

Google Politics: The Political Determinants of Internet Censorship in Democracies*

STEPHEN A. MESERVE AND DANIEL PEMSTEIN

The expansion of digital interconnectivity has simultaneously increased individuals' access to media and presented governments with new opportunities to regulate information flows. As a result, even highly democratic countries now issue frequent censorship and user data requests to digital content providers. We argue that government internet censorship occurs, in part, for political reasons, and seek to identify the conditions under which states censor. We leverage new, cross-nationally comparable, censorship request data, provided by Google, to examine how country characteristics co-vary with governments' digital censorship activity. Within democracies, we show that governments engage in more digital censorship when internal dissent is present and when their economies produce substantial intellectual property. But these demand mechanisms are modulated by the relative influence that democratic institutions provide to narrow and diffuse interests; in particular, states with proportional electoral institutions censor less.

The explosion of the internet and commensurate increase in citizens' ability to create, disseminate, and access media presents states with a host of regulatory challenges that did not exist before the rise of widespread digital networks. Early observers of internet expansion believed that it fundamentally interfered with governments' ability to regulate political and economic activity (Castells 1996) by encouraging politicians to reduce regulation to placate increasingly mobile firms (Tonnelson 2000), introducing complexities that undermined the clarity and legitimacy of legal jurisdictions (Post and Johnson 1996), and by frustrating geographically restricted authorities' attempts to police the flow of information across fluid, multi-national, digital networks (Barlow 1996; Haufler 2001; Rosenau and Singh 2002). This logic was summed up by John Gilmore's oft-cited comment that, "... the Net interprets censorship as damage and routes around it" (Elmer-DeWitt 1993).¹

But more recent scholarship notes the surprising flexibility with which governments adapted to the rapid expansion of digital interconnectivity and the fundamental role that states played in shaping the legal boundaries of information technology (Reidenberg 1998; Geist 2003; Newman and Zysman 2005; Goldsmith and Wu 2006).² While multiple authors highlight the role that private actors—and the communications protocols and infrastructure that they develop and control—have in regulating digital content (Reidenberg 1998; Lessig 1999; DeNardis 2009; DeNardis 2012), others emphasize the tendency of governments to use these private

* Stephen A. Meserve is an Assistant Professor in the Department of Political Science, Texas Tech University, 113 Holden Hall Boston & Akron Streets, Lubbock, TX 79409 (stephen.meserve@ttu.edu), Daniel Pemstein is an Assistant Professor in the Department of Criminal Justice and Political Science, North Dakota State University, 1661 12th Avenue North, Fargo, ND 58108 (daniel.pemstein@ndsu.edu). The authors would like to thank seminar participants at Texas Tech University, the European Political Science Association conference, and the annual meeting of the Midwest Political Science association, and two anonymous reviewers for helpful comments. Sivagaminathan Palani provided invaluable research assistance. The authors contributed equally to this work, which was supported, in part, by a Google Faculty Research Award. To view supplementary material for this article, please visit <https://doi.org/10.1017/psrm.2017.1>

¹ Quoted in Newman and Zysman (2005), Farrell (2006), and Breindl (2013).

² See Breindl (2013) for a review of the literature on internet content regulation in democracies.

actors—notably internet service providers (ISPs) and online content providers (OCPs)—as “points of control” to regulate the digital behavior of a variety of other private actors (Zittrain 2003; Farrell 2006; Adler 2011; MacKinnon 2012). Thus, by focusing their regulatory efforts on large content and connectivity providers that cannot easily uproot their businesses to avoid state oversight, governments developed an internet regulation strategy that takes advantage of the structure of digital networks, partially overcoming the fluidity of the content that resides on those networks.

