailand, finding that "Australian e-government sites do not significantly differ from Thai sites in their vulnerability level" (Thompson, Mullins and Chongsutakawewong, 2020, p.1) and many websites expose security vulnerabilities.

Overall, government cybersecurity is being increasingly recognized as an important field of study. However, previous research is limited by its high-level treatment of cybersecurity and its lack of comparative work across countries. Current research largely treats cybersecurity as a "black-box" phenomenon, without closer investigation of the actual technologies and infrastructure that governments use for their digital systems. Where researchers do look closer, studies focus on within-country comparisons between local or state governments. Additionally, researchers often rely on survey data – either from government employees or other experts. I will return to these shortcomings in subsection 1.2. In the following sections of this chapter I will first discuss literature from three different fields related to how governments use and engage with the Internet. I will then provide an overview over the gaps in the literature and the contributions of this dissertation and finally conclude with a synopsis of this thesis.

### 1.1.2 Governments and Information Control

As the Internet has become the main venue for political discourse, it has been recognized as both a potential "liberation technology" (Diamond, 2010) as well as a tool for repression and censorship (Lessig, 1997; Morozov, 2011). Governments across the world not only use to the Internet to provide services or make their administrations more efficient, but also surveil their citizens, broadcast propaganda, and censor online content critical of government performance. The latter aspect is usually linked to governments' desire to secure their power. If citizens are free to express their discontent and coordinate, governments may be held accountable for not delivering on their promises. This provides an incentive for governments to control information and restrict online communication between citizens (Roberts, 2018, p. 21f) and results in a global rise of efforts in information control, both in democracies as well as autocracies (Sundara Raman et al., 2020).

How governments attempt to censor online content is a lively field of literature. In their extensive review of the field, Keremoglu and Weidmann (2020) outline three distinct "layers" of Internet technology on which governments can interfere with online expression. Most work on online censorship has been conducted on the highest of these layers, the application layer (Keremoglu and Weidmann, 2020). Studies here predominantly rely on social media data for their analysis, ranging from experimental evidence on how social media posts are reviewed and censored (King, Pan and Roberts, 2014) to observational data on how sudden censorship of one social media network can lead users to circumvent censorship technology and increase their access to censored information Hobbs and Roberts (2018). On the second, network, layer, studies have investigated how governments use filtering technology to block access to websites (Hellmeier, 2016; Deibert et al., 2008) or use tools such as Denial-of-Service attacks to temporarily take down websites critical of the government (Lutscher et al., 2020; Nazario, 2009).

Generally, previous research has built a good understanding of when and how governments use tools for digital information control. However, there remain at least two shortcomings of the literature worth highlighting. First, that the dominance of social media data has created what Keremoglu and Weidmann (2020) call "tool-dependence", where researchers rely on the features, dynamics and access points specific to particular social networks as well as restricting cross-country comparisons. Because the content on digital social networking sites comes in different languages, large-scale and cross-language comparisons are generally infeasible absent new tools for automatic translation. However, studies like Lutscher et al. (2020) and Hellmeier (2016) show that comparative research is possible when focusing on the network layer. Second, that most studies (across network layers) focus on the use of a single censorship tactic (Keremoglu and Weidmann, 2020), which may lead to misleading results: do we see little usage of a tactic because the government sees no need to censor, or because they have opted for another tactic?

### 1.1.3 The Internet and International Relations

For governments, the Internet poses a new set of challenges distinct from previous new technologies. More than any other technological innovation, the Internet challenges conceptions of sovereignty, state power, and international governance. Over the last thirty years, a diverse and expansive literature has developed around these challenges and how governments deal with them.

While the concept of sovereignty has been controversial in political science (Oppenheim, 1912), a common definition emphasizes that state sovereignty requires "supreme legitimate authority within a territory" Philpott (1995, p. 357). This emphasis on territory is challenged by the technical reality of the inherently internationally networked nature of the Internet. While websites may be physically located, or rather originate, in a particular country, they are accessible from anywhere in the world without crossing borders like physical goods would. Governments and theorists have responded to this challenge by "territorializing cyberspace" in different ways (Lambach, 2020). This territorialization has been met with critique – Mueller (2020) argues that "attempts to apply sovereignty to cyberspace governance are [...] inappropriate to the domain" and "threaten existing transnational governance institutions in the private sector" (Mueller, 2020, p.2). This argument has, however, not deterred governments from expanding effort to increase control over "their" parts of the Internet.

One particular way in which governments have sought to carve out "territory" on the Internet has been through data localizations laws in an attempt to protect their national data sovereignty (Irion, 2012). Nigeria, for example, requires that all government data is hosted in datacenters located within Nigeria (Bowman, 2015), as does India (PTI, 2017). Countries like Brazil, France, Russia, Malaysia and others have enacted laws in a similar spirit (Bowman, 2015; Chander and Lê, 2015; Ewing, 2018). In addition to securing data from signal intelligence of other countries (Selby, 2017) and tighter control over general security, some governments also see such laws as opportunities to strengthen their local IT sector (Chander and Lê, 2015). Further, efforts to territorialize the Internet are not only carried out through enacting domestic laws requiring local data storage, but also through efforts to influence Internet governance. Here, governments "[insert] themselves into the technical and operational networks and [attempt] to shape standards and practices in a multistakeholder environment" (Mueller, Schmidt and Kuerbis, 2013, p.100).

Even though some researchers have stated that "cyber war will not take place" (Rid, 2012) or cast doubt that cyberwar can actually work (Gartzke, 2013), much of the discussion around government use of the Internet and cybersecurity revolves around the threat of cyberwar. Highly public cases of intrusions into government systems by rival nation states, such as the CIA's system for communicating with covert sources (Dorfman and McLaughlin, 2018), the Office for Personnel Management (McLaughlin and Dorfman, 2019), or the hack of the German Bundestag in 2015 by Russia (Von Der Burchard,

2020), have touched multiple questions covered in the literature. Can such attacks be deterred (Nye Jr, 2016)? When should states publicly attribute breaches of their systems (Egloff and Smeets, 2021)? How should states judge the potential threat from escalating cyber operations (Fischerkeller and Harknett, 2019)? And why has there not emerged a global governance regime for cyberweapons (Stevens, 2018)?

The main shortcoming of the literature on international relations and the Internet is the dominance of theoretical discussions over empirical accounts. While some work has investigated policy choices and compared these across countries, the *technical* choices of ICT implementation and their correlates have been mostly ignored. We know little about the technical choices governments make in their digital infrastructure, and how these choices translate into how governments defend themselves against attacks.

## 1.2 Gaps and Contributions

In this section I briefly summarize the main gaps in the previous literature and how my dissertation addresses them.

### 1.2.1 Political Behavior on the Network Layer

The preceding review of the literature has surfaced that previous work on the nexus of politics and the Internet predominantly relies on application layer data. While these studies have certainly made important contributions to our understanding of the nexus between governments and the Internet, we still lack understanding in many questions that go beyond what is immediately visible *on* the Internet: who do political actors, and governments in particular, trust with providing their digital infrastructure and hosting their websites? How secure are these websites comparatively? How do governments choose between different technical measures of information control? In short, how does political behavior play out on the network layer? I remedy this shortcoming in this dissertation by demonstrating how researchers can use Internet measurement techniques to "go deeper" and use new data to answer questions about political behavior on the network layer.

### 1.2.2 Disparate Empirical Consideration of the Underlying Technology

The preceding review of the literature has also shown that different fields vary widely in their empirical consideration of the underlying technology of the Internet. While all fields *problematize* the technology, *empirical investigations* remain relatively sparse. Studies from the information control literature rely the most on empirical data, across all three layers described by Keremoglu and Weidmann (2020). In contrast, the international relations literature remains firmly in the grasp of theorists, with some work that compares policy choices between countries but still treats the technology itself as a "black box".

The literature on cybersecurity also often remains on the level of theory, with limited empirical accounts that mostly rely on survey data over technical, observational data. In sum, the literature lacks a consistent empirical consideration of the underlying technology of the Internet. By introducing Internet measurement techniques, I demonstrate how researchers can avoid treating Internet technology as a "black box" phenomenon, investigate existing questions with observational data, and answer new questions.

### 1.2.3 Missing Comparative Studies

Where studies in the different literatures *do* conduct empirical research, they often remain constrained to case studies or within-country comparisons. However, this ignores the inherently international nature of the Internet: how a government uses and relies on Internet technology does not occur in isolation or a vacuum, but is necessarily connected to the behavior of other governments. This means that to understand cybersecurity, information control, and the role of the Internet in international relations comparative research is necessary. In my dissertation I contribute three such comparative studies. In the first paper (Chapter 2), I provide an example application of network measurement data analysis to investigate government cybersecurity across countries. In the second paper (Chapter 3), we empirically study the interplay of different online censorship techniques in autocracies and provide evidence that autocrats select tactics from their censorship repertoire depending on the current level of contention. In my third paper, I use similar data to investigate where governments host their official websites.

## 1.3 Dissertation Synopsis

In this section I summarize the arguments, findings and implications of the three papers.

### 1.3.1 Paper 1 – More to Measure: Internet Measurement Data for the Social Sciences

The goal of the first paper (Chapter 2) is twofold: I first provide an introduction to Internet measurement techniques and data, and how political scientists can use these to investigate previously unanswered questions. I then provide an application of this approach to the topic of cybersecurity.

Because the Internet is a network of networks, consisting of billions of individual endpoints that provide data on how they connect to each other and what they are used for, network measurement techniques can provide a significant amount of information relevant for researchers. This data has a number of advantages to what would otherwise be available for analysis: data gathering can be automated, provides much higher fidelity than expert interviews or manual collection, and is unaffected by human biases such as coder fatigue.

Using Internet measurement data, I analyse government cybersecurity across countries. I construct a new observational indicator of defensive government cybersecurity capability and compare it to an indicator based on expert surveys. My analysis shows that the observational indicator plausibly measures the same concept as the indicator based on expert surveys and that expert surveys might be biased by media coverage of security breaches in a way observational indicators are not.

### 1.3.2 Paper 2 – Attack or Block? Repertoires of Digital Censorship in Autocracies

The goal of the second paper (Chapter 3), written with Nils B. Weidmann and Alberto Dainotti, is to study the interplay of different online censorship techniques. Focusing on autocracies, we study the relationship between *website blocking* and *cyberattacks* (Denial-of-Service) using Internet measurement techniques.

We address a gap in the existing literature resulting from a focus on single tactics by studying how autocrats might choose different online censorship tactics depending on the current situation. We argue that there are two possible relationships between website blocking and Denial-of-Service (DoS) attacks. Autocrats can choose between "reinforcement," where they combine different tactics, and "substitution," where the use of one of the tactics goes along with a reduced reliance on the other. Because tactical reinforcement or substitution may not be constant over time, we study the relationship between different censorship tactics depending on the political situation on the ground, distinguishing between periods with mass protest and those without.

For our empirical analysis, we rely on Internet measurement data to contribute a new measurement to better map Denial-of-service Attacks to possible targets. By creating a new measurement we call a "stable IP period," based on a large, existing corpus of historical web-crawls, we are able to connect data on DoS attacks to potential targets, alleviating a constraint in previous studies. Together with data on the blocking of opposition websites in autocracies, we test whether the presence of one type of censorship (blockings) is statistically related to the occurrence and magnitude of the other (attacks).

The results of our analysis provide first evidence that autocrats select tactics from their censorship repertoire depending on the current situation. In weeks with protest, observing the presence of website blocking is associated with *fewer* DoS attacks against opposition websites, while in weeks without protest it is correlated with *more* DoS attacks. This confirms our theoretical expectation that autocrats choose between tactical reinforcement and tactical substitution when deciding how to employ the tactics in their repertoire of techniques.

### 1.3.3 Paper 3 – Trust and Safety: Where do Governments Host Official Websites?

The goal of the third paper (Chapter 4) is to investigate an empirical puzzle related to the choices governments make when hosting their official websites. Even though previous research strongly suggests that governments should host their official websites domestically, motivated by concerns over supply chain risks and considerations of national data sovereignty, most countries do host official government websites with businesses outside of their jurisdiction. This raises the question which factors might influence a government's choice to host in a particular country? I argue that this choice depends on how trustworthy the "hosting" country is seen.

The empirical analysis leverages measurements of the hosting location choices of official government websites and models trust between countries in two ways. The first indicator is an explicit signal of trust between two countries in the form of shared alliance membership. The second indicator models trust through the democratic credibility advantage, coding whether the hosting country is more democratic than the country that owns the government websites. I show that governments indeed choose to host their official websites with countries they trust and that this holds for both operationalizations of trust.

# 2

# More to Measure: Internet Measurement Data for Political Science

## 2.1 Introduction

With the dominance of the Internet as the main venue for political interaction, a lot of research attention has shifted towards analysing political behavior online. Research has investigated the possibly dangerous effects of algorithmic filter bubbles on political polarization, the spread and impact of disinformation and "fake news", or the use of social media as a communications tool by political leaders, among other things. Most of this research has focused on the "visible" part of the Internet where users engage with each other, the *application layer* (Keremoglu and Weidmann, 2020). Research on the application layer has been aided by the relative ease with which political scientists can gather large amounts of data. On the *network layer*, gathering such data has been more difficult. In this paper, I use government cybersecurity as an example to demonstrate how political science can use *Internet measurement data* to learn more about how political actors use the Internet beyond what is immediately visible to citizens.

Internet measurement data is collected by either passively observing or actively probing data flows and hardware connected to the Internet. Collecting Internet measurement data relies on the public interfaces which websites, servers, and other devices use to make themselves and their data available on the Internet. This in contrast to application layer data, which is most often collected through proprietary programming endpoints provided by, for example, social media websites. Using Internet measurement techniques,

researchers can collect data on who owns a particular webserver, where that server is located geographically, what software runs on the server and how secure it is, whether a website saw any interruptions in availability, the content available on the website, and much more.

Previous research on the network layer has often focused on censorship and information control. Gohdes (2015) uses data on Internet network outages to investigate whether they are used to weaken opposition groups in the Syrian civil war and restrict the reporting of fatalities. Dainotti et al. (2014) uses data on Internet access interruptions in Egypt and Libya to investigate the strategic use of online censorship. Beyond case studies, there are also comparative studies focusing on cross-country contexts: Weidmann and Rød (2019) use Internet penetration data in autocracies to investigate the occurrence and persistence of protest as a consequence of Internet availability, and Lutscher et al. (2020) use sophisticated data measuring DDoS attacks to research the occurrence of cyberattacks during election periods. In addition to these important questions, data from the network layer can also aide in investigating other topics.

In this paper, I demonstrate how political science research can use Internet measurement data to answer questions that go beyond what is directly visible to Internet users. Who do political actors, and governments in particular, trust with providing their digital infrastructure and hosting their websites? How secure are these websites comparatively? Early efforts to map political presence online, such as Norris' 2001 investigation into the types of government departments that had an online presence, relied on manual data collection. Since then, the exponential growth of the Internet has meant that manual collection of this kind of data has become infeasible, which has stifled further research. But because the Internet is a network of networks, it is possible, for example, to collect data from the servers that host government websites that can tell us much about government presence and behavior online. By using automated collection of Internet measurement data, such as website content, IP addresses and other measurements, we can investigate more questions, including those that focus on latent concepts, in a rigorous manner.

To demonstrate the use of Internet measurement data, the main focus of this paper is one such latent concept: defensive government cybersecurity capability. Recent incidents show the importance of this concept. Examples include the SolarWinds attack (Fireeye, 2020) that infiltrated government supply chains in the US, the Office of Personell Management data breach (Chaffetz, Meadows and Hurd, 2016), or the hack of the German Bundestag in 2015 by Russia (Von Der Burchard, 2020). Previous research has used expert surveys to create an indicator for this latent concept (Mechkova et al., 2019*a*), yet expert surveys come with drawbacks: heavily reported incidents might bias the expert indicator more than is warranted, for example. Additional, objective data-based measurements can counter this potential bias. In this paper, I compare the results of an analysis of a new indicator for government's defensive cyber security capabilities to an indicator published by the V-Dem Institute (Mechkova et al., 2019*a*). My analysis

relies on new Internet measurement data on the severity of security vulnerabilities found on government-owned servers of 154 countries. Through my analysis, I show that the observational indicator plausibly measure the same concept as the indicator based on expert surveys. I also show that expert surveys might be biased by media coverage of security breaches in a way observational indicators are not.

## 2.2 Related Literature

Political Science has a rich history of embracing new data sources to study political behavior. Importantly, this history shows that using new data sources is not an end in itself, but advances our understanding of substantive research questions. Bulmer, Bales and Sklar (1991) trace the use and growing sophistication of survey data as one of the first tools of quantitative social science research, the first use of GIS data by Richardson (1960) established a new way of investigating questions around conflict (Gleditsch and Weidmann, 2012), and newspaper data has become a central source for insights into collective action around the world and other topics (Earl et al., 2004). It is no surprise, then, that data gathered on the Internet has been equally embraced by political scientists.

Closely linked to the growing computational capabilities around processing text data, the communications literature in particular has embraced data originating on the Internet. As Theocharis and Jungherr (2020) point out, this data is used to investigate questions that have been "always asked: how can we reliably measure the reach of specific media outlets or political actors, identify the often-overlapping media and information diets of people, and estimate the effects of information – especially in new, noisy and deeply confusing information environments?" (Theocharis and Jungherr, 2020, p. 1f). Studies in this area look into how Chinese propaganda needs to employ clickbait tactics to drive engagement online (Lu and Pan, 2020), how rebels use social media networks to create international support for their cause (Jones and Mattiacci, 2019), or how political leaders adopt social media as a communication tool when they face political pressure (Barberá and Zeitzoff, 2018).

But research fields outside of traditional political communication also leverage online data to answer long-existing questions from new vantage points. In international relations research, for example, Duncombe (2017) provides a case study of the use of social media platforms as a tool of diplomacy and negotiation – a new vantage point for a traditionally hidden process. In comparative politics research, particularly studies focused on autocracies or collective action and protest, Internet data has also attracted considerable interest to answer existing questions: King, Pan and Roberts (2013) investigate censorship of social media platforms in China, Enikolopov, Makarin and Petrova (2020) study the role of social media in protests in Russia, and Larson et al. (2019) investigate real-world protest attendance in France in 2015 using Twitter data.

However, close inspection reveals that a large majority of the studies using Internet

data rely on only one slice of the available data: application layer data (Keremoglu and Weidmann, 2020) and more specifically *content* data of social networks. In essence this means that research is mostly using data *available on* the Internet compared to data *about* the Internet. This is not to say that other data is never used: Dainotti et al. (2014) use Internet network interruption data to investigate censorship in Egypt and Libya during the Arab Spring, Gohdes (2015) uses network blackouts data based on Google service interruptions to investigate the connection between their occurrence and military action in the Syrian civil war, Weidmann and Rød (2019) use Internet penetration data to investigate protest in autocracies, and Lutscher et al. (2020) uses sophisticated data measuring DDoS attacks to research the occurrence of cyberattacks during election periods.