Fundamentally, governments—especially in liberal democracies—now substantially rely on private actors to regulate public internet access. And crucially, while the decentralized structure of the internet is often seen as an important bulwark against state interference, scholars routinely worry that regulation by proxy provides states a route around formal speech and privacy protections, and that—at least in states with strong rule of law—privately provided limits on internet content are subject to fewer legal constraints than outright government censorship, undermining civil liberties and consumer protections (Boyle 1997; Adler 2011; Marsden 2011). Nonetheless, because studies of internet censorship tend to focus on firewalls and filtering tactics in authoritarian regimes (see e.g., Deibert et al. 2008; King, Pan and Roberts 2013), we have only limited understanding of, and even less comparative empirical evidence about, the circumstances under which authorities in democratic states censor digital sources (Breindl 2013). In particular, we do not know when and why democratic governments take advantage of private points of control (PPCs) and ask service providers to remove content from their networks.³ Exploring this question is fundamental to building a comprehensive empirical understanding of the politics of internet regulation. There are strong practical and normative reasons to examine the conditions under which democratic states pressure ISPs and OCPs to remove content from the internet. We observe vast differences in this form of digital censorship across democracies. Why do some democracies make extensive use of such censorship mechanisms while others rarely bother service providers with content removal requests?

Crucially, we argue that even democratic states seek to curtail content dissemination in response to demands to restrict speech, either to influence public opinion, reduce criticism of public officials, limit citizens’ access to media, and other sources of information, or, more benevolently, to bolster national security or protect individuals’ reputations or privacy. In particular, using both large-*N* panel data and a synthetic case-control design, we show that democratic governments respond to internal opposition, criticism, or unrest by stepping up their digital censorship activities. While our understanding of authoritarian repression hinges on how regimes control information to limit dissent, this work shows that, by leaning on PPCs, even democracies circumscribe digital speech when they face internal unrest.

While our empirical focus is on the relationship between internal unrest and digital censorship, we also highlight two other factors that encourage the construction and maintenance of digital content regulation regimes. First, regulators may seek to protect the intellectual property (IP) of their citizens. States that house firms that hold extensive IP portfolios have an incentive to protect knowledge industries, and should use PPCs to cater to IP-producing firms. Second, demands for content removal often pit the preferences of concentrated actors—politicians, individuals or firms with reputations to protect, and IP producers—against broad societal and consumer interests in unfettered access to media, open information flows, and freedom of speech. Thus, we expect governments operating under political institutions that provide comparative advantages to narrow interests to make more intensive use of PPCs. Building on

³ But see Farrell (2006) for an examination of how (dis)similarity in regulatory preferences across states, and the presence or absence of such points of control, affects governments’ broad regulatory policy choices.

the trade protection literature (Rickard 2012, for example), we argue that governments with less proportional electoral systems will regulate content more aggressively.

This paper represents one of the first attempts to systematically examine internet censorship cross-nationally.⁴ We use data on government requests for content removal, furnished by the Google Corporation, to test our arguments. Google provides a common set of services—search, YouTube, Google+, and so forth—to consumers around the world.⁵ These data represent a unique opportunity for scholars of censorship, because it is difficult to directly observe most forms of censorship activity by governments.⁶ It is rarer still for data on censorship activities to be collected comprehensively across countries. In this context, however, governments rely on a single third party to execute their wishes. Indeed, since Google provides similar products to individuals within multiple countries, these data provide an ideal opportunity to test the determinants of censorship from a comparative perspective. They provide a consistent, cross-national window into when—and how often—governments use PPCs to regulate their citizens' access to digital content.

DEFINING INTERNET CENSORSHIP

We define internet censorship as actions taken by a government to remove or obscure internet content from citizens, or to limit the digital transmission of information to a broad audience. Our conceptualization is intentionally devoid of any of the normative content sometimes associated with censorship, adhering to the dictionary definition of a censor: “a person who examines books, movies, letters, etc., and removes things that are considered to be offensive, immoral, harmful to society, etc.” (Merriam-Webster 2014).