What the studies leveraging data on Internet penetration, network traffic and DDoS attacks point to is that data *about* the Internet allows us to research many more questions, old and new, than just those that can be answered through data available *on* the Internet. Are cyberattacks an *effective* tool for censorship or weapon for the weak? Do autocrats use different attacks in different contexts, employing a digital strategy (Keremoglu and Weidmann, 2020)? Where do governments around the world host their websites – do they keep them on domestic servers or do they trust other jurisdictions than their own with their data? How secure are these websites, comparatively? How large are government websites, how many pages do they have, and are there consistent differences between departments? How are civil society organizations represented online? How interlinked are they – within one country and internationally?

What can we do to make progress in this wider field of questions? The reason that studies relying on application layer data are so dominant is likely that gathering content, and in particular social media data, is much easier to collect than other kinds of Internet data. Using human-coded data is only practical for very high-level concepts and questions. Early efforts to manually collect and categorize all government websites on the Internet like carried about by Norris (2001) would be impossible today, given the sheer size of the Internet. To answer new questions, then, we need to look to the field of network measurement, a subfield of computer science that is concerned with, among other things, investigating the performance, topology and usage of the Internet. So far, this subfield has found applications in the social sciences mostly when investigating social networks – measures like network centrality, density and the like were first heavily used to understand how the Internet works before being applied to accounts on Twitter or Facebook. However, network measurement applied to the Internet encompasses much more – it can provide information on all levels of the "network stack", from the hardware that runs a given website over the ownership structures of the infrastructure between the hardware, to which programs run there, how safe they are, and how content and linking between websites changes over time.

Still, even when using network measurements to gather data for political science pur-

poses, challenges arise that make gathering and using the data difficult. In the next section, I will explain these challenges, how to overcome them, and then present an example application of a dataset that can serve as a starting point for investigating new questions.

### 2.2.1 Internet Measurement for the Social Sciences

The Internet is a network of networks, consisting of billions of devices ranging from cellphones and personal computers to servers in datacenters or smart-home appliances like Internet-connected security cameras, weather stations or fridges. Each of these devices is reachable through its Internet Protocol (IP) address, a numerical label roughly analogous to a postal address. Like postal addresses, IP addresses belong to administrative groups, called Autonomous Systems (AS): if we take the street address “4 Pine Road, Philadelphia" the IP address would be “4 Pine Road" and the AS would be “Philadelphia". Autonomous Systems are nested (like cities in states in countries) and connect to each other, eventually making up the Internet. Understanding *how* the ASes and the individual machines in them connect, the patterns in the flow of traffic between them, and what they are used for is the what the field of network measurement specializes in.

Because the digital nature of the Internet removes many physical challenges (“pinging" every public IP address to see what information is returned is possible, whereas sending a letter to each postal address in the world is impossible), network measurement can provide a significant amount of information relevant for researchers. Each website has a registration date attached to it, and what domain registrar manages the registration. The server that hosts the content of a website has an IP address, which is administered by its AS, which has publicly listed ownership information. By sending generic data requests to the “ports" of a server (dedicated communication endpoints), it is possible to learn what software runs on a given server: Wordpress, email software, or some kind of database? This in turn allows insights into the security of the server: is the software outdated and has security vulnerabilities? And, of course, the content of the websites themselves: what language is used? how many pages are there? where do these pages link? how do content and links change over time? Network measurement allows answering these and many more questions about individual connected devices, as well as higher-level questions: how does traffic flow between ASes, and countries? How centralized is the ownership structure of Internet providers in a country, is there evidence for censorship?

The data generated through network measurement techniques has a number of advantages to what would otherwise be available for analysis: first, gathering the data can be automated, requiring no human intervention after the initial setup. Second, it can be much more encompassing – in many cases it is possible to observe *all* devices, and not rely on manually selected subsets or sampling. Third, because it is machine-collected

observational data, there are no human biases involved[1]. Taken together, these factors present a significant advantage even when investigating latent concepts that would otherwise require expert surveys or other proxy measurements.

The biggest challenge when using Internet data for research is gathering historical data. Because the Internet is so big and ever-expanding, keeping historical records for all of it is impossible. For researchers, this is a problem: link-rot, the disappearance of websites others link to, means that only what is recorded and archived in the moment can later be reliably analyzed. Retrospective coding of data, like often used with newspaper articles available in extensive archives, is impossible unless a given website is archived. This challenge of missing historical data also goes deeper than the surface-level content of websites: the mapping from domain to IP address can change at a moments notice, and because the tables that contain these mappings are so large in their own right, no historical records are available even for this comparatively small subset of Internet data. This means that researchers have to gather the data themselves, restricting their analysis to cross-sectional designs looking at the present moment, or starting data gathering far in advance, which comes with its own challenges of data management, questions of copyright, and other headaches. The other option is to rely on data gathered by others.

A number of institutions are dedicated to gather network measurement data and making it available to researchers. By combining these datasets and only gathering additional data where needed, political scientists can answer many relevant questions without needing to build dedicated infrastructure, and with the option to carry out historical analyses even today. The by far largest data source is the Common Crawl archive. Every four weeks, Common Crawl generates a new dataset that contains content and metadata of billions of webpages based on following the links on a large sample of websites from the whole Internet. This dataset is made freely accessible to anyone and provides historical data from 2015 onwards. Another data source is the Center for Applied Internet Data Analysis (CAIDA) at the University of California, San Diego. CAIDA generates multiple datasets available to researchers, including historical data for which IP addresses belong to which Autonomous System, who owns these ASes, DDoS attacks, and more. Where the previous to examples generate general datasets that can be used for political science purposes, the Open Observatory Network Initiative (OONI) focuses explicitly on politically relevant websites and provides historical data on online censorship.

## 2.3 Application: Government Cybersecurity

To demonstrate the use of Internet measurement for data collection relevant to the social sciences, I compare the results of an analysis of government's defensive cyber security

[1]Naturally, human biases still play a role in what data is gathered in the first place and when it is analysed. During the automated data gathering process, however, issues like coder fatigue, boredom or other human failings do not play a role.

capabilities to an indicator published by the V-Dem Institute. Government cybersecurity provides an instructive example because it is both a highly salient topic as well as a generally latent concept that is hard to measure. In recent years, cyber attacks against government assets connected to the Internet, ranging from websites to personell data or election infrastructure, have been an increasingly common topic of media reports, in particular when these attacks were alleged to be "nation-state attacks", i.e. carried out or directed by other governments. Recent examples for this include the SolarWinds attack (Fireeye, 2020) that infiltrated government supply chains in the US, the Office of Personell Management data breach (Chaffetz, Meadows and Hurd, 2016), or the hack of the German Bundestag in 2015 by Russia (Von Der Burchard, 2020).

Given the high salience of government cybersecurity, finding ways to study how capable governments are in this domain comparatively is a worthwhile pursuit for political scientists. However, defensive capabilities in cybersecurity are also a highly latent concept that is difficult to evaluate from the outside. The reason for this difficulty is intuitive: governments are highly incentivized to keep information on how they secure their infrastructure secret, and commonly used channels to infer similar information, such as reports on military spending, also offer little data. Hence, the best currently available data relies on expert surveys collected by the V-Dem project and described in the next section. As an alternative to expert surveys I construct a new indicator that relies on observational data.

### 2.3.1 V-Dem Indicator and Methodology

The V-Dem Project provides a global dataset that provides disaggregated, multidimensional codings of the concept of democracy, relying on expert surveys and a unique measurement model (Coppedge et al., 2016). Since version 9 of the dataset, it also provides data on digital societies, including an indicator for government cybersecurity (Mechkova et al., 2019*a*). Based on the overall V-Dem methodology, it collects expert ratings for the question "Does the government have sufficiently technologically skilled staff and resources to mitigate harm from cyber-security threats?" (Mechkova et al., 2019*a*,*b*). Expert responses to this question, ranging from '0' (No) to '4' (Yes), are aggregated according to the V-Dem Methodology (Pemstein et al., 2019).

The benefit of the V-Dem methodology is that the measurement model (Pemstein et al., 2019) accounts for differences in how experts apply rating scales as well as the differences in their substantive ratings, allowing V-Dem to provide estimates for latent concepts that are otherwise hard or impossible to measure. As Marquardt et al. (2019) show, these estimates are reliable and it is generally safe to use aggregated expert-coded data. However, there remains a weakness to expert-coded data that even sophisticated item-response theory models are unable to circumvent. While Marquardt et al. (2019) show that the V-Dem methodology can effectively control for biases *among* a set of experts

on the level of individual-to-expert-consensus and the expert-to-perception-of-scale level, *accuracy* is impossible to evaluate given the nature of expert-coded surveys (Marquardt et al., 2019, p. 7).

Yet, absolute accuracy is not the most crucial factor when looking at indicators over time: the stability of the perception-accuracy error, i.e. relative accuracy, is. Expert perception might be accurate in one year, but dramatic news events can widen that error because it biases the perception of all the experts. This too is unavoidable – the premise of expert surveys is that their collective judgment is influenced by changes all of them perceive. Accordingly, it is important to not exclusively rely on expert-surveys as a data source, but augment it with observational indicators wherever possible. I construct such an indicator for government cybersecurity in the next section.

### 2.3.2 An Internet Measurement-based Indicator

Before we can create an Internet Measurement-based indicator for government cybersecurity, we need to address one challenge: defining the IP addresses and websites we are interested in. While network measurement studies in the computer sciences often target the Internet as a whole, this paper requires a narrow subset: IP-Addresses used to host government websites and services. Some researchers use manual collection to identify websites relevant to their research question: to investigate censorship, the Citizenlab project (Citizen Lab and Others, 2014) manually collects lists of potentially censored websites for each country in the world. In one of the earliest surveys of government presence online, Norris (2001) also relied on a manually collected list of government websites. As the Internet has grown exponentially, this has become infeasible – even the UN E-Government Survey (United Nations, 2020*b*), which aims to evaluate the e-government performance of countries, restricts itself to a single website per country for its evaluation (United Nations, 2020*a*).

The challenge of identifying government-owned websites can be largely solved by relying on a feature of the domain name system used to translate between human-readable website names (domains) and IP addresses. Most websites on the Internet are accessible through the '.com', '.net', '.org' or country-specific top-level-domains such as '.de', '.mx', and '.fr'. Many governments have chosen to identify official government websites through a special domain suffix: the White House in the United States is available through 'whitehouse.gov', with '.gov' being reserved for US government websites. In France, the ministry for foreign affairs is available at 'diplomatie.gouv.fr', with '.gouv.fr' as the suffix that identifies all french government websites. By manually coding all publicly available suffixes for whether they signify a government website, I have a reliable indicator to classify domains. To gather the data for this dataset, I manually coded the publicly available list of domain suffixes to identify those that signify government ownership of a given domain. Of the 249 country-level domains, 64% used such suffixes. Focusing on

politically relevant countries as coded by Gleditsch and Ward (1999), this results in 136 out of 173 countries (79%) using government-only domain suffixes[2].

Based on the coded domain suffixes, it is now possible to collect a list of government websites for each country. To collect this list, I rely on data collected by the Common Crawl project, a non-profit organization that regularly crawls the Internet and makes this crawl data available. By searching the historical archives of these crawls for domains that end in a government suffix, it is possible to collect a list of all government websites that were ever visited by Common Crawl. To note is that this list is not necessarily exhaustive: Common Crawl is not able to crawl the *whole* Internet every month, but through repeated sampling the breadth of its reach is currently the best available data for this purpose. Common Crawl also collects the IP address, and thus server address, of each website. Using this information plus data gathered from a MassDNS[3] scan for all websites gives a list of all IP addresses on which at least one government website was observed since 2014.

The IP address alone is already very useful for analyzing government presence on the Internet. Based on the address, it is possible to find out who owns and operates the server to which this address points, and geo-locate where the server is physically. Yet to enable insights into the security of these servers, more data is needed. The most common way through which the average user interacts with a server is through the browser. This browser requests website data from the server, which is delivered through a *port* dedicated to such websites requests. However, the same server can at the same time also communicate with other servers or programs through other *ports* on the same machine. For example, websites are most often delivered over port 80, while database connections use port '5432' or '27017'. By scanning these ports, i.e. requesting data from them, security researchers and hackers determine what services are available on a given server. Such a scan does not expose vulnerable data itself, but it returns so called "banner data" about what service is available on a given port, and often also *what version* of the service is used. This information can be used to look for security vulnerabilities associated with a given version of a service and then be exploited to gain access to the server.

While port-scanning can in principle be done from any computer connected to the Internet, this restricts analyses to the time at which the scanning started and going forward. To gather historical data, I therefore rely on Shodan.io, which regularly carries out port-scans of the whole Internet and makes that data available as a service. By submitting all the IP addresses obtained from the previous steps, as well as searching for domains ending in the coded domain suffixes, I receive banner data going back to 2016.

The banner data available through Shodan contains much more information than nec-

[2]Whether countries use suffixes to distinguish government websites does not seem to be driven by systemic factors.

[3]MassDNS is a DNS resolver which allows associating large amounts of domain names to their IP addresses.

| Date | IP | Port | CVE | CVSS |
|---|---|---|---|---|
| 2019-09-10 | 143.196.144.137 | 443 | CVE-2019-0220 | 5.0 |
| 2019-09-10 | 143.196.144.137 | 443 | CVE-2017-3169 | 7.5 |
| 2019-09-10 | 143.196.144.137 | 443 | CVE-2017-7679 | 7.5 |
| 2019-09-10 | 143.196.144.137 | 443 | CVE-2019-0196 | 5.0 |
| 2019-09-10 | 143.196.144.137 | 443 | CVE-2017-7659 | 5.0 |

Table 2.1: Example of Shodan.io banner data for the host dac-infos.dg.aviation-civile.gouv.fr

essary for the purpose of this paper, so Table 2.1 only gives a short excerpt to illustrate. In the table, we see that on 2019-09-10, the hostname `dac-infos.dg.aviation-civile.gouv.fr`, owned by the *Direction des Services de la Navigation Aerienne*, the French Aviation authority, was associated with the IP address '143.196.144.137'. The version of the service available on port 443 (Apache Webserver) was associated with a number of security vulnerabilities (18, five shown in the table), for which the severity of the vulnerabilities are also given. Each vulnerability is identified by its CVE ID Number, a unique identifier assigned to publicly disclosed cybersecurity vulnerabilities (The MITRE Corporation, 2021). How severe each vulnerability is is measured by its Common Vulnerability Scoring System (CVSS) score. This score is an open industry standard developed to grade vulnerabilities according to their severity, ranging from 0 (least severe) to 10 (most severe) (First, Inc., N.d.). The most commonly used version of the score is the CVSS Base score, although modifications exist that allow an organization to take into account factors in their specific environment. The analysis in this paper relies on the CVSS v2 Base score.

Overall, the initial dataset contains roughly 5.9 million Country-IP-Vulnerability observations from 2016 to 2020, covering 135 countries and 40703 government domains across 62164 distinct IP addresses. Based on these observations, I construct a new indicator of government cybersecurity in the following way: I first average the CVSS scores for each IP address across all observations in a given year, giving a 'mean_cvss' score for the IP. IP addresses that have zero vulnerabilities across the observations in a year receive a 'mean_cvss' score of 0, which indicates a severity of "None" according to CVSS v2 (First, Inc., N.d.). I then average across the 'mean_cvss' score for all IP addresses in a given year, creating a 'mean_mean_cvss' score for each country-year. In the following section, I compare this indicator to the V-Dem indicator described in the previous section.

### 2.3.3 Comparison of Indicators

First, I compare the CVE-based indicator to the V-Dem indicator. Because the two indicators are constructed using different scales, a simple comparison of the indicator values is difficult. To ease the comparison between the two indicators, I therefore rescale

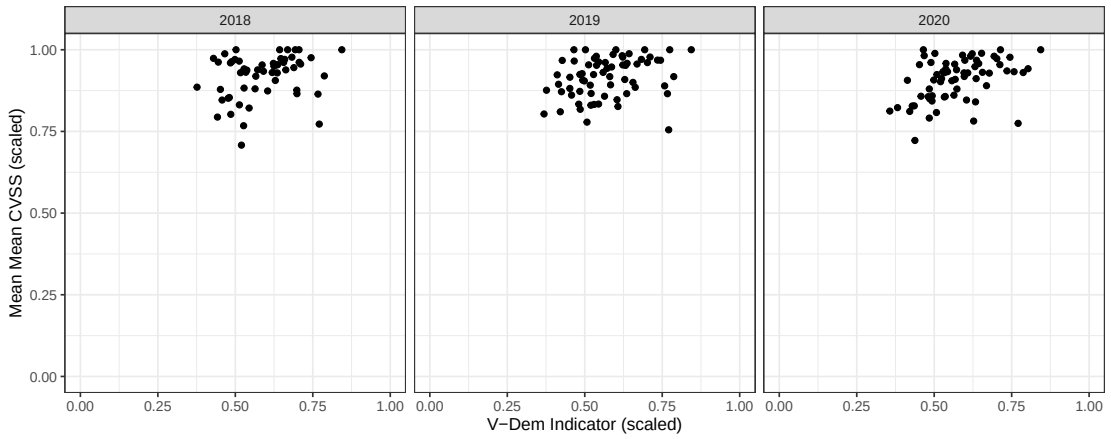


Figure 2.1: Scatterplot of both indicators

both the V-Dem indicator as well as the CVE-based indicator to values between 0 and 1, with higher values signifying a higher security rating[4].

While V-Dem has collected rating data from 2019 on going back to the year 2000, the data available through Shodan is more limited. The historical search in Shodan only covered about 1000 IP addresses total in 2016, about 13,000 in 2018 and 62,000 in 2020. Accordingly, the following comparison will focus on the years 2018 to 2020. Because the number of observed IP addresses is not equally distributed across countries and years, the analysis will only cover countries for which at least 10 IP addresses were observed in all years[5]. Additionally, V-Dem recommends to only consider country-years with at least four expert ratings (Mechkova et al., 2019*a*, p. 3). After removing country-years below this threshold, the resulting dataset consists of 188 observations of 66 countries.

Figure 2.1 shows scatter plots of both indicator values for the period 2018 to 2020. Descriptively, it's noticeable that the measurement-based indicator rates countries generally higher than the expert-coded indicator. Figure 2.2 provides a closer look at four countries for the year 2018 that show the same pattern. For all three countries, the measurement-based indicator shows a higher level of security than the V-Dem indicator. Considering the whole dataset, Table 2.2 shows the result of a simple linear regression between the two indicators for 66 countries across the three years. As the model indicates, there is a statistically significant positive relationship between the two measurements.

#### 2.3.3.1 Change

How do changes in the two indicators compare for particular countries? Figure 2.3 illustrates comparative ratings and changes between 2018 and 2020 for the four previously shown countries: Australia, Czech Republic and France. Table 2.3 cross-tabulates the changes for the three countries and the two indicators. All three countries show notable

[4] A deeper technical explanation for the procedure can be found in the appendix.
[5] This is a common procedure in computer science to avoid false positives.

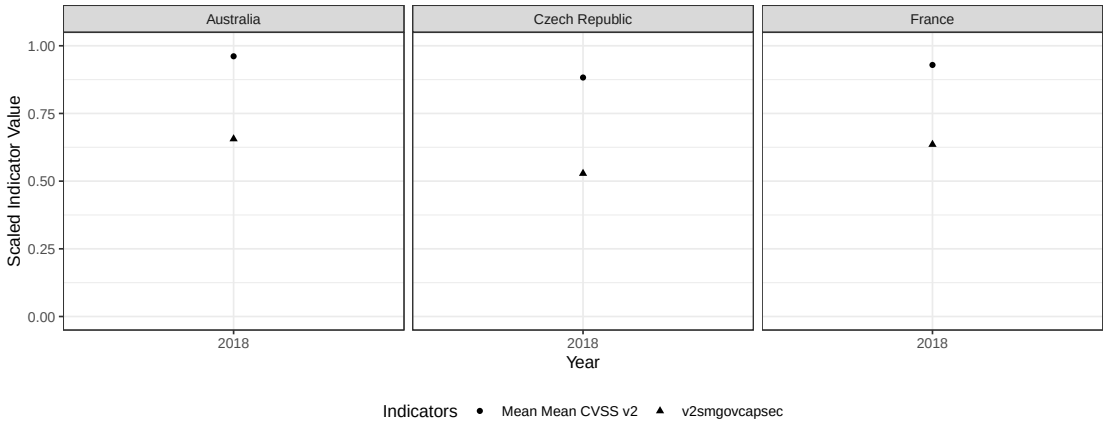


Figure 2.2: Indicator values in 2018 for Australia, Czech Republic and France.

| | Model 1 |
|---|---|
| (Intercept) | 0.08 |
| | (0.11) |
| Mean Mean CVSS Score (scaled) | $0.55^{***}$ |
| | (0.12) |
| $R^2$ | 0.11 |
| Adj. $R^2$ | 0.10 |
| Num. obs. | 188 |

$^{***}p < 0.001$; $^{**}p < 0.01$; $^{*}p < 0.05$

Table 2.2: Linear Regression between CVSS-based indicator and V-DEM indicator

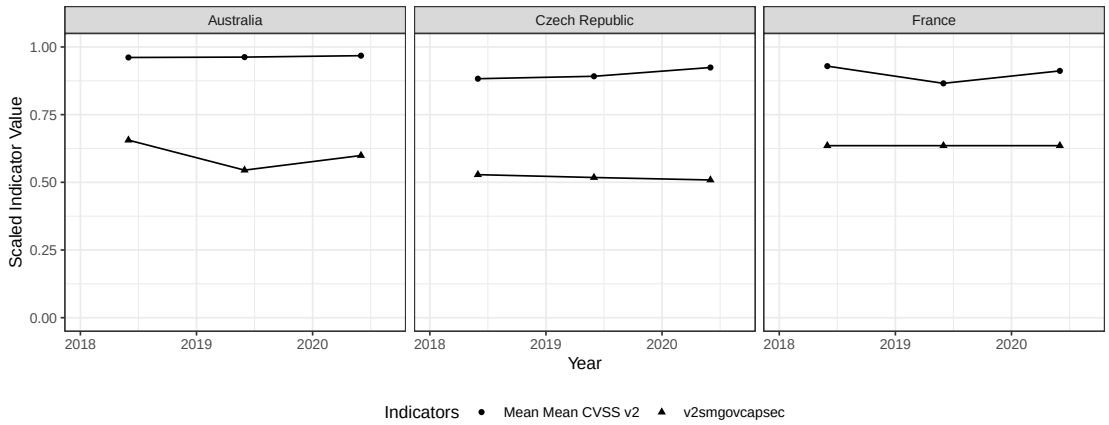


Figure 2.3: Indicator values over time for Australia, Czech Republic and France.

| Country | Year | Change in CVSS-based indicator | Change in V-Dem indicator |
|---|---|---|---|
| AUS | 2019 | positive | negative |
| AUS | 2020 | positive | positive |
| CZE | 2019 | positive | negative |
| CZE | 2020 | positive | negative |
| FRA | 2019 | negative | no change |
| FRA | 2020 | positive | no change |

Table 2.3: Changes between 2018 and 2020 respective to previous year

differences between the indicators and years that are worth investigating.

In Australia, the expert-coded indicator classifies the country as less secure than the observational indicator and shows a steep decline from 2018 to 2019, while the observational indicator remains relatively stable. The steep decline in the expert rating is likely due to a large, state-sponsored cyberattack against Australia publicised in February of 2019 (BBC News, 2019). In the Czech Republic, both indicators diverge over time, with the observational indicator improving, while expert ratings decline. The decline in expert ratings is likely due to news coverage of large cyberattacks in 2019 and 2020. In 2019, the Czech Foreign Ministry was reportedly attacked by a nation state adversary (Reuters, 2019), and in April 2020 two large hospitals reported attacks, shortly after a Czech government agency had warned of impending attacks against the country's critical infrastructure (Reuters, 2020). In France, the expert rating remains stable over the three years, while the observational indicator declines from 2018 to 2019 and then recovers. This does not seem to be driven by changes in available data – the number of observed IP addresses increases from 65 in 2018 over 144 in 2019 to 207 in 2020.

Given that Table 2.2 shows a positive association between the two indicators, the question arises whether the *changes* between the indicators are also linked. More specifically, does a change in the observational indicator predict a change in the expert-coded indica-

| | Change in V-Dem indicator | |
|---|---|---|
| | negative | positive |
| (Intercept) | −1.61 | −1.61 |
| | (1.10) | (1.10) |
| CVSS change (neg) | 0.48 | 0.39 |
| | (1.15) | (1.15) |
| CVSS change (pos) | 0.81 | 1.13 |
| | (1.14) | (1.14) |
| AIC | 250.63 | 250.63 |
| BIC | 267.46 | 267.46 |
| Log Likelihood | −119.32 | −119.32 |
| Deviance | 238.63 | 238.63 |
| Num. obs. | 122 | 122 |
| K | 3 | 3 |

$^{***}p < 0.001$; $^{**}p < 0.01$; $^{*}p < 0.05$

Table 2.4: Multinomial Regression for change in CVSS-based indicator and change in V-DEM indicator between years

tor? Table 2.4 presents a multinomial model with the V-Dem indicator as the dependent variable and “no change" as the baseline for the indicators. Contrary to the linear regression presented in Table 2.2 we see no statistically significant relationship – changes in the observational indicator do not predict changes in the V-Dem indicator.

The preceding analysis shows three things: First, that both indicators plausibly measure the same thing and are in relatively high agreement on the subject matter. Considering the very different data sources, this is encouraging for the use of both indicators. Second, that these indicators come with important caveats that researchers need to keep in mind and which point towards possible improvements. Third, that high-level concepts such as “government cybersecurity" can be investigated beyond aggregated expert opinions by relying on network measurement data.

The second and third point deserve further elaboration. For the indicator based on network measurements, and really all such indicators, it is clear that for actual use in research more data is required. For example, large changes in observed IP addresses might lead to very different scores. This is particularly problematic for historical data, where gathering the required data is now impossible. Going forward, however, consistently collecting the data showcased in this paper is feasible, as is broadening the covered countries.

The cases of Australia and the Czech Republic highlight shortcomings of both indicators. For the expert-coded indicator it seems reasonable to assume that the dip in ratings in Australia was due to the presumably state-sponsored attack against political parties in February 2019 (BBC News, 2019). But what does this rating say in absolute terms about Australia’s cyber-defense capabilities? It’s not unreasonable to think that the security vulnerabilities exploited in the attack were present in 2018 already given how long in advance state-actors in particular plan their attacks (Fireeye, 2020).

That would mean the indicator is lagging and ultimately bound to how highly publicised potential security breaches are – not how secure or vulnerable a country actually is. Similarly, the divergence of indicator values in the Czech Republic might be explained by highly reported cyberattacks, while the government invested in improving its defensive capabilities. Both cases then show that the stability or even improvement of the network-measurement based indicator is also not an absolute measure of security. Even low-rated security vulnerabilities can turn into a large data breach when used together, meaning that a country that would be generally rated as very secure can still fall victim to an attack. An indicator constructed from network-measurements as done in this example can thus only serve as a comparative measure between countries, not as an absolute indicator of security.

Nevertheless, using objective, observational data such as network measurements can be used to create indicators for high-level concepts such as a country's or governments "cybersecurity" and comes with advantages over expert-coded data. As discussed, such measurements are insensitive to changes in public perception due to media coverage, but they also allow higher resolution with regards to time (monthly indicators instead of yearly indicators, for example), scope (scores for individual departments) and fidelity (are eGovernment portals affected?), while also easier to collect after an initial setup due to automation.

## 2.4 Conclusion

In this article, I show that political science researchers can answer a range of questions going beyond the "visible" parts of the Internet by expanding data collection and incorporating network measurement data. Previous research has heavily relied on application layer data, limiting the scope of insights into how governments and other political actors use the Internet for their purposes. I provide an example application of network measurement data analysis to investigate government cybersecurity across countries. Since defensive government cybersecurity capability is a latent concept, researchers so far had to rely on aggregate measures based on expert interviews (Mechkova et al., 2019*a*). As the analysis shows, the observational indicator plausibly measure the same concept as the indicator based on expert surveys. I also show that expert surveys might be biased by media coverage of security breaches in a way observational indicators are not.

While the measure for government cybersecurity I construct in this paper provides a first step towards an observational indicator, it comes with some limitations. First, any indicator of government cybersecurity can only provide comparative measurements between countries, not absolute levels of defensive capabilities – this is true both for observational as well as expert survey-based indicators. Second, the data used to derive the observational indicator is limited in both historical depth as well as breadth of coverage. Future work should focus on continuous active measurement of vulnerabilities on

government servers to remedy this.

These limitations also highlight some practical challenges in using Internet measurement data in political science. The biggest challenge is the limited availability of historical data. Any research project will need to rely on third-party datasets like CommonCrawl and be subject to their constraints, unless researchers implement their own dedicated data collection efforts and maintain them over long periods of time. This in turn requires skills not traditionally taught in political science curricula, like programming, database- and system administration, and knowledge about the technologies that power the Internet. Additionally, researchers need access to technical infrastructure that enables storing and processing the large volumes of data they are likely to collect – collecting and analysing Internet measurement data is generally not feasible on personal computers but requires dedicated servers. Finally, researchers need to be aware of and navigate challenges in redistributing the Internet measurement data they collect. Copyright and data privacy laws as well as Terms-of-Service requirements can hamper the reproducibility of studies dealing with the *content* of websites, whereas studies dealing with more "technical" measurements are less likely to be subject to such constraints.

Where these practical challenges can be overcome, however, future research should look to further expand our understanding of the nexus of politics and the Internet beyond social media analysis. On a technical level, political scientists should seek to widen active measurement of political presence online, by collecting data across all network layers, from the IP addresses of servers that host websites of political actors, the programs that run on the servers, to the content of the websites themselves. Based on such continuous measurement, it will then be possible to carry out research into a wide array of topics, touching many subfields of political science: which cyberattacks, like Distributed Denial-of-Service attacks or website defacements, are used in a political context, and how? Who do governments and other political actors trust with their websites and data? Which departments have the biggest presence online, possibly indicating citizen demand for information? And how does the content of government websites change over time, reflecting different policy priorities?

# 3

# Censor or Block? Repertoires of Digital Censorship in Autocracies

*Co-authored with Nils B. Weidmann and Alberto Dainotti.*
*Currently under review at the Journal of Information Technology & Politics.*

## 3.1 Introduction

Without a doubt, the Internet is one of the technologies that has had the most profound impact on humanity in recent decades. The ability to communicate with many others around the globe, or to post messages that can instantly be seen my millions of people has changed the lives of many. Not surprisingly, much of political discussion and debate is also happening online, which is why research is increasingly focusing on the political repercussions of online news and social media. Early research along these lines has emphasized that online communication and social media may have "liberating effects" and empower oppressed groups in illiberal societies. More recently, however, scholars have examined the use of digital communication by governmental actors, in particular with regards to their desire to curtail freedom of expression and exert control over information. Governments across the world have increased their efforts to control information, often under the guise of preventing the spread of "fake news" or mis-information. Governments can use digital technology for several different purposes: to identify and track citizens

for the purpose of surveillance, to spread governmental propaganda online, but also to censor online content that is deemed problematic by the government. This latter aspect is the focus of this paper

Online censorship can take a variety of forms. The most drastic way to disable online communication is a complete Internet shutdown that stops all traffic in and out of a country (Dainotti et al., 2014; Gohdes, 2015). However, besides the extensive disruption this causes to all citizens, Internet shutdowns can have negative repercussions for a variety other outcomes, for example economic activity. This is why they are employed by governments only in particularly severe situations, for example when massive protests were expected in Cairo in January 2011 (Hassanpour, 2014). More common, every-day attempts to control information and the freedom of expression typically employ much less pervasive and more targeted measures.

In this paper, we compare two different types of censorship: Censorship through (1) the *blocking* of particular websites (Deibert et al., 2008; Filasto and Appelbaum, 2012), and (2) attempts to shutdown of particular servers through *cyberattacks* (Lutscher et al., 2020). For a government, the former is easier to implement through, e.g., filtering at the gateways in and out of a country. However, Internet filtering can be bypassed by citizens with simple technical tools (such as Virtual Private Networks, VPNs), which is why governments may prefer to disable the entire server that distributes the content to be censored. This can be done with the second type of censorship through Denial-of-service attacks, a brute-force type of cyberattack that requires little technical expertise to be carried out.

Existing research has produced insights into each of these types of censorship independently. As argued by Keremoglu and Weidmann (2020), this approach has been able to generate new and interesting results on the use and effects of particular censorship technologies. However, governments rarely ever rely on a single type of censorship alone; rather, they employ different ones in combination with each other. This *repertoire* of censorship technologies used by governments is something we know less about, and it is the main focus of this paper. We study how blocking and attacks are used in combination with each other, and how this relationship may be affected depending on the political situation on the ground. Because this has important consequences for the work of activists and dissidents in non-democratic regimes in particular, we focus our analysis on autocracies. Our analysis relies on new and refined data on the use of censorship in autocracies. We rely on Internet measurement techniques and large existing datasets to observe attacks and blockings and to compare them to each other. In the following, we briefly review the literature and present our theoretical argument, before describing these datasets and our analysis in detail.

## 3.2 Related Literature and Theoretical Argument

From the early days of the Internet in the 1990s until today, its meteoric rise in importance for everyday life as also been connected to hopes for political change (Barlow, 1996). Its potential as a "liberation technology" (Diamond, 2010) has been highlighted in connection to political upheaval and social movements around the world such as the "Arab Spring" in 2011 (Hassanpour, 2014; El-Baradei, 2011). Others, however, have cautioned against this idea and have argued that the Internet has the potential to further increase repression and censorship (Lessig, 1997; Morozov, 2011). These diverging predictions have resulted in a broad body of research investigating how governmental actors use the Internet and for what purposes. Beyond providing government services online and using the Internet to gauge citizens preferences, there is extensive evidence that governments also use the Internet for less benevolent reasons. This includes identifying and tracking citizens for surveillance, using the Internet to broadcast government propaganda at home and abroad and censoring online content that is critical of the government.

The usual assumption is that governments employ these strategies to secure their power. If governments do not live up to the promises they have made to win support, they may be held accountable and their support erodes. For this reason, governments try to shape and restrict what information reaches their citizenry to prevent negative outcomes for those in power (Roberts, 2018, p. 21f). While efforts to control information are on the rise in both democracies and autocracies (Sundara Raman et al., 2020), these incentives are still held in check in democracies by strong institutions and the notion that the right to communicate and organize freely is the very core of what constitutes a democracy (Merkel, 2004). In contrast, autocracies have weaker institutions, and the regime's entire survival often depends on keeping threats through public mobilization in check (Friedrich and Brzezinski, 1965; Wintrobe, 2000; Geddes et al., 2018).

The tools autocrats might use to censor online content can roughly be distinguished as belonging to one of three layers of Internet technology as Keremoglu and Weidmann (2020) describe. The first layer is the infrastructure layer, where governments control whether the Internet is accessible at all (Keremoglu and Weidmann, 2020, p. 3). Here, governments have been shown to interfere in access provision to politically excluded ethnic groups (Weidmann et al., 2016) or shut down access to the Internet for the whole country during times of contention (Belarus and Iran are recent examples, but also Egypt and Libya during the Arab Spring, see Dainotti et al., 2014). Second, where infrastructure is in place and accessible, the network layer allows governments to impose restrictions at a more granular level (Keremoglu and Weidmann, 2020, p. 3). Through tools based on keyword filters, governments can not only control access to information on a continuing basis (Hellmeier, 2016), but also carry out surveillance over their population and identify dissidents (Deibert et al., 2008). Where installing such censorship measures is too costly or citizens frequently circumvent the censorship by using tools such as VPNs, govern-

ments can use blunter tools like so-called denial-of-service (DoS) attacks to react quickly to emerging threats. DoS attacks overwhelm a website's server with network packets or data requests, in order to prevent citizens from reaching a website (Lutscher et al., 2020; Nazario, 2009). As a third and most researched layer, Keremoglu and Weidmann (2020) identify the application layer, which is where network applications such as browsers, e-mail programs or social media clients operate. Here, autocrats choose between four tactics to use the Internet to their advantage: explicit content controls, information manipulation, surveillance and the provision of "controlled venues of preference articulation" (Keremoglu and Weidmann, 2020, p. 5).

Each of the governmental tactics for information control is associated with its own set of costs. The most costly choice is shutting down the Internet completely within a country. This choice comes with severe repercussions such as reduced economic activity (Khrennikov, Kudrytski and Sazonov, 2020), and is therefore only employed in particularly severe situations, for example when massive protest were expected in Cairo in January 2011 (Hassanpour, 2014). Usually, censorship in autocratic countries is therefore done through less pervasive and more targeted measures. Two of these measures are (1) censorship through blocking of particular websites (Filasto and Appelbaum, 2012; Deibert et al., 2008) and (2) the targeted shutdown of servers hosting potentially objectionable content through cyberattacks. On a technical level, the former is achieved by installing tools in the infrastructure of Internet Service Providers in a country which inspect each website a user wants to visit and block access if the website is on a government-controlled blocklist. This is particularly effective at the gateways that connect one country to another since these gateways create centralized points of control that all traffic has to flow through. However, this type of filtering can be circumvented relatively easily by technologically proficient users through the use of simple technical tools like Virtual Private Networks. There is ample anecdotal evidence both for the use of blocking as a censorship measure (Okunoye et al., 2018; Deibert et al., 2008) and for the use of VPNs and other means to circumvent this blocking (Hobbs and Roberts, 2018).

While blocking is generally sufficient to dissuade the average citizen from accessing information critical of or dangerous for the regime (Roberts, 2018), the fact that politically engaged people can circumvent this measure relatively easily leads to the need for tools that can prevent access even to those people. In this case, governments can use targeted cyberattacks to take information entirely offline, instead of merely blocking access for their citizens. Denial-of-service attacks are a brute-force tool to take information offline by overwhelming the server with data requests (or other network traffic) through, e.g., the use of large collections of hacked computers and devices connected to the Internet called "botnets." Attacks via botnets can be purchased cheaply on the darkweb, and have the additional benefits that there is no visible indication of censorship to the end user and that the attacks themselves are not attributable to a specific actor. Despite this attribution problem, there is extensive anecdotal as well as some quantitative evidence

for the use of DoS attacks as a tool for censorship. Examples include DoS attacks against independent news websites before the 2011 election in Russia (Jagannathan, 2012) and in Turkey in 2015 (Kelly et al., 2017). These are not isolated events; a study by Lutscher et al. (2020) shows increased levels of DoS activity during politically contentious periods across a large sample of autocratic countries.