While the mechanisms that produce internet takedown requests may vary in their levels of institutionalization and legitimacy, we argue that it is inappropriate to make *a priori* distinctions about content removal requests here. First, we simply lack the tools and information to measure validity and process, and thus cannot discriminate between requests objectively. More importantly, we do not find such distinctions helpful when predicting censorship activities cross-nationally. Most people agree that muzzling domestic criticism of government is both a form of censorship and normatively unappealing. But activities that some would consider simple applications of law—such as curbing defamation and aggressively protecting IP rights—remain contested and intensely political. For example, while limiting defamation serves a useful societal purpose, it also restricts freedom of speech, and the value that society places on these two goals varies both across and within countries. Similarly, the appropriate scope of IP protection is a hotly debated question (Electronic Frontier Foundation 2010); reasonable people disagree about how to balance consumers' interests in the free flow of information, or low-cost access to medicine and technology, against producer protections designed to ensure just compensation for investment in research and creative activity, and to spur innovation. Some people consider IP protection, at least certain forms, “censorship,” while others do not. Moreover, defamation and IP law create winners and losers and the application and scope of these laws are political decisions that vary over time and space (Baldwin 2014).

⁴ But see the ONI project on internet filtering (Deibert et al. 2011).

⁵ While a graduate research assistant for this project was supported by a Google Faculty Research Award (GFRA), Google provided no input on any aspect of the research reported here. Indeed, the award took the form of an unrestricted gift to Texas Tech University. All analysis in this paper was conducted using publicly available data and do not reflect privileged access to Google data or internal sources. To maximize transparency, we include a copy of our GFRA application and the award letter in the replication package for this article.

⁶ See King, Pan and Roberts (2013) for a notable exception in the authoritarian context.

Thus, because common standards for what constitutes censorship are inherently normative, vary broadly, and are shaped by political conflict, we opt for an inclusive definition: for our purposes, censorship is simply the act of restricting public access to content.

POLITICAL INCENTIVES AND INTERNET CENSORSHIP

Internet censorship is a political activity. That is, while a variety of incentives—ranging from information and speech control, to the maintenance of privacy, copyright and IP protection, and national security concerns—may all drive governments to remove digital content, these activities are filtered through political lenses. We therefore identify and test three political determinants of censorship through PPCs: the need to muzzle opposition from internal challengers, the demand for censorship emanating from IP interests, and democratic institutions that encourage political responsiveness to concentrated interests.

Internal Unrest

Individual citizens, businesses, and politicians can all generate demands for government censorship that are based on personal, or political, motivations. Content on the internet may defame individuals, criticize politicians, impinge on personal privacy, or violate national security statutes. Attempts to limit the availability of such content generate removal requests.

However, incentives to restrict contested, controversial, or politically sensitive speech varies across democracies. While this variety is driven by a number of factors, we focus on a particular issue, internal unrest, because of its political salience and importance to the quality of democratic governance. A democracy with passive internal rivals and strong internal stability creates limited demand for censorship from the regime itself. Political motivations to censor are most prominent when a country is shaken by internal dissent. Because free speech is fundamental to democratic politics, democratic governments will face legal obstacles to direct assaults on political speech. Thus, they are likely to appeal to concerns, such as national security, or defamation, when attempting to silence rivals. Riots, protests, terrorism, and other forms of large-scale, or violent, anti-regime activity may incentivize even democratic politicians to muzzle opposition; they also provide a rationale for—or even legitimize—speech curtailment. And, in reasonably democratic states, especially those with strong rule of law, PPCs may provide an especially attractive venue for quashing political criticism, specifically because private actors have substantial legal leeway to govern the content that resides on their networks. Indeed, it will often be financially and politically expedient for ISPs and OCPs to submit to government pressure in circumstances under which end users might appeal to legal speech protections.⁷ Finally, governments may react to internal opposition not only by requesting more censorship, but by creating legal and regulatory frameworks to facilitate digital censorship.

HYPOTHESIS 1: Democracies will lodge more digital content removal requests when they face higher levels of internal unrest.

⁷ Clearly, ISPs and OCPs must tread carefully on this issue to avoid antagonizing their customers. Yet, while such firms have incentives to resist large-scale government intervention, their preferences are unlikely to align with those of their users on a takedown-by-takedown basis. Furthermore, end users typically enter into contracts with ISPs and OCPs that leave them with few options should a firm decide to honor a government's takedown request.