If governments want to censor online, should they resort to website blocking or DoS attacks? The two tactics have different characteristics and implications. For one, they vary in the extent to which there is a visible indication of censorship. Website blocking is much more obvious to the end user, and is oftentimes even clearly indicated with a message alerting the user that they intended to visit a forbidden website. A DoS attack, if successful, temporarily disables a server, which results in an error message that would be the same in the case of a technical malfunction that is not politically motivated. Also, the attribution of the censorship intervention is much more obvious in the case of website blocking, where the end user knows that blocking happened at the level of a regional or national ISP. A DoS attack, in contrast, usually cannot be attributed to a particular actor and offers plausible deniability because "patriotic hackers" can also carry out such attacks without direct links to the government. Finally, the two types of tactics differ with respect to their probability of success. Website blocking can be implemented such that ordinary users cannot circumvent it (while the more technically skilled can). At the same time, DoS attacks sometimes fail to reach their goal of disabling a server, but if they do, nobody has access. Overall, blocking is a relatively overt type of censorship, visibly preventing access to particular sites with guaranteed success for large, but not all, groups of users. DoS attacks is a covert strategy that has a certain probability of failure, but if successful, can prevent access from everyone.

The question of whether governments choose one tactic over the other is one about their censorship *repertoire*, i.e. the combination of different ways to interfere in online communication. Similar to repertoires of conventional repression (**?**), autocrats likely combine different online tactics depending on the current situation and on the goals they would like to achieve. For researchers, this means that studying one of these tactics alone can lead to misleading results: do we see low levels of website blocking, because (i) the government sees no need to censor, or because (ii) they have opted for other tactics, such as DoS? The focus on single tactics in empirical work is a common shortcoming of the literature on online censorship (Keremoglu and Weidmann, 2020), and we make a first step to remedy it in this article. In particular, there are two possible relationships between website blocking and DoS attacks. The first one is what we call "reinforcement", where governments combine different tactics. This is an approach of limited selectivity; if governments see the need to censor, they employ different tactics at the same time. If governments indeed use tactical reinforcement in online censorship, we should see that

the number of DoS attacks and the number of observed blockings in a country

are *positively* correlated (H1).

Alternatively, we can also expect that governments are more selective in their use of censorship tactics, carefully balancing their strengths and weaknesses, and employing the one that is considered to be the most useful. This is what we call "substitution," and it means that the use of one of them goes along with a reduced reliance on the other:

> The number of DoS attacks and the number of observed anomalies are *negatively* correlated (H2).

In addition, tactical reinforcement or substitution may not be constant over time. In fact, governments may choose to opt for an approach of sweeping censorship (reinforcement) at certain times, while favoring a more selective one (substitution) at other times. Therefore, in our empirical analysis, we study the relationship between different censorship tactics depending on the political situation on the ground, distinguishing between periods with mass protest and those without.

## 3.3 Data on Censorship Tactics

For our analysis, we require systematically measured data of two censorship tactics: website blocking and DoS attacks. In order to keep our measures of these tactics comparable, we keep the list of censorship targets constant, and then observe whether they were affected by either type of censorship. We rely on lists maintained by Citizenlab, which publishes categorized lists of potential censorship targets for 141 countries. For our analysis, we restrict potential targets to websites categorized as belonging to one of the following categories: religion, political criticism, human rights, militants or terrorism, news media, host and blogging platforms, or intergovernmental organizations[1]. For the websites in these categories, we measure both DoS attacks and website blocking. We describe the data gathering approach in the following sections.

### 3.3.1 Website Blocking

For website blocking, we rely on data gathered by the Open Observatory of Network Interference (OONI) web connectivity tests (OONI, 2020). OONI collects measurements of potential Internet censorship around the world through crowd-sourced network measurements. Residents of a country can download an app on their phone or computer, a so-called "OONI probe," and run network measurement tests. Because OONI measurements are only collected when a user decides to manually initiate such a measurement, measurements are not carried out at regular intervals, but depend on the number of active participants in each country.

[1]The full list of categories can be found here: `https://github.com/citizenlab/test-lists/blob/master/lists/00-LEGEND-new_category_codes.csv`

Each time a user initiates their probe, the program samples from the Citizenlab list of potentially censored websites for the respective country and performs automated visits to these websites (OONI, 2020). At the same time, automated visits are performed from OONI-owned infrastructure outside the user's country. The probe then records whether it received data for each website and sends these measurements as a "report" back to the OONI infrastructure. By comparing both measurements, OONI determines whether there is potential censorship occurring: if both results match, there is likely no censorship but if the user's results differ, the website being tested is likely censored.

On a technical level, the web connectivity test performs three checks for potential censorship (OONI, 2020). As a first step, the test checks whether a requested website is blocked via DNS tampering, i.e. the user's Internet service provider maps the request to a website to the wrong Internet Protocol (IP) address. Once the probe has received the IP address of the website, it will try to connect to that IP address through a TCP/IP request. Finally, the probe will send an HTTP GET request to the website, to which websites usually respond with their web content. If the HTTP request fails or the HTTP status codes do not match between the probe and the test run by the OONI infrastructure, this can indicate censorship. If any of the three checks fail, the report will record an "anomaly" for the particular website, the number of which is used in our statistical analysis below.

### 3.3.2 DoS Attacks

For DoS attacks, we contribute new data that allows us to compare the use of this tactic to blocking in a way that was previously not possible. In contrast to website blocking, which is amenable to active testing as carried out by OONI, DoS attacks cannot be observed by such active probing. This meant that previous studies had to rely on media reports (Asal et al., 2016; Jagannathan, 2012) even though media reports come with the significant risk of potential reporting bias: only successful and highly salient attacks are reported, and the targets of cyberattacks vary substantially in their salience for reporting, with attacks on human rights and other non-governmental organizations being under-reported (Hardy et al., 2014, p. 527).

To overcome reporting bias in the study of DoS attacks and allow detailed comparison to website blocking, we leverage passively collected data from the Center for Applied Internet Data Analysis (CAIDA, UC San Diego, 2019). This data permits a high-resolution perspective on denial-of-service attacks targeted at victims all around the world, irrespective of their salience and coverage in English speaking media. Through the UCSD network telescope, it is possible to capture one of the most frequently used DoS attack types, which are the so-called "randomly spoofed" attacks. Following the approach presented in Moore et al. (2006), it is possible to detect the IP addresses of attacked systems worldwide. Using this data, previous work by Lutscher et al. (2020) was able to show that DoS attacks are used for political purposes, with more DoS attacks recorded during

election periods in autocracies.

We develop a new measurement to better map Denial-of-service Attacks to possible targets to address the question who the actual targets of DoS attacks are. Because CAIDA can only capture the IP address that was attacked but not the intended host/website, it is not directly possible to identify attacks against websites on the Citizenlab list of hosts: we do not know which IP address a particular host was using at a particular point in time, since the mapping of host names to addresses via the Domain Name System (DNS) changes frequently. To solve this, we combine the data collected by CAIDA with historical data on websites and IP addresses collected by CommonCrawl (`https://commoncrawl.org`). CommonCrawl is a US-based non-governmental organization that crawls a large portion of the public Internet at four week intervals, scraping the content and IP addresses of the websites it comes across, and makes this content publicly available for research purposes. When a given host is visited by CommonCrawl multiple times and has the same IP address in two adjacent observations, we assume that the host had the same address on every day between those two observations. We call this period a "stable IP period," and only use DoS attacks that fall within a stable IP period when counting DoS attacks against the hosts in our sample[2].

Our new measurement improves on existing efforts to measure censorship in several ways. First, our measurement provides event-level data on the use of a specific censorship tactic. Existing research often relies on aggregate measures based on expert surveys, coded on the level of country-years (Coppedge et al., 2020). Second, our measurement is not biased by the salience of targets often present in news coverage (Hardy et al., 2014, p. 527) or constrained by language barriers in reporting. Finally, the global coverage of our measurement allows comparative research that goes beyond individual case studies.

## 3.4 Research Design

We set up regression models to test whether the presence of one type of censorship (blockings) is statistically related to the occurrence and magnitude of the other (attacks). It is important to mention here that these models do not serve to test a causal relationship; we do *not* claim that one type of censorship *causes* an increase (or decrease) in the use of the other. Rather, in line with our hypotheses, our regression models allows us to estimate partial correlations between the two types of censorship, removing country and time-specific trends, and allowing us to check whether correlations change depending on the political situation on the ground. Our analysis covers the period from 2015 to 2019, since this is the maximum period for which we have data from the OONI database that our analysis relies on. In line with our focus on autocracies, we include those countries with a Polity IV score of less than six in 2015, the first year of our analysis period (**?**). All countries are observed in weekly intervals. For each country, we include all years

[2]We provide a more detailed description of our method in the Appendix

for which we have at least three weeks of data available, both for OONI reports and CommonCrawl/DoS data.

In our main models, we use the number and occurrence (0/1) of DoS attacks during stable IP periods for the websites in the categories listed above. The DoS measurement is obtained by constructing stable IP periods for all websites in these categories. For each of these periods, we then record the number of DoS attacks against this IP address during that week. Since we can only observe potential attacks if we have at least one stable IP period in a country, we include only those country-weeks for which we have at least one stable IP period. In order to estimate correlations both for the number, but also the occurrence, of DoS attacks, we use estimate models with the (log-transformed) number of occurrences per week as well as models with a binary dependent variable ("any attack", with values of 0/1). Importantly, our measurement of attacks will be strongly dependent on how many stable IP periods we have – if we have few stable IP periods, the amount of attacks we can potentially pick up in our data will be small. The number of stable IP periods varies substantially between and within countries over time, with is why we include it as a control variable in our analysis.

To see how DoS attacks depend on the occurrence of website blockings, we include data on recorded anomalies as the independent variable. For this, we use all OONI Probe reports of web connectivity tests that fall into our analysis period. The reports are generated when a user of the OONI Probe software in a given country manually initiates a test, which means there is no consistent interval at which websites are tested. From these reports, we create a binary variable that takes the value of 1 if there is at least one "anomaly" in the reports for the selected website categories, which indicates potential censorship. Because we can only observe anomalies when users run the probe and a report is generated, country-weeks with no reports necessarily show no anomalies. To prevent this from biasing our results, we exclude country-weeks during which no reports were generated. As stated above, we compare periods of high political contention (protest) to those without, to see if there is a potential switch in the governmental strategy towards censorship. For this, we rely on protest data collected in the ICEWS dataset (Boschee et al., 2015). From ICEWS, we select all protest events directed at the government[3] to code whether a given country-week had any anti-regime protest.

For our analyses, we use a panel data approach. Because we are interested in substantive correlations within countries rather than across them, we include country fixed effects in our models. In addition, since the global level of DoS attacks varies between years, we additionally include year fixed effects to net out secular changes in the overall number of DoS attacks. We use OLS to estimate our models, which in the case of the binary dependent variables correspond to linear probability models (LPMs). To account

[3]We select events where the ICEWS variable 'target sectors' mentions at least one of the following: "Executive", "Executive Office", "Government", "Government Major Party (In Government)", "Ministry", "Legislative / Parliamentary", "Lower House", "Municipal", "Police", "Upper House", "Cabinet", "Elite", "Legislative / Parliamentary", "Army", "Military"

for serial correlation, we include attack presence as an additional predictor, lagged by one week.

## 3.5 Analysis

### 3.5.1 Case Illustrations

In Figure 3.1, we start with a descriptive look at four case examples for the use of website blocking and DoS Attacks. The top row displays the number of anomalies reported per week in each country (log-transformed), while the bottom row shows the number of DoS attacks per week, also log-transformed. The grey vertical lines indicate weeks in which the ICEWS data records at least one protest event. For all countries, we see substantial variation of anomalies and attacks over time that suggest long-term trade-offs as well as more tactical interactions. As a general pattern, we see that some countries rely predominantly on one tactic over another. Iran, Ukraine and Venezuela, for example, seem to shift most of their censorship towards blocking, beginning in 2017 with Iran. Malaysia, in contrast, seems to use both tactics consistently. At the same time, Malaysia in particular shows patterns consistent with short-term tactical choices. Both in 2017 and 2018, we see short drops in reported anomalies in Malaysia, followed by quick rises in such reports that coincide with an increase in DoS attacks. This pattern points towards tactical reinforcement and is repeated in early 2019: after a period of no recorded DoS attacks at the end of 2018, DoS attacks increase for a few weeks, adding to existing censorship efforts through blocking. Similar short term dynamics are visible in Iran in the summer of 2019 or Ukraine in the first half of the same year. Ukraine also shows indicators for possible tactical substitution as well, however. During a period of protest in the first half of 2017, we see no DoS attacks while reports of blocking go up. At the end of the year, a short spike in anomalies is followed by a short spike in DoS attacks, and followed again by a larger spike in reported anomalies in the first weeks of 2018.

These patterns provide anecdotal evidence for our hypotheses but do not yet allow us to draw any firm conclusions, in particular because Figure 3.1 also shows a generally high variance in the numbers of attacks and anomalies both within a country over time as well as between countries. We therefore proceed to a more systematic comparison by means of statistical analysis, which is able to separate out country-specific and time-specific levels of censorship.

### 3.5.2 Regression Analysis

Our main analysis is concerned with a possible tactical interaction between blocking and DoS Attacks. The results are shown in Table 3.1. Models 1-3 in the table use the log-transformed number of attacks as the dependent variable, while Models 4-6 use a binary variable coding the presence of any attack as their dependent variable. For each

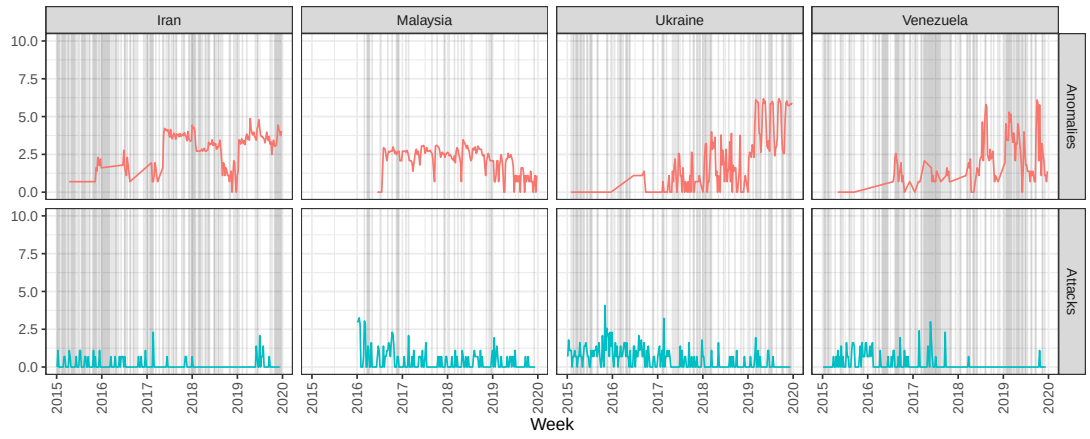


Figure 3.1: Blocking and DoS Attacks in Iran, Malaysia, Ukraine and Venezuela, ICEWS protest weeks in grey.

set of models, we proceed in three steps. First, we include a model that only contains the occurrence of anomalies as independent variable, as well as the standard controls (attacks in previous week, number of active SIPs). Second, we use the same model, but only include the occurrence of protest as the main independent variable. Third, to test whether the relationship between the two censorship tactics varies between times of protest and times of peace, we estimate an interaction effect between DoS attacks and the presence of protest.

In Models 1 and 4 in Table 3.1, we see no correlation between presence of anomalies and the number or occurrence of DoS attacks. Not surprisingly, attacks have a strong serial correlation, as the positive and significant coefficient of the lagged attack indicator shows. Similarly, Models 2 and 5 provide no evidence of a relationship between protest and attacks. Hence, there is no indication in our data that autocratic governments have a higher likelihood to resort to DoS attacks in times of political turmoil, as we may have expected based on similar work (Lutscher et al., 2020). However, the patterns change considerably once we interact the anomalies indicator with the protest variable. Here, we see clearly discernible correlations in Models 3 and 6: the protest indicator alone receives a positive coefficient, while the interaction is negative. This means that in a week with protest, the presence of an anomaly is negatively associated with the number of DoS attacks (Model 3). The same pattern holds when DoS attacks are coded as a binary variable (Model 6).

To make the interpretation of these results easier, we visualize the interaction models 3 and 6 by plotting the average marginal effects of our main variables of interest in Figure 3.2. In the figure, we plot the marginal effect of anomalies on the number of attacks (left panel) and on the probability of an attack (right panel). The figures show that when autocracies experience political contention (i.e. protest), observing any anomaly is associated with *fewer* DoS attacks, while the opposite is true in times without protest:

| | Num Attacks (Log) | | | Any Attack | | |
|---|---|---|---|---|---|---|
| | Model 1 | Model 2 | Model 3 | Model 4 | Model 5 | Model 6 |
| (Intercept) | $0.2626^{*}$ | $0.3185^{***}$ | $0.2305^{*}$ | $0.2806^{**}$ | $0.3213^{***}$ | $0.2493^{*}$ |
| | (0.1172) | (0.0832) | (0.1175) | (0.1061) | (0.0697) | (0.1063) |
| Any Anom | 0.0118 | | $0.0271^{*}$ | 0.0123 | | $0.0300^{*}$ |
| | (0.0124) | | (0.0133) | (0.0113) | | (0.0121) |
| Protest | | −0.0056 | $0.0790^{**}$ | | −0.0052 | $0.0784^{***}$ |
| | | (0.0101) | (0.0256) | | (0.0085) | (0.0232) |
| Any Anom x Protest | | | $-0.0909^{**}$ | | | $-0.1027^{***}$ |
| | | | (0.0282) | | | (0.0255) |
| Any Attack (prev. week) | $0.1755^{***}$ | $0.2454^{***}$ | $0.1733^{***}$ | $0.1392^{***}$ | $0.1784^{***}$ | $0.1367^{***}$ |
| | (0.0189) | (0.0143) | (0.0189) | (0.0172) | (0.0120) | (0.0171) |
| Active SIPs (log) | −0.0009 | −0.0176 | −0.0019 | 0.0038 | −0.0015 | 0.0029 |
| | (0.0258) | (0.0217) | (0.0258) | (0.0234) | (0.0182) | (0.0233) |
| Country Fixed Effect | Yes | Yes | Yes | Yes | Yes | Yes |
| Year Fixed Effect | Yes | Yes | Yes | Yes | Yes | Yes |
| $R^2$ | 0.2092 | 0.2195 | 0.2119 | 0.2302 | 0.2422 | 0.2340 |
| Adj. $R^2$ | 0.1968 | 0.2135 | 0.1989 | 0.2181 | 0.2364 | 0.2215 |
| Num. obs. | 3290 | 6708 | 3290 | 3290 | 6708 | 3290 |

$^{***}p < 0.001$; $^{**}p < 0.01$; $^{*}p < 0.05$

Table 3.1: Relationship between the presence of anomalies and number of DoS Attacks (logged), protest, and an interaction between anomaly presence with protest (ICEWS, Country/Week)

here, observing any anomaly is correlated with *more* DoS attacks. We emphasize that even though these plots are called marginal *effects* plots, we do not imply causality in this association – anomalies do not *cause* more or fewer DoS attacks. Rather, they reflect decisions about tactical choices that governments make when it comes to using different ways of information controls.