IP

Many governments aggressively protect the IP of their citizens and businesses. Indeed, the United States has an “IP Tsar” (formally, the Intellectual Property Enforcement Coordinator) tasked with developing and evaluating IP protection policy (Espinel 2012). The US case illustrates a key point: states that generate significant IP, relative to consumption, have higher incentives to protect that property than do other regimes (Sell 1998). Public efforts to reduce IP theft are a subsidy to knowledge industries. They benefit specific industries and firms, at cost to the taxpayer.⁸

The development and enforcement of IP law are political acts. Indeed, Baldwin (2014) extensively describes the inherent political conflict surrounding copyright, noting how IP laws in the United States and Europe changed over time, varying with political and economic conditions. Initially a net importer of creative works, and an inveterate intellectual pirate state, the United States refused to recognize foreign copyright until the late 1800s. Now the United States is a champion of IP protection worldwide. American politicians’ appetite for IP protection has grown in tandem with domestic creative industries. And continental Europe, traditionally more copyright friendly than the anglophone world, now houses multiple political parties representing content “pirates.”

In sum, the size of local knowledge industries should play an important role how aggressively states regulate IP. Furthermore, the political influence wielded by IP producers should also facilitate increased non-IP-related censorship by altering legal institutions to ease the production of takedown requests.

HYPOTHESIS 2: Democracies housing firms that produce substantial IP will pursue more digital content removal requests than those with small knowledge industries.

Domestic Political Institutions

Democracies sporting institutions that empower diffuse interests, like consumers, will generate fewer takedown requests than those with institutions that particularly respond to concentrated interests, such as IP-producing businesses, defamation targets, and politicians themselves. We highlight one such institution: the electoral system. We describe two mechanisms for generating content removal requests that should be modulated by district magnitude.⁹ The key actors supporting both mechanisms are elected politicians; while both firms and individuals play a role in the processes that we describe, our theoretical focus is on how politicians translate personal incentives, and pressures placed on them by firms and individuals, into behavior that should predict government use of PPCs.

First, because politicians’ reputations become more important as district magnitude shrinks, small district magnitudes should enhance politicians’ appetites to quash personal criticism, generating a direct incentive for actors in government to restrict digital speech through tools such as defamation claims. Many of the vignettes that Google provides as part of its transparency report, some of which we describe in third and fifth sections, reflect such behavior, which ranges from demands to remove content issued by executive or legislative organs, to

⁸ Of course, some IP protection provides a public good, by stimulating innovation. Nonetheless, creative industries benefit from regulatory regimes that over-protect IP, and may support levels of IP protection that encourage rent-seeking. Furthermore, non-IP producers have little incentive to contribute, individually, to the provision of IP protection. Thus, the median voter (representative consumer) is likely to support less IP protection than would producers.

⁹ District magnitude is simply the number of legislative seats elected in each district.

defamation charges filed with the courts (Google Incorporated 2013). While large-magnitude systems can generate incentives for personal vote cultivation when co-partisans compete for preference votes in open lists (Carey and Shugart 1995), candidate name recognition shrinks with district magnitude. Even in open-list systems, incumbents in large-magnitude districts have substantially lower name recognition advantages vis-à-vis challengers than those that compete in single member districts (Samuels 2001). Similarly, in low-magnitude systems, incumbents all have an interest in protecting themselves from less known challengers, and to collude in the provision of tools for reputation protection, while incumbents in larger magnitude elections that do encourage personal vote cultivation (e.g., open list) may worry that their counterparts will use speech limitations to restrict their ability to criticize one another. Therefore, political demand for personality-driven speech regulation should decrease as district magnitude grows.