What can we learn about tactical choices from our analysis? First, tactical choices are not constant over time: in models without interaction effects, we see no consistent correlation between anomalies and attacks. Only if we include an indicator for political contention in our models as a moderator variable, we start to see statistically significant relationships between our variables of interest. Second, we do find evidence for both our hypotheses: in the absence of protest, observing any anomalies is associated with more DoS attacks, suggesting tactical reinforcement (H1). In other words, during "normal" times, autocrats do not seem to be very selective in their tactical choices. If they censor, they do so by different means, relying at the same time on blockings and attacks. During times of political contention, however, observing any anomalies is correlated with fewer DoS attacks, thus indicating tactical substitution (H2). In other words, when times politically threatening for autocratic governments, they become more selective in their censorship tactics and use one of them, but not both at the same time.

Overall, this finding is counter-intuitive: one would naively assume that autocrats "double down" during periods of contention and "ease up" on censorship when there is no protest. Clearly, autocrats do *not* generally lower censorship efforts during protest

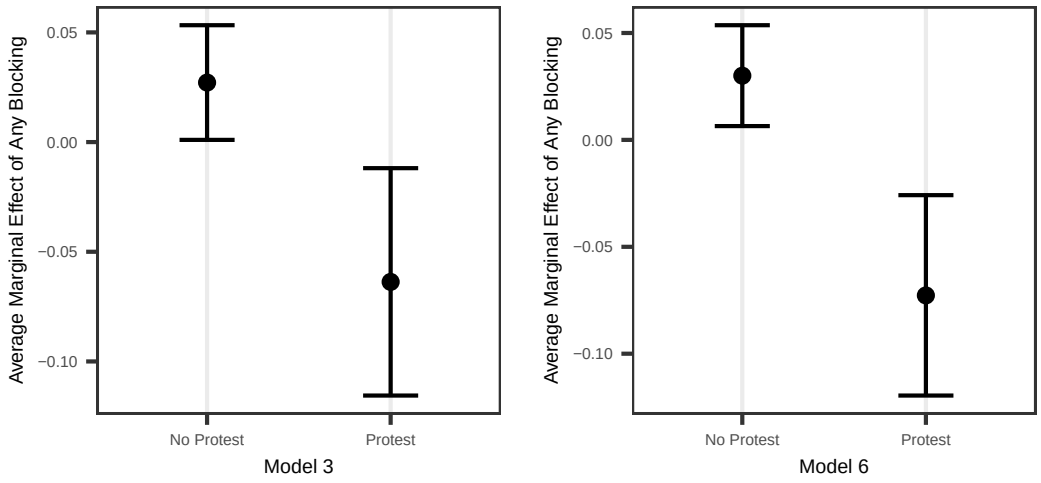


Figure 3.2: Marginal Effect of observing any anomaly on the logged number of DoS attacks (Model 3) and any DoS attacks (Model 6)

periods – if they did, we would have seen this in the models with protest (Models 2 and 5). Rather, they seem to be temporarily reducing the tactical diversity in their censorship repertoire. One possible reason for this finding might be the fear of a backlash if censorship becomes too pervasive during a period of high tension. While DoS attacks are impossible to detect from a user's perspective, the owners of a website immediately notice and have evidence for such attacks. Autocrats might therefore hold off on DoS attacks during times of protest, in order not to give dissidents evidence of repression that can be shared with foreign media and be used to increase foreign attention to the protests.

### 3.5.3 Robustness Tests

We conduct two tests to check the robustness of our results. First, it might be the case that our selection of autocratic countries influences the results. As recommended by Kasuya and Mori (2019), we conduct analyses using countries with a score below 0.42 of the indicator for electoral democracy in the V-Dem Project data (Coppedge et al., 2020) instead of the Polity IV score. These results are reported in Table B.1 in the appendix. Using a different selection method for autocratic countries leads to the same patterns as our main models. Second, we reverse the relationship between our variables, using *anomalies* as our dependent variable and DoS attacks as our independent variable. We report our results in Table B.2 in the appendix. Models A7-9 use the log-transformed number of anomalies as the dependent variable, while Models A10-12 use a binary variable coding the presence of any anomaly as their dependent variable. While missing statistical significance, the estimated coefficients display the same directions as

in the analysis above.

## 3.6 Conclusion

With this article, we have sought to investigate how autocratic governments use their *repertoire* of censorship techniques to control online communication and how this usage may be affected by the political situation on the ground. Previous work has investigated online censorship tactics independently, leading to a lack of insight into the interplay between these tactics (Keremoglu and Weidmann, 2020). Relying on Internet measurement techniques and large existing datasets, we provide first evidence that autocrats select tactics from their censorship repertoire depending on the current situation. In weeks with protest, observing the presence of website blocking is associated with *fewer* DoS attacks against opposition websites, while in weeks without protest it is correlated with *more* DoS attacks. This confirms our theoretical expectation that autocrats choose between tactical reinforcement and tactical substitution when deciding how to employ the tactics in their repertoire of techniques.

Our finding also raises new questions. That autocrats seem to decide *against* tactical reinforcement during times of protest is counter-intuitive. One possible explanation might be that while DoS attacks are a covert censorship technique from the perspective of potential website visitors, the owners of a website receive evidence for interference. Do autocrats take this into account and employ DoS attacks less so that evidence can not be shared with media outlets to increase attention to the situation? Is there evidence for tactical choices being influenced by media coverage? Further, autocrats might learn from the success and failures of other autocrats in deploying specific tactics. Are there patterns of tactics or combinations of tactics increasing and decreasing in popularity among autocrats over time?

Despite improving upon the existing literature, our analysis also comes with limitations that further work can address. First, our analysis still relies on historical data collected by third parties. This results in relatively sparse data, and future work should aim to introduce continuous and active measurement of politically relevant websites. In particular, our measurement of stable IP periods can be improved upon by recording the IP address of relevant websites more frequently going forward. Second, the data we use to detect website blocking relies on infrequent measurements by individual users in autocratic countries. Future work should find ways to detect of website blocking in a more automated way that allows better continuous monitoring and thus higher-resolution analyses. Third, the data we use to detect DoS attacks is based on only one—even if popular—way to carry DoS attacks. Further analysis would benefit from supplementing such dataset with data on attacks performed through other techniques, such as, e.g., amplification attacks. Finally, our analysis focuses on two censorship tactics while the repertoire of techniques available to autocrats contains other tactics as well. Frequent

active measurement of website *content* could also allow researching additional tactics, such as website defacements (replacing the content of a website with a message that the website has been hacked).

# 4

# Trust and Safety: Where do Governments Host Official Websites?

## 4.1 Introduction

The Internet has profoundly changed how economies and governments around the world operate. This includes how governments provide information to and interact with their citizens, with more and more government services shifting online. Early research into this nexus found that even 20 years ago, while the World Wide Web was still very new, 55% of all parliaments in the world had a website (Norris, 2001). Since then, government presence online has grown and governments increasingly rely on the World Wide Web to provide citizens easier access to government services and the Internet more generally to make administrative processes more efficient and less costly (Ni and Bretschneider, 2007). As the importance of the Internet has grown for governments, new challenges and considerations for their presence online have arisen. How governments deal with these new challenges is the focus of this paper.

The most important challenges for governments have been the growing importance of cybersecurity and questions of national data sovereignty. Recent incidents such as the SolarWinds attack (Fireeye, 2020) or the Office of Personnel Management data breach (Chaffetz, Meadows and Hurd, 2016) illustrate the consequences of cyberattacks on the national level. While governments use the Internet to provide public information to citizens, government servers connected to the Internet also store highly confidential data, such as personal identifiable information of citizens, personnel data, or provide commu-

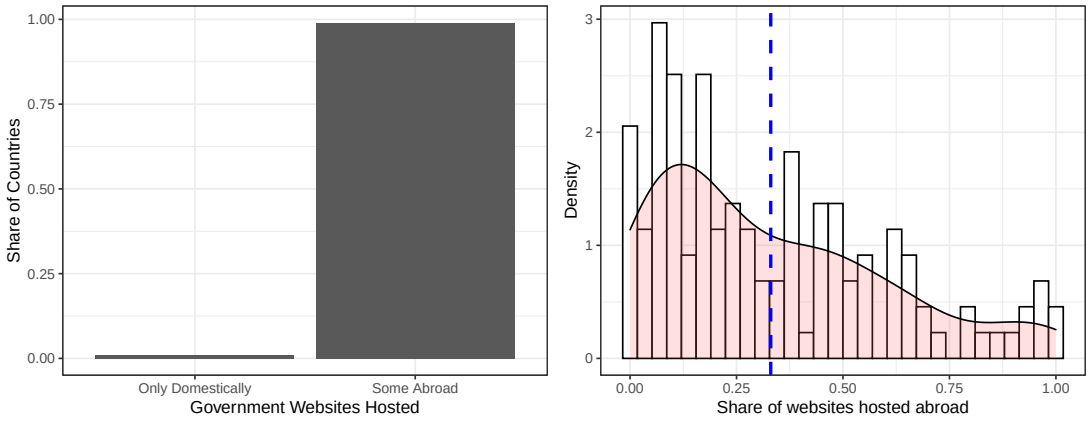


Figure 4.1: Distribution of hosting locations for government websites

nication services to government bodies like parliaments – all of which have been targeted by cyberattacks (Norris et al., 2019; McLaughlin and Dorfman, 2019; Von Der Burchard, 2020). While cyberattacks are an immediate and obvious threat to government data, governments increasingly also consider where their data is hosted, and by whom. Previous research has argued that this emphasis is driven by both economic factors, i.e. the desire to boost their own Internet data hosting industry, as well as political factors such as the desire to reduce exposure to foreign countries through signals intelligence (Chander and Lê, 2015) and considerations of national data sovereignty (Irion, 2012)

In this paper, I investigate a puzzling observation given the importance of cybersecurity and national data sovereignty for governments. Despite the need to secure their websites from cybercriminals, activists, and interference from strategic competitors, the vast majority of countries (94%) host at least some of their government websites with businesses they do not have jurisdiction over, i.e. business incorporated abroad and controlled by foreign entities (see Figure 4.1). Together with the substantial variance in *how many* of their government websites countries host abroad, this raises a question: what influences which foreign countries a government chooses to host official websites? Governments have generally strong incentives to host their websites domestically for both security as well as economic reasons, i.e. to boost their local IT sector. I argue that this choice depends on how how trustworthy the "hosting" country is seen. Using a novel dataset of hosting location choices for 123 countries, I find that countries generally host higher shares of their abroad-hosted websites with countries they share military alliances with or that are more democratic than they are.

## 4.2 Related Literature and Theoretical Argument

As the importance of the Internet has grown for economies generally, governments have increasingly moved parts of their communication infrastructure and citizen-facing ser-

vices online. These efforts have been motivated by desires to make administrative processes more efficient and less costly and to expand access to services for citizens (Ni and Bretschneider, 2007). Where websites 20 years ago were mainly static representations of content, advancing e-Government initiatives mean that today they often are portals for citizens to directly interact with government services. At the same time, this shift has come with an increase in vulnerabilities for governments falling into two distinct threat models.

First, both political activists and criminals are motivated to attack government websites. For activists, government websites are interesting targets because taking these websites down or "defacing" them, i.e. replacing a website with a political message, allows to reach every normal visitor of that government website. Examples include the defacement of the main e-Government portal in the Philippines to protest against the worsening human rights situation under President Duterte (Neil, 2021) or in Myanmar as protest against the military coup (Cimpanu, 2021). Additionally, government websites are often portals to sensitive administrative data. In 2020, for example, activists hacked 251 law enforcement websites in the United States, releasing personal data of more than half a million police officers in the wake of the police murder of George Floyd (Lee, 2020). Criminals are equally interested in hacking government websites, since they often store personal identifiable data such as social security numbers of citizens that can be sold on the black market (Norris et al., 2019).

Second, government websites are targets in international conflicts as well. Whether or not outright "cyberwar" is a realistic threat (Rid, 2012), breaches of the CIA's system for communicating with covert sources (Dorfman and McLaughlin, 2018) or the Office for Personnel Management (McLaughlin and Dorfman, 2019) show that even government systems storing highly classified data are targeted and often vulnerable to online hacks from adversarial nations. Additionally, incidents like the Snowden leaks have demonstrated that countries surveil online data flows, putting government data at risk (Selby, 2017). Importantly, adversarial nations have much larger capabilities than activists or criminals. Where the latter are likely to focus their efforts on software vulnerabilities, nation states can also leverage their jurisdiction over businesses for their purposes.

The second threat model in particular has important implications. Morrison (2013), for example, notes the particular relevance of securing control over the entire supply chain of technology providers enabling government use of digital technologies. This supply chain starts at the hardware level, with governments even considering in which country IT equipment is being produced out of a fear that adversarial nations can covertly introduce backdoors (Morrison, 2013, p. 760f). A recent example for this fear is the ban of Huawei, a Chinese IT equipment provider, in the United States and efforts to convince US allies to enact similar bans (Whalen, 2021). The same considerations extend from the hardware up to the software level, with the US government banning software from the Russian cybersecurity firm Kaspersky Labs after allegations the business worked with

the Russian Federal Security Service (Lubold, 2017).

Supply chain security considerations thus force governments to choose whether to acquire and run the technical infrastructure, such as data centers, themselves or whether to contract their operation out to private businesses. Each choice comes with distinct trade-offs. Building and operating data centers is costly, but reduces the number of parties involved in the supply chain. Contracting out is potentially cheaper, but adds parties to the supply chain that need to be trusted.

The desire to protect a countries national data sovereignty (Irion, 2012) has lead countries to build government-controlled data centers and enact data localization laws that require data to be stored domestically. In 2021, Senegal announced a new government-owned data center to increase digital data sovereignty and become less reliant on US and other providers (Swinhoe, 2021*b*), with similar efforts reported from Zimbabwe (Swinhoe, 2021*a*). With regards to data localization, Nigeria requires that all government data is hosted in data centers located within Nigeria (Bowman, 2015), as does India (PTI, 2017). Countries like Brazil, France, Russia, Malaysia and others have enacted laws in a similar spirit (Bowman, 2015; Chander and Lê, 2015; Ewing, 2018). In addition to securing data from signal intelligence of other countries (Selby, 2017) and tighter control over general security, some governments and local firms also see such laws as opportunities to strengthen the local IT sector (Chander and Lê, 2015; Dawn-Hiscox, 2018).

At the same time, previous research on governments' decision to contract out services has shown that cost savings are an important factor in government choices on whether to operate their digital infrastructure themselves. As Ni and Bretschneider (2007) lay out in their framework for government contracting decisions, governments face pressure to provide services at lower costs, which includes e-Government and government IT services. One example is the "Data Center Optimization Initiative", part of the "Cloud Smart" strategy of the US government, which has the explicit goal of reducing the number of government-run data centers in efforts to reduce costs (Office of Management and Budget, 2016, 2019) and resulted in over 200 closed data centers as of August 2021 (Data Center Optimization Initiative, 2021). Instead of operating data centers themselves, governments then rely on businesses offering cloud services, sometimes also called "Infrastructure as a Service".

When governments make the decision to contract out the operation of their technical infrastructure, they are presented with the choice of contracting domestic businesses or international firms. Contracting to foreign businesses is in direct opposition to national data sovereignty and supply chain concerns, something governments are keenly aware of. Israel, for example, required Amazon and Google in a recent contract to "set up local Israeli companies that will be in charge of building and then operating the data centers, and that will be subject to local laws" in order to guarantee continued service even if the companies face political pressure to boycott Israel (Solomon, 2021).

Despite such concerns, governments choose foreign businesses to host their digital in-

frastructure, citing cost concerns and faster service procurement (Hanada and Tobita, 2020; Jones, 2012). *Which* countries a government chooses to host official websites remains an open question, however.

### 4.2.1 Hosting Location Choice and Bilateral Trust

The central problem for governments deciding to rely on foreign businesses for their digital infrastructure is trust. Given the highly sensitive nature of government digital infrastructure, the "procuring" government must be able to trust that the government in which the "hosting" business resides does not use its jurisdiction over the business to exploit the "procuring" government[1]. Because the direct parties are a government and a business instead of the usually assumed two governments, this presents in essence a collaboration problem *by proxy*. The field of International Cooperation Theory (ICT) provides a rich literature investigating under what circumstances countries cooperate with each other (for an overview see Dai, Snidal and Sampson, 2010). As Dai, Snidal and Sampson (2010) point out, ICT has paid particular attention to two mechanisms for facilitating cooperation: reciprocity and reputation. If the tit-for-tat incentives for cooperation of a repeated interaction are not present or immediate enough, reputation can serve as a stand-in (Dai, Snidal and Sampson, 2010; Keohane, 1986). I contend that this logic extends to the case of digital infrastructure procurement, even though the actors are not two governments like ICT usually assumes.

One way to operationalize reputation is by modeling trust. Different approaches to modeling trust exist in the literature, ranging from building trust through costly signals (Kydd, 2007) to distinguishing between trust and trusting relationships (Hoffman, 2002). For the present case, however, a simpler operationalization of trust seems adequate: membership in military alliances. It stands to reason that countries who commit to mutual defense and cooperate by exchanging defense technology have a high level of trust, sufficient enough to permit their government websites to be hosted in countries that are also members in the alliance. If this indeed the case, we would expect that

> the share of all websites hosted abroad that is hosted in a given country is positively correlated with mutual alliance membership (H1)

A more general way to think about trust between countries is discussed in the literatures touching the democratic credibility advantage. In particular research on international commitments has investigated which regime characteristics might influence possible treaty violations and whether countries reneg on their alliance promises. In this

[1]Note that this does not require perfect control over the "hosting" business. Export restrictions that cover digital services (Motamedi, 2019), or legislation such as the "Telecommunications and Other Legislation Amendment (Assistance and Access) Act" that grants the Australian government the power to compel businesses to provide access to any kind of digital information if in the government's interest (Stilgherrian, 2018) are enough to present a danger of exploitation.

context, Leeds (2003) and Leeds, Mattes and Vogel (2009) show that democracies are more likely to honor international alliance commitments, Von Stein (2005) shows that democracies are more likely to honor human rights treaties and Mattes and Rodríguez (2014) find that autocracies that are more similar to democracies in their institutions are more likely to cooperate with democracies. Extending this reasoning to the hosting choice of government websites, countries might generally be interested in hosting their websites in countries that are more democratic than they are themselves. If this is true, we would expect that

> the share of all websites hosted abroad that is hosted in a given country is positively correlated with whether this country is more democratic (H2)

## 4.3 New Data on Government Website Hosting Locations

For my analysis I rely on new data on the hosting decisions of 128 governments during the period 2014-2018. In order to collect this data, I first create a list of government domains. Because there is no global list of government domains, I identify government domains by their URL suffix. Most countries use dedicated URL suffixes like `.gov`, `.gov.mx` or `.gouv.fr` to make it easy for citizens to distinguish legitimate government websites from others on the Web. By manually coding the Public Suffix List maintained by the Mozilla Foundation (2020), I am able to identify government websites of 128 countries based on their URL suffix[2].