The second mechanism that we describe is less direct. Mirroring a long-standing thread of the trade literature, we argue that electoral systems determine how effectively politicians translate the preferences of narrow interests—which here will tend to favor increased digital regulation—into policy. Politicians in low-magnitude systems will be more receptive to such lobbies and will tend to pass legislation, and engage in bureaucratic oversight, that facilitates protection of firms' IP and digital content regulation more broadly.¹⁰ Indeed, many authors argue that plurality-based, or low district magnitude elections, encourage politicians to cater to concentrated interests and overlook the diffuse preferences of the majority (see e.g., Magee, Brock and Young 1989; Rogowski 1989; Grossman and Helpman 2005; Persson and Tabellini 2005).¹¹

HYPOTHESIS 3: States that use high district magnitude elections will generate fewer digital content removal requests than low district magnitude systems.

DATA AND METHODS

Dependent Variable: Google Takedown Requests

We draw our dependent variable from Google's online "Transparency Report." Google has publishes counts of takedown and user data requests lodged by governments worldwide (Google Incorporated 2013). Google began reporting these data in the second half of 2009 and issues their data in the form of half-year summaries, by country, of censorship attempts by government sources. Google reports currently list the number of content takedown requests issued by 58 different governments. These removal requests can be related to any of Google's many services (Search, YouTube, Gmail, Google+, etc.). Each individual request by a government may identify one or more pieces of digital content for takedown (e.g., multiple related defamatory images), but represents a single instance of attempted government censorship. Multiple attempts to censor the same item are counted as multiple requests. Thus, we measure censorship attempts in terms of government contacts rather than in terms of individual pieces of content. Takedown request counts omit activity that Google performs at its own initiative, regardless of local law, particularly removal of child pornography. Finally, the data

¹⁰ The Digital Millennium Copyright Act (DMCA) is the classic example of a law that provides sweeping legal provisions that facilitate digital content regulation, but which critics have argued had the unintended consequence of allowing for government interference in the transmission of digital speech in ways that have little to do with the protection of IP (Electronic Frontier Foundation 2010).

¹¹ See Bawn and Thies (2003), Rogowski and Kayser (2002), and Chang et al (2010) for counter-arguments. Space considerations preclude us from adequately engaging this literature here. Therefore, we leave a more thorough theoretical development of this argument to future work.

TABLE 1 *Google Takedown Request Categories*

Category	2010–2013 Distribution (%)
Adult content	2
Copyright	3
Defamation	35
Electoral law	3
Government criticism	3
Hate speech	2
Impersonation	2
National security	1
Other	10
Privacy and security	15
Religious offense	2
Trademark	2
Violence	2

Note: Takedown percentage is percent of total from the July 2010 to July 2013 period, omitting reasons more recently introduced in 2012 (18 percent), calculated from recent Google provided data (Google Incorporated 2016).

contains requests related to IP when they take the form of successful court proceedings and actions taken by government agencies but these data do not contain copyright requests that firms and other rights holders (e.g., Recording Industry Association of America) issue directly to Google, which Google fields using a different system. The data contain any takedowns coming from court orders, executive branch interventions, and other direct government activity.¹²

Google now categorizes takedown requests, as we show in Table 1. These breakdowns are only available after July 2010, making it difficult to use this information in over time analyses, such as those presented in this paper. Nevertheless, the descriptive data provide a picture of why governments censor. Takedown requests run the gamut between reasonably clear speech-suppressing activity, like government criticism, to more regulatory activities like copyright or trademark violations. The largest categories are “defamation,” “privacy and security,” and “other,” which contains a mix of anti-speech and regulatory activities. Furthermore, while many of the censorship attempts that fall into these three categories would be generally regarded as legitimate governance, they provide rationales for speech limitation that governments could potentially misuse, and that worry critics of the rise in state leverage over PPCs. Unfortunately, Google does not provide full details for all individual takedown requests. Nevertheless, it does provide descriptive vignettes for selected requests.¹³ In general, Google appears to prefer to highlight non-regulatory, politically motivated, takedown requests by governments. Consider this example of defamation from Italy in early 2011: “We received a request from the Central Police in Italy to remove a YouTube video that satirized Prime Minister Silvio Berlusconi’s lifestyle. We did not remove content in response to this request.” Nevertheless, many requests are, in fact, more mundane, regarding economic regulation or legal requirements: “[In Norway] two requests resulted in the removal of 1814 items from AdWords for violating Norwegian marketing laws.” It is also worth noting that governments sometimes attempt to censor political criticism under the guise of regulatory activity; for example, Google determined that a blog post

¹² For more information on the details of Google takedown and compliance data, see Google’s transparency report FAQ (<https://www.google.com/transparencyreport/removals/government/faq/>).