Based on this list of country-suffix pairs, I collect data on individual government websites for each country. For this purpose I leverage the CommonCrawl archive. CommonCrawl is a non-profit organization that regularly collects data from a large portion of the public Internet and makes it available for public use. Every four weeks, CommonCrawl collects a sample of over 2.5 billion webpages across 100 million domains (CommonCrawl, 2019). The generated archives contain the content of the visited websites and, among other things, the IP address from which the sites were served. Because CommonCrawl started its data collection in 2014, it is possible to observe which website was hosted on which IP-Address over time.

While CommonCrawl records which website was served from which IP-Address, it does not record who *owned* the IP-Address at that point in time. On the Internet, IP-Addresses are "owned" by so-called "Autonomous Systems", sub-networks owned by government departments, Internet Service Providers, and other businesses which are connected to form the Internet as a whole. To connect historical observations of IP addresses to the businesses that owned them at the time, I rely on data collected by the Center of Applied Internet Data Analysis (CAIDA) at the University of California, San Diego

[2]Systemic factors do not seem to influence whether countries use suffixes to distinguish government websites.

| Hostname | Country | ASN Country |
|---|---|---|
| acarape.ce.gov.br | Brazil | Brazil |
| acarape.ce.gov.br | Brazil | United States |
| altinopolis.sp.gov.br | Brazil | Brazil |
| aragoiania.go.gov.br | Brazil | Brazil |
| aragoiania.go.gov.br | Brazil | United States |
| arroiodosratos.rs.gov.br | Brazil | United States |

Table 4.1: Sample URL Host, Country and AS Country Mapping

(CAIDA, 2020*a*,*b*). The "Prefix to AS Mapping" dataset contains historical data on which IP-Address was owned by which Autonomous System (CAIDA, 2020*b*), while the "AS to Organization Mapping" dataset contains historical data on which Autonomous System was owned by which business, and the business' country of registration (CAIDA, 2020*a*). Combining the IP-Address observations from CommonCrawl with the data provided by CAIDA, I am able to generate a dataset that contains for every observation of a government website which business or other entity owned the IP from which the website was served, and which country had jurisdiction over said business or entity. Table 4.1 shows a simplified example of the resulting data.

By comparing the origin of a given website using its Top Level Domain (such as `.mx` for Mexico or `.fr` for France) to the country of the Autonomous System, I create a dataset which contains information on how many websites from country A (the origin country) were hosted in country B (the destination country) in any given year. This dataset also allows me to record what percentage of a country's government websites are hosted domestically and abroad, and how large the share of websites hosted abroad is that each destination country receives. While origin countries are restricted to countries that use government websites identifiable through their domain suffix (`.gob.mx`, `.gouv.fr`), the destination country of the directed dyads can be any politically relevant country as coded by Gleditsch and Ward (1999). For example, Germany does not use a variant of `.gov` to distinguish government websites, but hosts the government websites of multiple other countries.

## 4.4 Research Design

To answer the question whether countries host higher shares of their abroad-hosted websites with countries they trust, I create a directed dyadic panel dataset. This means that for each country pair, two observations are recorded per year. The first observation records the share of abroad-hosted websites that country A hosts in country B, while the second observation records the share of abroad-hosted websites that country B hosts in country A. Each observation also contains indicators of trust between the two countries. For the explicit modeling of trust per the first hypothesis, each observations records

whether both countries were members of a military alliance with obligations to defend each other in 2012, taken from the dataset published by Gibler (2008). For the modeling of trust through the democratic credibility advantage, I rely on the electoral democracy index provided by the V-DEM project (Coppedge et al., 2019). This index describes how competitive and free of systemic irregularities elections in a country are, whether there is an independent media and whether political and civil society organizations can operate freely. Each observation records whether the hosting country is more democratic than the country that owns the website by substracting the electoral democracy index score of the country that owns the website from the index score of the hosting country. If this variable is positive, the websites are hosted in a country that is more democratic than the owner, if it is negative, the hosting country is less democratic than the owner.

To test the hypotheses described in the theoretical section, I use OLS to estimate the statistical models. The models for both hypotheses use the share of abroad-hosted government websites hosted in a particular country as the dependent variable. The main independent variables are mutual alliance membership and relative democracy status. The standard errors for these models are clustered at the country level. Because the share of abroad-hosted websites hosted in a particular country should not be influenced by broad technological trends, these models do not include year-fixed effects. I also exclude the United States as a hosting location given its outlier status with regards to Internet technology and particularly hosting providers such as Amazon, Microsoft and Google. As control variables I include control variables for (logged) population size (The World Bank, 2019*b*), per-capita GDP (Coppedge et al., 2019) and Internet penetration (The World Bank, 2019*a*) and add whether the countries share a contiguous border (Weidmann and Gleditsch, 2010) to control for spatial effects. All variables are recorded both for the government country as well as for the hosting/ASN country. Because the number of observed government websites for each country is not equally distributed across years, I exclude country-years with less than 10 observations[3].

## 4.5 Results

In this section, I examine how trust between two countries is related to the share of abroad-hosted government websites hosted in each.

Table 4.2 reports the results of the analyses for the relationship between trust and hosting decisions. Figures 4.2 and 4.3 show the coefficient plots for both models. Because the effect sizes are very small, the dependent variable, the share of abroad-hosted government websites, is scaled by a factor of 100. In line with the second hypothesis, shared alliance membership is positively correlated to the share of abroad-hosted government websites hosted in a given country. This also holds when trust is modeled as the relative democracy status: Model 2 in Table 4.2 shows this variable is positively

[3]This is a common procedure in computer science to avoid false positives.

| | Model 1 | Model 2 |
|---|---|---|
| (Intercept) | $-0.92^{***}$ | $-1.06^{***}$ |
| | (0.21) | (0.21) |
| Shared Alliance Membership | $0.49^{***}$ | |
| | (0.03) | |
| Host More Democratic | | $0.13^{***}$ |
| | | (0.02) |
| GDPpc (Gov, log) | −0.02 | −0.02 |
| | (0.02) | (0.02) |
| GDPpc (Host, log) | −0.00 | 0.00 |
| | (0.02) | (0.02) |
| Cont. Border | $1.13^{***}$ | $1.33^{***}$ |
| | (0.06) | (0.06) |
| Total Pop. (Gov, log) | −0.00 | −0.01 |
| | (0.01) | (0.01) |
| Total Pop. (Host, log) | $0.06^{***}$ | $0.07^{***}$ |
| | (0.01) | (0.01) |
| Int. Pen. (Gov) | 0.00 | $0.00^{*}$ |
| | (0.00) | (0.00) |
| Int. Pen. (Host) | $0.00^{***}$ | $0.00^{***}$ |
| | (0.00) | (0.00) |
| $R^2$ | 0.02 | 0.01 |
| Adj. $R^2$ | 0.02 | 0.01 |
| Num. obs. | 59439 | 59439 |

$^{***}p < 0.001$; $^{**}p < 0.01$; $^{*}p < 0.05$

Table 4.2: Relationship between share of all abroad-hosted government websites and indicators for trust

correlated to the share of abroad-hosted government websites in a given country. This means that countries host higher shares of their abroad-hosted websites with countries that have higher democracy scores than they have themselves.

### 4.5.1 Robustness Tests

One potential challenge for the results regarding trust between countries is that countries in the European Union have advanced technology sectors, are part of a single market and most member states are also members of NATO. To prevent potential bias from this, I conduct the same analysis as in Models 2 and 3 and include a control variable that codes whether the hosting country is a member of the EU. Table 4.3 presents the results of this analysis. Including this control does not change our main results – shared alliance membership is positively correlated to the share of abroad-hosted government websites hosted in a given country, as is a higher democracy score value of the hosting country.

A second potential challenge for the results regarding trust between countries is that autocracies might be less inclined to host official websites with more democratic countries out of a fear of sanctions. To remedy this potential source of bias, I include an interaction effect between a control variable that codes whether the country owning the websites is an autocracy and the main independent variables. In my coding of this control variable

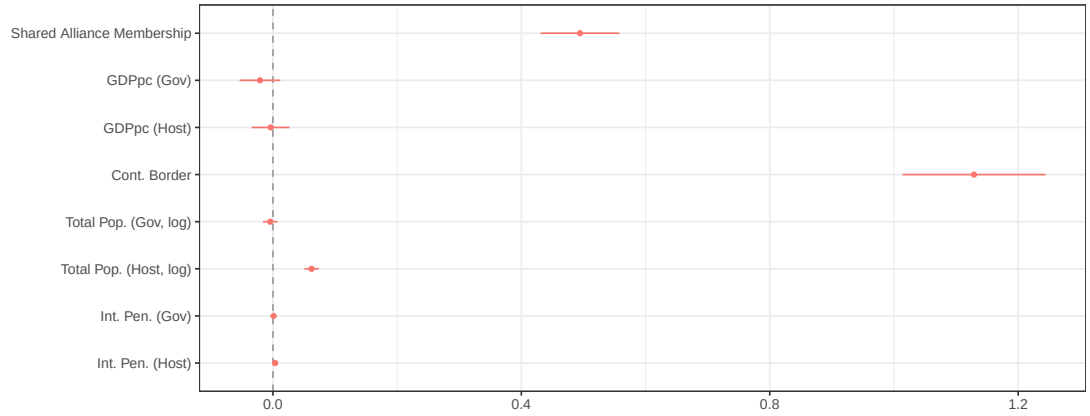


Figure 4.2: Coefficient Plot for Model 1

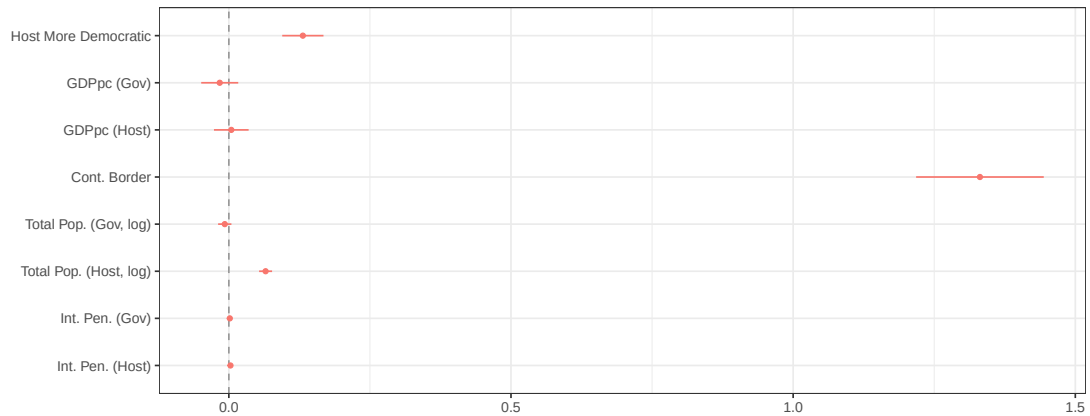


Figure 4.3: Coefficient Plot for Model 2

| | Model 3 | Model 4 |
|---|---|---|
| (Intercept) | $-0.91^{***}$ | $-1.08^{***}$ |
| | (0.21) | (0.21) |
| Shared Alliance Membership | $0.50^{***}$ | |
| | (0.03) | |
| Host More Democratic | | $0.14^{***}$ |
| | | (0.02) |
| Host EU Country | 0.02 | $-0.04$ |
| | (0.03) | (0.03) |
| GDPpc (Gov, log) | $-0.02$ | $-0.02$ |
| | (0.02) | (0.02) |
| GDPpc (Host, log) | $-0.00$ | 0.01 |
| | (0.02) | (0.02) |
| Cont. Border | $1.13^{***}$ | $1.33^{***}$ |
| | (0.06) | (0.06) |
| Total Pop. (Gov, log) | $-0.00$ | $-0.01$ |
| | (0.01) | (0.01) |
| Total Pop. (Host, log) | $0.06^{***}$ | $0.07^{***}$ |
| | (0.01) | (0.01) |
| Int. Pen. (Gov) | 0.00 | $0.00^{*}$ |
| | (0.00) | (0.00) |
| Int. Pen. (Host) | $0.00^{***}$ | $0.00^{***}$ |
| | (0.00) | (0.00) |
| $R^2$ | 0.02 | 0.01 |
| Adj. $R^2$ | 0.02 | 0.01 |
| Num. obs. | 59439 | 59439 |

$^{***}p < 0.001$; $^{**}p < 0.01$; $^{*}p < 0.05$

Table 4.3: Relationship between share of all abroad-hosted government websites and indicators for trust, controlling for EU countries

| | Model 5 | Model 6 |
|---|---|---|
| (Intercept) | $-0.88^{***}$ | $-1.08^{***}$ |
| | (0.21) | (0.21) |
| Shared Alliance Membership | $0.43^{***}$ | |
| | (0.04) | |
| Autocracy | 0.03 | $-0.08^{*}$ |
| | (0.02) | (0.04) |
| Host More Democratic | | $0.12^{***}$ |
| | | (0.02) |
| Shared Alliance x Gov Autocracy | $0.22^{**}$ | |
| | (0.07) | |
| Host More Democratic x Gov Autocracy | | 0.08 |
| | | (0.04) |
| GDPpc (Gov, log) | $-0.03$ | $-0.02$ |
| | (0.02) | (0.02) |
| GDPpc (Host, log) | $-0.00$ | 0.01 |
| | (0.02) | (0.02) |
| Cont. Border | $1.12^{***}$ | $1.34^{***}$ |
| | (0.06) | (0.06) |
| Total Pop. (Gov, log) | $-0.01$ | $-0.01$ |
| | (0.01) | (0.01) |
| Total Pop. (Host, log) | $0.06^{***}$ | $0.07^{***}$ |
| | (0.01) | (0.01) |
| Int. Pen. (Gov) | $0.00^{*}$ | $0.00^{*}$ |
| | (0.00) | (0.00) |
| Int. Pen. (Host) | $0.00^{***}$ | $0.00^{***}$ |
| | (0.00) | (0.00) |
| $R^2$ | 0.02 | 0.01 |
| Adj. $R^2$ | 0.02 | 0.01 |
| Num. obs. | 59439 | 59439 |

$^{***}p < 0.001$; $^{**}p < 0.01$; $^{*}p < 0.05$

Table 4.4: Relationship between share of all abroad-hosted government websites and indicators for trust, with interaction effect of autocracy

I follow the recommendation by Kasuya and Mori (2019) and code countries with a score below 0.42 of the indicator for electoral democracy in the V-Dem Project data as autocracies. Table 4.4 shows the results of this analysis. Again, including additional controls does not change the substantial results. The interaction effect between shared alliance membership and autocracy status of the owning country does, however, show a positive association with the share of abroad-hosted websites hosted in a given country. This means that autocracies host higher shares of their abroad-hosted websites with members of the same alliance. Considering the interaction between the owning country's autocracy status and whether the hosting country is more democratic presented in Model 6, we see no significant association. in this model, autocracies are, however, generally hosting a lower share of their abroad-hosted websites in a given country compared to countries that are not autocracies.

## 4.6 Conclusion

With this article, I have sought to investigate a puzzle presented by the observation that the vast majority of countries host government websites with businesses outside their own jurisdiction even though considerations of cybersecurity and national data sovereignty should make this unlikely. Relying on new data about the hosting locations of government websites, I show that in cases where governments choose to host their websites abroad, they do so with countries they trust. This holds for two operationalizations of trust, either as shared alliance membership or as relative democracy status.

While this paper contributes to our understanding of how governments navigate the challenges presented by the importance of the Internet coupled with its inherently networked international nature, it comes with limitations. In particular, the data on government websites is limited in its historical depth as well as in the coverage of countries. Future work should emphasize continuous data collection and increasing the number of covered countries beyond those with government-specific URL suffixes.

This paper also points towards new questions at the nexus of international relations and the Internet. While the analysis presented here focused on the technical layer of website hosting, the *content* of these websites has not been explore at all in the literature. How do governments link to international organisations like the United Nations and the EU, and how does this change over time? Are government positions to international treaties or other issues reflected on their websites? How accessible are government websites to citizens from other countries?

# 5 Conclusion

Over the last thirty years, the Internet has become the main venue for political interaction. This has meant new challenges and opportunities for governments, touching questions of sovereignty and security, information control and international trust. In my dissertation I have provided comparative studies on three different facets of the nexus of governments and the Internet, and show how social scientists can leverage Internet measurement techniques to advance our understanding of the government-Internet nexus. In this chapter I discuss the contributions of my dissertation and point towards avenues for future research.

## 5.1 Contributions

My dissertation provides three main contributions to the literature. First, I address a gap in the literature regarding political behavior on the network layer of the Internet. Second, I show how using Internet measurement data can provide a way to consistently consider the underlying technology of the Internet for political science research questions. Third, I contribute *comparative* studies to address the dominance of case studies or within-country comparisons.

### 5.1.1 Political Behavior on the Network Layer

Most political science research has focused on the "visible" part of the Internet where users engage with each other, the *application layer* (Keremoglu and Weidmann, 2020). This has limited the range of questions researchers have asked and meant that there

are many facets of how political actors, and governments in particular, are navigating the digital revolution that we know little about. We still lack understanding in many questions that go beyond what is immediately visible *on* the Internet: who do political actors, and governments in particular, trust with providing their digital infrastructure and hosting their websites? How secure are these websites comparatively? How do governments choose between different technical measures of information control? In short, how does political behavior play out on the network layer? In my dissertation, I remedy this shortcoming by providing three empirical studies that speak to different fields. As comparative studies, they also address the dominance of case studies or within-country comparisons. In the following I summarize the results of the studies and discuss their respective contributions.

My first paper contributes to the literature on cybersecurity. Previous work in this literature has begun to empirically investigate how secure government websites are and built theory around the risks associated with relying on digital systems. However, empirical studies have been constrained to within-country comparisons of cybersecurity or comparing small numbers of countries with regards to their practices. Relying on Internet measurement data of security vulnerabilities found on servers that host government websites, I analyse government cybersecurity across countries. I use the data on security vulnerabilities to construct a new observational indicator of defensive government cybersecurity capability and compare it to an indicator based on expert surveys. My analysis shows that the observational indicator plausibly measures the same concept as the indicator based on expert surveys and that expert surveys might be biased by media coverage of security breaches in a way observational indicators are not. The main contribution of this paper to the literature on government cybersecurity lies in providing a first observational indicator of defensive cybersecurity capabilities across countries and time.

In the second paper, we contribute to the literature on information control by providing insight into the interplay between online censorship tactics. Previous work has considered tactics in isolation, which others have observed to be a common shortcoming of the literature on online censorship (Keremoglu and Weidmann, 2020) and we argue may lead to misleading results: do we see low levels of a tactic because the government sees no need to censor, or because they have opted for another tactic? In our paper we make a first step to remedy this shortcoming by studying the interplay between two online censorship tactics (website blocking and DoS attacks) in autocracies. Relying on Internet measurement data we construct a new measurement that lets us identify the victims of DoS attacks among potential censorship targets in autocracies and compare attacks against these targets with blocking of their websites. Our results show that autocrats select tactics from their censorship repertoire depending on the current situation. In weeks with protest, observing the presence of website blocking is associated with *fewer* DoS attacks against opposition websites, while in weeks without protest it is correlated

with *more* DoS attacks.

With my third paper I contribute to the international relations literature by empirically investigating where governments host their official websites. Previous research suggests that countries have strong incentives to host official government websites and infrastructure domestically. Yet despite considerations of supply chain risks to cybersecurity and the motivation to host their websites domestically to secure their national data sovereignty, governments still host official websites abroad. This raises the question which factors might influence a government's choice to host in a particular country? I argue that this choice depends on how how trustworthy the "hosting" country is seen. By using Internet measurement data on the hosting location choice of governments, I show that governments choose to host their official websites with countries they trust and that this holds for two operationalizations of trust, either as shared alliance membership or as relative democracy status.

### 5.1.2 Internet Measurement Data for Political Science

Previous work across different subfields investigating the nexus of governments and the Internet has varied widely in the (empirical) consideration of the technology underlying the Internet. As a technological network of networks, the Internet relies on network members to expose certain technical information to the world to facilitate communication. This data is the "ground truth" of how participants in the network behave, whether individual users, businesses, or governments. Crucially, this data can be observed and collected by anyone willing to expend the necessary resources. This creates significant opportunities for analysis, but also requires a certain technical level of understanding of how the Internet works. The latter point is likely the reason application layer data has so far dominated empirical studies of digital politics: social media networks, for example, are highly visible to researchers, in contrast to the technologies that make the Internet itself work. Additionally, accessing data from social media networks is often relatively easy because they offer public endpoints to collect data that require only a modicum of technical experience to use.

In my dissertation I go beyond application layer data and show how researchers can access and use different types of Internet measurement data to answer substantive research questions in different fields. I do so by using data from the network layer of the Internet, which resides "below" the application layer (Keremoglu and Weidmann, 2020) where most research using Internet data has so far been carried out. In my first paper, I use data on security vulnerabilities found on servers that host government websites. This data is collected through Internet measurement techniques that rely on how the computers in the network communicate with each other. Every device connected to the Internet exposes dedicated communication endpoints, so-called *ports*. By requesting data from these ports, it is possible to determine what services, such as websites, databases, or email, are

available on a given device. Often, the returned data also includes information on *what version* of a given service is used, and can therefore be connected to security vulnerabilities in the respective service. Using this data, I create an observational measurement of defensive cybersecurity capability.

In the second and third paper I leverage two other kinds of Internet measurement data to investigate questions related to information control and international relations. When users visit a website, they enter a domain which their browser translates into an IP address from which it requests the data. In the second paper, we use historical observations of these domain-IP pairs to identify Denial-of-Service victims among potential censorship targets in autocracies. In my third paper, I use similar data to investigate where governments host their official websites. Because the number of IP addresses is limited, there are institutions that assign usage rights of address ranges to businesses. These businesses then provide these numbers to individual users or website owners. By connecting the IP-address of known government websites to the "owners" of that IP-address at a given point in time, it is possible to infer which businesses host government websites and whether the businesses are domestic or not. This allows insight into how governments navigate the inherent international structure of the Internet.

While the Internet measurement techniques I use in my papers have some limitations, they allow me to conduct comparative, empirical analyses using observational data where previous work was mostly constrained to case studies, reliance on survey measurements, or not able to investigate at all. The main limitation of my approach lies in my need to rely on third parties for historical data – a limitation I address in the next section.

## 5.2 Future Research

While my dissertation makes a significant contribution to the literature, there are some limitations that future work can address. In the following, I will first discuss the limitations of the empirical studies, before I point towards ways in which the use of Internet measurement data generally can be improved in political science and which new questions can be addressed as a consequence.

The main limitation of the first paper is that any indicator of government cybersecurity can only provide comparative measurements between countries, not absolute levels of defensive capabilities – this is true both for observational as well as expert survey-based indicators. Future work can improve on this by connecting defensive capabilities to actual incidents and closely examining the technical and institutional pathways that enable these incidents. One limitation of the second paper is that our analysis focuses on two censorship tactics while the repertoire of techniques available to autocrats contains other tactics as well. Future work can address this through investigating more tactics and considering additional influences to tactical choices such as media attention and possible learning effects across countries. The third paper is limited by its modeling of

factors that might impact the choice of hosting locations for government websites. While international trust likely plays a role, future work can improve on this study by modeling this choice in more detail by including factors such as bilateral trade agreements or other economic variables.

Even though using Internet measurement techniques improves upon the current literature, all three papers and potential future work faces limitations connected to these techniques. The first limitation is the limited availability of historical data collected by third parties. Due to the volume of data generated on the Internet every day, keeping a "full record" akin to newspaper archives is impossible. This means that even large scale data collection efforts by organizations such as Common Crawl or the Internet Archive can only provide data going back to the day the data collection started – anything before that is unrecoverable. An additional consequence is that third-party data collection efforts will always remain sparse when used for specific research questions. For example, our measurement of stable IP periods in the second paper can be improved upon by recording the IP address of relevant websites more frequently going forward. Equally, the datasets used in this dissertation remain limited in their coverage of government websites – only a dedicated data collection effort focused on government websites will be able to remedy this. Second, to use Internet measurement techniques researchers need access to technical infrastructure that enables storing and processing the large volumes of data they are likely to collect – collecting and analysing Internet measurement data is generally not feasible on personal computers but requires dedicated servers. Finally, using Internet measurement techniques requires skills not traditionally taught in political science curricula, like programming, database- and system administration, and knowledge about the technologies that power the Internet.

Where the limitations of Internet measurement techniques can be addressed, however, researchers have the opportunity to answer a range of new questions also going beyond the nexus of governments and the Internet. With frequent active measurements across all network layers, researchers can investigate which cyberattacks, like Distributed Denial-of-Service attacks or website defacements, are used in a political context, and how? Who do governments and other political actors trust with their websites and data? Which departments have the biggest presence online, possibly indicating citizen demand for information? And how does the content of government websites change over time, reflecting different policy priorities?

# A

# Declaration of Authorship

I hereby declare that I am the sole author of the introduction, the first paper, the third paper, the conclusion, and the accompanying material. The second paper is co-authored with Nils B. Weidmann and Alberto Dainotti. In the following, I outline the division of labor.

Alberto Dainotti provided most of the computational resources required to gather the data used in the paper, as well as providing access to the knowledge of the team at CAIDA, in particular Alistair King. Using the technical infrastructure at CAIDA, I gathered data on historical observations of government websites and post-processed it to create a database that allowed me to combine this data with the data on DoS attacks provided by CAIDA. To be able to leverage the data for our analysis, I came up with the idea to create the measurement of "stable IP periods" and carried out the required data processing. After many discussions with Nils B. Weidmann about the most useful frame of analysis for the paper, I completed a first draft of the paper in 2020. While the data collection, post-processing and empirical analysis was exclusively carried out by me, Nils B. Weidmann provided important feedback and direction for the empirical part of the paper. Over multiple rounds of revisions, Nils B. Weidmann supported me significantly in framing and re-writing parts of the paper and contributed the largest share of the theoretical argument.

Konstanz, August 2021

Lukas Kawerau

# B

# Supplementary Material For Chapter 3

## B.1 Stable IP Periods and DDoS Attacks

By itself, the data gathered by CAIDA does not yet allow a direct comparison of DoS attacks and website blocking for censorship. Because CAIDA can only capture the IP address that was attacked, not the intended host/website, it is not directly possible to identify attacks against websites on the Citizenlab list of hosts: we do not know which IP address a particular host was using at a particular point in time, since the mapping of host names to addresses via the Domain Name System (DNS) changes frequently. To solve this problem, we leverage data collected by CommonCrawl (`https://commoncrawl.org`). CommonCrawl is a US-based non-governmental organization that crawls a large portion of the public Internet at four week intervals, scraping the content of the websites it comes across, and makes this content publicly available for research purposes. During each crawl, CommonCrawl collects the data of approximately 2.5 billion web pages from roughly 100 million domains (CommonCrawl, 2019). The collected data contains the content of the visited website, but also the IP address from which the content was delivered and the time and date at which the site was visited.

Combining the data collected by CommonCrawl and CAIDA, we develop a new measurement to better map Denial-of-service Attacks to possible targets. When a given host is visited by CommonCrawl multiple times and has the same IP address in two adjacent observations, we assume that the host had the same address on every day between those two observations. We call this period a "stable IP period," and it is illustrated in Figure B.1: A hypothetical host is observed on July 30th and August 25th. On both of these

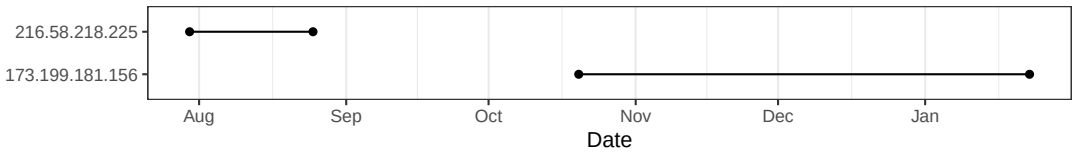


Figure B.1: Stable IP Periods

dates it has the same IP address (216.58.218.225), and we now assume that the host was reachable under this IP address on every day between those two dates. The next time the same host is observed, it has a different IP address: on October 20th it is found on 173.199.181.156, and it keeps this address until the next crawl on January 23rd the next year. From these four observations of the same host, we now get two stable IP periods: between July 30th and August 25th, and between October 20th and January 23rd. For the period between August 25th and October 20th, we do **not** have a stable IP period, because we do not know on what date the host switched between the two addresses. When counting DoS attacks against the hosts in our sample, we only use those that fall within a stable IP period.

In creating this measurement, we balance two constraints of the available data. The most conservative way of combining the CommonCrawl and CAIDA data would be to map only observations that fall on the same day, i.e. DoS attacks that happened on the same day as the crawler visited the website. Given the size of the Internet and the frequency of visits at most every four weeks, this approach would result in very little usable data. At the same time, our approach also introduces some amount of uncertainty: if a host was briefly switched to a different IP address and back in between two observations, we would wrongly attribute any attack that happened during this time to this host.

## B.2 Additional Regression Tables

| | Num Attacks (Log) | | | Any Attack | | |
|---|---|---|---|---|---|---|
| | Model A1 | Model A2 | Model A3 | Model A4 | Model A5 | Model A6 |
| (Intercept) | $0.3469^{**}$ | $0.3114^{***}$ | $0.3149^{**}$ | $0.3411^{***}$ | $0.3276^{***}$ | $0.3123^{***}$ |
| | (0.1128) | (0.0836) | (0.1130) | (0.0671) | (0.0296) | (0.0675) |
| Any Anom | 0.0076 | | $0.0259^{*}$ | 0.0090 | | $0.0288^{*}$ |
| | (0.0118) | | (0.0129) | (0.0105) | | (0.0115) |
| Protest | | −0.0028 | $0.0805^{***}$ | | −0.0020 | $0.0776^{***}$ |
| | | (0.0097) | (0.0240) | | (0.0081) | (0.0214) |
| Any Anom x Protest | | | $-0.0933^{***}$ | | | $-0.1002^{***}$ |
| | | | (0.0261) | | | (0.0233) |
| Any Attack (prev. week) | $0.1816^{***}$ | $0.2439^{***}$ | $0.1785^{***}$ | $0.1496^{***}$ | $0.1786^{***}$ | $0.1463^{***}$ |
| | (0.0192) | (0.0145) | (0.0192) | (0.0171) | (0.0120) | (0.0171) |
| Active SIPs (log) | −0.0070 | −0.0096 | −0.0066 | | | |
| | (0.0247) | (0.0220) | (0.0247) | | | |
| Country Fixed Effect | Yes | Yes | Yes | Yes | Yes | Yes |
| Year Fixed Effect | Yes | Yes | Yes | Yes | Yes | Yes |
| $R^2$ | 0.2279 | 0.2316 | 0.2309 | 0.2525 | 0.2556 | 0.2567 |
| Adj. $R^2$ | 0.2156 | 0.2256 | 0.2182 | 0.2408 | 0.2499 | 0.2446 |
| Num. obs. | 3327 | 6713 | 3327 | 3327 | 6713 | 3327 |

$^{***}p < 0.001$; $^{**}p < 0.01$; $^{*}p < 0.05$

Table B.1: Relationship between the presence of anomalies and number of DoS Attacks (logged), protest, and an interaction between anomaly presence with protest (ICEWS, Country/Week) for V-Dem Autocracies (v2x_polyarchy lower than 0.42)

| | Num Anom (Log) | | | Any Anomaly | | |
|---|---|---|---|---|---|---|
| | Model A7 | Model A8 | Model A9 | Model A10 | Model A11 | Model A12 |
| (Intercept) | $-1.1654^{***}$ | $-1.1769^{***}$ | $-1.1805^{***}$ | $-0.5355^{***}$ | $-0.5451^{***}$ | $-0.5386^{***}$ |
| | (0.2293) | (0.2294) | (0.2302) | (0.1517) | (0.1521) | (0.1522) |
| Any Attack | 0.0447 | | 0.0704 | 0.0266 | | 0.0547 |
| | (0.0403) | | (0.0460) | (0.0267) | | (0.0304) |
| Protest | | 0.0325 | 0.0446 | | 0.0237 | 0.0303 |
| | | (0.0263) | (0.0276) | | (0.0174) | (0.0182) |
| Any Attack x Protest | | | −0.0962 | | | −0.1065 |
| | | | (0.0852) | | | (0.0563) |
| Any Anom (prev. week) | $0.3024^{***}$ | $0.3046^{***}$ | $0.3025^{***}$ | $0.2426^{***}$ | $0.2458^{***}$ | $0.2419^{***}$ |
| | (0.0284) | (0.0281) | (0.0284) | (0.0185) | (0.0183) | (0.0185) |
| Num Reports (log) | $0.9076^{***}$ | $0.9105^{***}$ | $0.9064^{***}$ | | | |
| | (0.0120) | (0.0119) | (0.0120) | | | |
| Country Fixed Effect | Yes | Yes | Yes | Yes | Yes | Yes |
| Year Fixed Effect | Yes | Yes | Yes | Yes | Yes | Yes |
| $R^2$ | 0.8878 | 0.8886 | 0.8880 | 0.3837 | 0.3843 | 0.3848 |
| Adj. $R^2$ | 0.8857 | 0.8865 | 0.8858 | 0.3724 | 0.3733 | 0.3731 |
| Num. obs. | 2788 | 2840 | 2788 | 2788 | 2840 | 2788 |

$^{***}p < 0.001$; $^{**}p < 0.01$; $^{*}p < 0.05$

Table B.2: Relationship between the presence of DoS Attacks and number of anomalies (logged), protest, and an interaction between anomaly presence with protest (ICEWS, Country/Week)

# Bibliography


Asal, Victor, Jacob Mauslein, Amanda Murdie, Joseph Young, Ken Cousins and Chris Bronk. 2016. "Repression, Education, and Politically Motivated Cyberattacks." *Journal of Global Security Studies* 1(3):235–247.

Barberá, Pablo and Thomas Zeitzoff. 2018. "The New Public Address System: Why Do World Leaders Adopt Social Media?" *International Studies Quarterly* 62(1):121–130.

Barlow, John Perry. 1996. "A Declaration Of The Independence Of Cyberspace.".

BBC News. 2019. "Australian Political Parties Hit by 'state Actor' Hack, PM Says." *BBC News* .

Boschee, Elizabeth, Jennifer Lautenschlager, Sean O'Brien, Steve Shellman, James Starz and Michael Ward. 2015. "ICEWS Coded Event Data.".

Bowman, Courtney M. 2015. "A Primer on Russia's New Data Localization Law." https://privacylaw.proskauer.com/2015/08/articles/international/a-primer-on-russias-new-data-localization-law/.

Bulmer, Martin, Kevin Bales and Kathryn Kish Sklar. 1991. *The Social Survey in Historical Perspective, 1880-1940*. Cambridge University Press.

CAIDA, Center for Applied Internet Data Analysis. 2020*a*. "Inferred AS to Organization Mapping Dataset." https://www.caida.org/data/as-organizations/index.xml.

CAIDA, Center for Applied Internet Data Analysis. 2020*b*. "Routeviews Prefix to AS Mappings Dataset (Pfx2as) for IPv4 and IPv6." https://www.caida.org/data/routing/routeviews-prefix2as.xml.

CAIDA, UC San Diego. 2019. "Historical and Near-Real-Time UCSD Network Telescope Traffic Dataset.".

Caruson, Kiki, Susan A. MacManus and Brian D. McPhee. 2012. "Cybersecurity Policy-Making at the Local Government Level: An Analysis of Threats, Preparedness, and Bureaucratic Roadblocks to Success." *Journal of Homeland Security and Emergency Management* 9(2).

Chaffetz, J., M. Meadows and W. Hurd. 2016. "The OPM Data Breach: How the Government Jeopardized Our National Security For More Than a Generation." *Oversight and Government Reform, Tech. Rep.* .

Chander, Anupam and Uyên P. Lê. 2015. "Data Nationalism." http://law.emory.edu/elj/content/volume-64/issue-3/articles/data-nationalism.html.

Cimpanu, Catalin. 2021. "Malaysia Arrests 11 Suspects for Hacking Government Sites." https://www.zdnet.com/article/malaysia-arrests-11-suspects-for-hacking-government-sites/.

Citizen Lab and Others. 2014. "URL Testing Lists Intended for Discovering Website Censorship.".

CommonCrawl. 2019. "August 2019 Crawl Archive Now Available – Common Crawl.".

Coppedge, Michael, John Gerring, Carl Henrik Knutsen, Staffan I. Lindberg, Jan Teorell, David Altman, Michael Bernhard, M. Steven Fish, Adam Glynn and Allen Hicken. 2020. "V-Dem [Country–Year/Country–Date] Dataset V10." *Varieties of Democracy (V-Dem) Project* .

Coppedge, Michael, John Gerring, Carl Henrik Knutsen, Staffan I. Lindberg, Jan Teorell, David Altman, Michael Bernhard, M. Steven Fish, Adam Glynn, Allen Hicken, Anna Lührmann, Kyle L. Marquardt, Kelly McMann, Pamela Paxton, Daniel Pemstein, Brigitte Seim, Rachel Sigman, Svend-Erik Skaaning, Jeffrey Staton, Steven Wilson, Agnes Cornell, Lisa Gastaldi, Haakon Gjerløw, Nina Ilchenko, Joshua Krusell, Laura Maxwell, Valeriya Mechkova, Juraj Medzihorsky, Josefine Pernes, Johannes von Römer, Natalia Stepanova, Aksel Sundström, Eitan Tzelgov, Yi-ting Wang, Tore Wig and Daniel Ziblatt. 2019. "V-Dem [Country-Year/Country-Date] Dataset V9.".

Coppedge, Michael, Staffan Lindberg, Svend-Erik Skaaning and Jan Teorell. 2016. "Measuring High Level Democratic Principles Using the V-Dem Data." *International Political Science Review* 37(5):580–593.

Dai, Xinyuan, Duncan Snidal and Michael Sampson. 2010. International Cooperation Theory and International Institutions. In *Oxford Research Encyclopedia of International Studies*.

Dainotti, Alberto, Claudio Squarcella, Emile Aben, Kimberly C. Claffy, Marco Chiesa, Michele Russo and Antonio Pescapé. 2014. "Analysis of Country-Wide Internet Outages Caused by Censorship." *IEEE/ACM Transactions on Networking (TON)* 22(6):1964–1977.

Data Center Optimization Initiative. 2021. "Data Center Optimization Initiative." https://datacenters.cio.gov/.

Dawn-Hiscox, Tanwen. 2018. "Indonesia's Data Center Industry Protests Data Localization Reform." https://www.datacenterdynamics.com/en/news/indonesias-data-center-industry-protests-data-localization-reform/.

Deibert, Ronald, John Palfrey, Rafal Rohozinski, Jonathan Zittrain and Janice Gross Stein. 2008. *Access Denied: The Practice and Policy of Global Internet Filtering.* MIT Press.

Diamond, Larry. 2010. "Liberation Technology." *Journal of Democracy* 21(3):69–83.

Dorfman, Zach and Jenna McLaughlin. 2018. "The CIA's Communications Suffered a Catastrophic Compromise. It Started in Iran." https://news.yahoo.com/cias-communications-suffered-catastrophic-compromise-started-iran-090018710.html.

Duncombe, Constance. 2017. "Twitter and Transformative Diplomacy: Social Media and Iran–US Relations." *International Affairs* 93(3):545–562.

Dutton, William H. and Laura DeNardis. 2013. The Oxford Handbook of Internet Studies. Oxford University Press chapter The Emerging Field of Internet Governance.

Earl, Jennifer, Andrew Martin, John D McCarthy and Sarah A Soule. 2004. "The Use of Newspaper Data in the Study of Collective Action." *Annual Review of Sociology, Vol 36* 30:65–80.

Egloff, Florian J. and Max Smeets. 2021. "Publicly Attributing Cyber Attacks: A Framework." *Journal of Strategic Studies* 0(0):1–32.

El-Baradei, Mohamed. 2011. "Wael Ghonim – Spokesman for a Revolution." April 21, 2011.

Enikolopov, Ruben, Alexey Makarin and Maria Petrova. 2020. "Social Media and Protest Participation: Evidence From Russia." *Econometrica* 88(4):1479–1514.

Ewing, Thomas. 2018. "Inter-Nyet: The Difficulty of Technological Sovereignty in Russia." https://www.lawfareblog.com/inter-nyet-difficulty-technological-sovereignty-russia.

Filasto, Arturo and Jacob Appelbaum. 2012. OONI: Open Observatory of Network Interference. In *FOCI*.

Fireeye. 2020. "Highly Evasive Attacker Leverages SolarWinds Supply Chain to Compromise Multiple Global Victims With SUNBURST Backdoor." https://www.fireeye.com/blog/threat-research/2020/12/evasive-attacker-leverages-solarwinds-supply-chain-compromises-with-sunburst-backdoor.html.

First, Inc. N.d. "CVSS v2 Complete Documentation." https://www.first.org/cvss/v2/guide.

Fischerkeller, Michael P. and Richard J. Harknett. 2019. "Persistent Engagement, Agreed Competition, and Cyberspace Interaction Dynamics and Escalation." *The Cyber Defense Review* pp. 267–287.

Friedrich, Carl J and Zbigniew K Brzezinski. 1965. "Totalitarian Dictatorship." *Cambridge, MA: Harvard UP* .

Gartzke, Erik. 2013. "The Myth of Cyberwar: Bringing War in Cyberspace Back down to Earth." *International Security* 38(2):41–73.

Geddes, Barbara, Joseph George Wright, Joseph Wright and Erica Frantz. 2018. *How Dictatorships Work: Power, Personalization, and Collapse.* Cambridge University Press.

Gibler, Douglas M. 2008. *International Military Alliances, 1648-2008.* CQ Press.

Gleditsch, Kristian S. and Michael D. Ward. 1999. "Interstate System Membership: A Revised List of the Independent States since 1816." *International Interactions* 25.

Gleditsch, Kristian Skrede and Nils B. Weidmann. 2012. "Richardson in the Information Age: Geographic Information Systems and Spatial Data in International Studies." *Annual Review of Political Science* 15(1):461–481.

Gohdes, Anita R. 2015. "Pulling the Plug: Network Disruptions and Violence in Civil Conflict." *Journal of Peace Research* 52(3):352–367.

Hanada, Ryosuke and Rintaro Tobita. 2020. "Japan to Hire Amazon to Build Government Cloud." https://asia.nikkei.com/Business/Technology/Japan-to-hire-Amazon-to-build-government-cloud.

Hardy, Seth, Masashi Crete-Nishihata, Katharine Kleemola, Adam Senft, Byron Sonne, Greg Wiseman, Phillipa Gill and Ronald J Deibert. 2014. Targeted Threat Index: Characterizing and Quantifying Politically-Motivated Targeted Malware. In *23rd {USENIX} Security Symposium ({USENIX} Security 14)*. pp. 527–541.

Harknett, Richard J. and Max Smeets. 2020. "Cyber Campaigns and Strategic Outcomes." *Journal of Strategic Studies* 0(0):1–34.

Hassanpour, Navid. 2014. "Media Disruption and Revolutionary Unrest: Evidence From Mubarak's Quasi-Experiment." *Political Communication* 31(1):1–24.

Hellmeier, Sebastian. 2016. "The Dictator's Digital Toolkit: Explaining Variation in Internet Filtering in Authoritarian Regimes." *Politics & Policy* 44(6):1158–1191.

Hobbs, William R. and Margaret E. Roberts. 2018. "How Sudden Censorship Can Increase Access to Information." *American Political Science Review* 112(3):621–636.

Hoffman, Aaron M. 2002. "A Conceptualization of Trust in International Relations." *European Journal of International Relations* 8(3):375–401.

Irion, Kristina. 2012. "Government Cloud Computing and National Data Sovereignty." *Policy & Internet* 4(3-4):40–71.

Jagannathan, Malavika. 2012. "DDoS Attacks Disable Independent News Sites during Russian Protests.".

Jones, Benjamin T. and Eleonora Mattiacci. 2019. "A Manifesto, in 140 Characters or Fewer: Social Media as a Tool of Rebel Diplomacy." *British Journal of Political Science* 49(2):739–761.

Jones, Penny. 2012. "Hong Kong Lists Government Cloud Suppliers." https://www.datacenterdynamics.com/en/news/hong-kong-lists-government-cloud-suppliers/.

Kasuya, Yuko and Kota Mori. 2019. "Better Regime Cutoffs for Continuous Democracy Measures." *V-Dem Users Working Paper Series* (25).

Kelly, Sanja, Mai Truong, Adrian Shahbaz, Madeline Earp and Jessica White. 2017. "Freedom on the Net 2017 - Manipulating Social Media to Undermine Democracy." *Freedom House* .

Keohane, Robert O. 1986. "Reciprocity in International Relations." *International organization* 40(1):1–27.

Keremoglu, Eda and Nils B. Weidmann. 2020. "How Dictators Control the Internet: A Review Essay:." *Comparative Political Studies* .

Khrennikov, Iljya, Aliaksandr Kudrytski and Alexander Sazonov. 2020. "Tech Workers Flee Belarus as IT Haven Takes Authoritarian Turn." *Bloomberg.com* .

King, Gary, Jennifer Pan and Margaret E Roberts. 2013. "How Censorship in China Allows Government Criticism but Silences Collective Expression." *American Political Science Review* 107(02):326–343.

King, Gary, Jennifer Pan and Margaret E Roberts. 2014. "Reverse-Engineering Censorship in China: Randomized Experimentation and Participant Observation." *Science* 345(6199):1251722.

Kydd, Andrew H. 2007. *Trust and Mistrust in International Relations*. Princeton University Press.

Lambach, Daniel. 2020. "The Territorialization of Cyberspace*." *International Studies Review* 22(3):482–506.

Larson, Jennifer M., Jonathan Nagler, Jonathan Ronen and Joshua A. Tucker. 2019. "Social Networks and Protest Participation: Evidence from 130 Million Twitter Users." *American Journal of Political Science* 63(3):690–705.

Lee, Micah. 2020. "Hack of 251 Law Enforcement Websites Exposes Personal Data of 700,000 Cops." https://theintercept.com/2020/07/15/blueleaks-anonymous-ddos-law-enforcement-hack/.

Leeds, Brett Ashley. 2003. "Alliance Reliability in Times of War: Explaining State Decisions to Violate Treaties." *International Organization* pp. 801–827.

Leeds, Brett Ashley, Michaela Mattes and Jeremy S. Vogel. 2009. "Interests, Institutions, and the Reliability of International Commitments." *American Journal of Political Science* 53(2):461–476.

Lessig, Lawrence. 1997. "Tyranny in the Infrastructure." *Wired* .

Lu, Yingdan and Jennifer Pan. 2020. "Capturing Clicks: How the Chinese Government Uses Clickbait to Compete for Visibility." *Political Communication* 0(0):1–32.

Lubold, Shane Harris and Gordon. 2017. "Russia Has Turned Kaspersky Software Into Tool for Spying." *Wall Street Journal* .

Lutscher, Philipp M., Nils B. Weidmann, Margaret E. Roberts, Mattijs Jonker, Alistair King and Alberto Dainotti. 2020. "At Home and Abroad: The Use of Denial-of-Service Attacks during Elections in Nondemocratic Regimes." *Journal of Conflict Resolution* 64(2-3):373–401.

Marquardt, Kyle L., Daniel Pemstein, Brigitte Seim and Yi-ting Wang. 2019. "What Makes Experts Reliable? Expert Reliability and the Estimation of Latent Traits." *Research & Politics* 6(4):2053168019879561.

Mattes, Michaela and Mariana Rodríguez. 2014. "Autocracies and International Cooperation." *International Studies Quarterly* 58(3):527–538.

McLaughlin, Jenna and Zach Dorfman. 2019. "'Shattered': Inside the Secret Battle to Save America's Undercover Spies in the Digital Age." https://news.yahoo.com/shattered-inside-the-secret-battle-to-save-americas-undercover-spies-in-the-digital-age-100029026.html.

Mechkova, V., D. Pemstein, B. Seim and S. Wilson. 2019*a*. Digital Society Project (DSP) Codebook V1. Technical report Working Paper.

Mechkova, V., D. Pemstein, B. Seim and S. Wilson. 2019*b*. Measuring Internet Politics: Introducing the Digital Society Project (DSP). Technical report Working Paper.

Merkel, Wolfgang. 2004. "Embedded and Defective Democracies." *Democratization* 11(5):33–58.

Moore, David, Colleen Shannon, Douglas J Brown, Geoffrey M Voelker and Stefan Savage. 2006. "Inferring Internet Denial-of-Service Activity." *ACM Transactions on Computer Systems (TOCS)* 24(2):115–139.

Morozov, Evgeny. 2011. *The Net Delusion : The Dark Side of Internet Freedom*. PublicAffairs.

Morrison, Michael Ian. 2013. "The Acquisition Supply Chain and the Security of Government Information Technology Purchases." *Public Contract Law Journal* pp. 749–792.

Motamedi, Maziar. 2019. "Locked out: Why Is Amazon Blocking Iranians from Its Services?" https://www.aljazeera.com/economy/2019/10/2/locked-out-why-is-amazon-blocking-iranians-from-its-services.

Mozilla Foundation. 2020. "Public Suffix List." https://publicsuffix.org/.

Mueller, Milton, Andreas Schmidt and Brenden Kuerbis. 2013. "Internet Security and Networked Governance in International Relations." *International Studies Review* 15(1):86–104.

Mueller, Milton L. 2020. "Against Sovereignty in Cyberspace." *International Studies Review* 22(4):779–801.

Nazario, Jose. 2009. "Politically Motivated Denial of Service Attacks." *The Virtual Battlefield: Perspectives on Cyber Warfare* pp. 163–181.

Neil. 2021. "Main Philippine Gov't Portal Hacked after Death of 9 Activists." https://www.bworldonline.com/main-philippine-govt-portal-hacked-after-death-of-9-activists/.

Ni, Anna Ya and Stuart Bretschneider. 2007. "The Decision to Contract Out: A Study of Contracting for E-Government Services in State Governments." *Public Administration Review* 67(3):531–544.

Norris, Donald F., Laura Mateczun, Anupam Joshi and Tim Finin. 2019. "Cyberattacks at the Grass Roots: American Local Governments and the Need for High Levels of Cybersecurity." *Public Administration Review* 79(6):895–904.

Norris, Pippa. 2001. *Digital Divide: Civic Engagement, Information Poverty, and the Internet Worldwide*. Cambridge University Press.

Nye Jr, Joseph S. 2016. "Deterrence and Dissuasion in Cyberspace." *International security* 41(3):44–71.

Office of Management and Budget. 2016. "M-16-19 Data Center Optimization Initiative.".

Office of Management and Budget. 2019. "M-19-19 Update to Data Center Optimization Initiative (DCOI).".

Okunoye, Babatunde, Maria Xynou, Leonid Evdokimov, Sodiq Alabi and Chukwuzitere Okoli. 2018. *Measuring Internet Censorship in Nigeria.*

OONI, Open Observatory of Network Interference. 2020. "Web Connectivity Test." https://ooni.org/nettest/web-connectivity/.

Oppenheim, L. (Lassa). 1912. *International Law. A Treatise. Volume 1 (of 2) Peace. Second Edition.*

Paquette, Scott, Paul T. Jaeger and Susan C. Wilson. 2010. "Identifying the Security Risks Associated with Governmental Use of Cloud Computing." *Government Information Quarterly* 27(3):245–253.

Pemstein, Daniel, Kyle L. Marquardt, Eitan Tzelgov, Yi-ting Wang, Joshua Krusell and Farhad Miri. 2019. "The V-Dem Measurement Model: Latent Variable Analysis for Cross-National and Cross-Temporal Expert-Coded Data." *V-Dem Working Paper* 21.

Philpott, Daniel. 1995. "Sovereignty: An Introduction and Brief History." *Journal of International Affairs* 48(2):353–368.

PTI. 2017. "Govt IT Data on Cloud System Must Be Stored within India: Meity." https://www.livemint.com/Industry/OWeeqJSiHwcDFrOH9kJQoN/Govt-IT-data-on-cloud-system-must-be-stored-within-India-Me.html.

Reuters. 2019. "Foreign Power Was behind Cyber Attack on Czech Ministry: Senate." *Reuters* .

Reuters. 2020. "Czech Hospitals Report Cyberattacks Day after National Watchdog's Warning." *Reuters* .

Richardson, Lewis F. 1960. *Statistics of Deadly Quarrels.* Boxwood Press.

Rid, Thomas. 2012. "Cyber War Will Not Take Place." *Journal of strategic studies* 35(1):5–32.

Roberts, Margaret E. 2018. *Censored: Distraction and Diversion inside Chinas Great Firewall.* Princeton University Press.

Selby, John. 2017. "Data Localization Laws: Trade Barriers or Legitimate Responses to Cybersecurity Risks, or Both?" *International Journal of Law and Information Technology* 25(3):213–232.

Solomon, Shoshanna. 2021. "Israel Signs Deal for Cloud Services with Google, Amazon." https://www.timesofisrael.com/israel-signs-deal-for-cloud-services-with-google-amazon/.

Stevens, Tim. 2018. "Cyberweapons: Power and the Governance of the Invisible." *International Politics* 55(3):482–502.

Stilgherrian. 2018. "What's Actually in Australia's Encryption Laws? Everything You Need to Know." https://www.zdnet.com/article/whats-actually-in-australias-encryption-laws-everything-you-need-to-know/.

Sundara Raman, Ram, Prerana Shenoy, Katharina Kohls and Roya Ensafi. 2020. Censored Planet: An Internet-Wide, Longitudinal Censorship Observatory. In *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security*. CCS '20 New York, NY, USA: Association for Computing Machinery pp. 49–66.

Swinhoe, Dan. 2021*a*. "National Data Center in Zimbabwe Opens." https://www.datacenterdynamics.com/en/news/national-data-center-zimbabwe-opens/.

Swinhoe, Dan. 2021*b*. "Senegal to Migrate All Government Data and Applications to New Government Data Center." https://www.datacenterdynamics.com/en/news/senegal-to-migrate-all-government-data-and-applications-to-new-government-data-center/.

The MITRE Corporation. 2021. "CVE Program - About." https://cve.mitre.org/about/index.html.

The World Bank. 2019*a*. Individuals Using the Internet (% of Population). Technical report.

The World Bank. 2019*b*. Population, Total. Technical report.

Theocharis, Yannis and Andreas Jungherr. 2020. "Computational Social Science and the Study of Political Communication." *Political Communication* 0(0):1–22.

Thompson, Nik, Antony Mullins and Thanavit Chongsutakawewong. 2020. "Does High E-Government Adoption Assure Stronger Security? Results from a Cross-Country Analysis of Australia and Thailand." *Government Information Quarterly* 37(1):101408.

United Nations. 2020*a*. "E-Government Development Index - Methodology." https://publicadministration.un.org/egovkb/en-us/About/Methodology.

United Nations. 2020*b*. "E-Government Development Index - Overview." https://publicadministration.un.org/egovkb/en-us/Overview.

Von Der Burchard, Hans. 2020. "Merkel Blames Russia for 'Outrageous' Cyberattack on German Parliament." https://www.politico.eu/article/merkel-blames-russia-for-outrageous-cyber-attack-on-german-parliament/.

Von Stein, Jana. 2005. "Do Treaties Constrain or Screen? Selection Bias and Treaty Compliance." *American Political Science Review* pp. 611–622.

Weidmann, Nils B. and Espen Geelmuyden Rød. 2019. *The Internet and Political Protest in Autocracies*. Oxford Univ. Press.

Weidmann, Nils B. and Kristian S. Gleditsch. 2010. "Mapping and Measuring Country Shapes: The CShapes Package." *R J* 2:18–23.

Weidmann, Nils B, Suso Benitez-Baleato, Philipp Hunziker, Eduard Glatz and Xenofontas Dimitropoulos. 2016. "Digital Discrimination: Political Bias in Internet Service Provision across Ethnic Groups." *Science* 353(6304):1151–1155.

Whalen, Jeanne. 2021. "U.S. Campaign against Huawei Appears to Be Working, as Chinese Tech Giant Loses Sales Outside Its Home Market." *Washington Post* .

Wintrobe, Ronald. 2000. *The Political Economy of Dictatorship*. Cambridge University Press.

Zhao, Jensen J. and Sherry Y. Zhao. 2010. "Opportunities and Threats: A Security Assessment of State e-Government Websites." *Government Information Quarterly* 27(1):49–56.