¹³ All quotes taken from the notes section of Google’s transparency report (Google Incorporated 2013).

that the Bolivian legislative assembly requested be removed because it “infringed copyright” actually contained political speech.

While Google clearly fields takedown requests that are politically motivated, our inability to disaggregate requests is a limitation of the current analysis, because we cannot distinguish between “good” and “bad” censorship. It is also possible that the content of takedown requests co-varies with our key predictors. Nonetheless, we argue that understanding broad trends in digital takedown requests contributes significantly to our understanding of how governments use PPCs.

Independent Variables

We measure internal unrest three ways. First, we use the logged total number of terrorist incidents (National Consortium for the Study of Terrorism and Responses to Terrorism 2012).¹⁴ Terrorist attacks represent an explicit indicator of violent, anti-government activity. We also performed tests using a non-events indicator, the Worldwide Governance Indicators’ (WGI) Political Stability and Absence of Violence. WGI’s index intends to capture “the likelihood the government will be destabilized or overthrown by unconstitutional or violent means” (Kaufmann, Kraay and Mastruzzi 2012). Compared with terrorist events, however, this index is a less tangible and precise measure of internal opposition and stability, reflecting “hundreds” of individual underlying variables (Kaufmann, Kraay and Mastruzzi 2010).

We contend that states that produce IP at high rates are likely to house firms that put pressure on politicians to protect IP. We use patent applications as a proxy for IP production and measure cross-national patent production using the World Intellectual Property Organization’s IP database (World Intellectual Property Organization 2013), which runs through 2012. We adjust the number of patents for population using the World Development Indicators (WDI) data set (The World Bank 2013), creating a patents per capita variable. Finally, because the distribution of patent applications is highly skewed, we log the resulting indicator.

We operationalize the proportionality of the electoral system using a measure of average lower house district magnitude provided by the Database of Political Institutions (DPI) (Beck et al. 2001), logging observations to correct for substantial positive skew. We draw these data from the Quality of Government (QoG) data set. The QoG data contain yearly observations of district magnitude covering our entire 2009–2012 observation period.¹⁵

Turning to control variables, we use a measure of internet users per capita, drawn from the WDI, to account for the size of the digital information environment. We control for economic development—and government capacity—using logged gross domestic product (GDP) per capita, in 2005 US dollars, again drawing from the WDI. The WDI also record the average time it takes to start a business, including licensing delays and other red tape, which we use as a proxy for bureaucratic activity and intrusiveness. Finally, to control for the effects of Google’s market share on takedowns, we use Google’s percentage of internet search to proxy for market penetration (StatCounter 2013).

¹⁴ In model 10 in the Online Appendix, we substitute an events-based measure of unrest incorporating yearly counts of other types of events (e.g., protests, assassinations), weighted by their severity, replicating our terrorism-based result.

¹⁵ Bormann and Golder (2013) and Johnson and Wallack (2010) both provide alternate measures of district magnitude. The former runs through 2011 but provides district magnitude only for the lowest magnitude tier in an electoral system, regardless of the proportion of seats elected from that tier. The latter runs only through 2005. Thus, neither measure is well suited to use here. Nonetheless, we include these measures in multiple imputation models because they provide useful information about missing values in the DPI measure.

Our analysis focuses on democracies lodging at least one takedown request between July 2009 and June 2012. The Online Appendix discusses our sampling decisions, missing data, and also presents a zero-inflated negative binomial regression that models the selection process that determines if states lodge takedown requests with Google; this robustness check reinforces the results that we present in the main text. Furthermore, because many takedown requests are defamation based, legal institutions and culture may play an important role. In the Online Appendix we examine the robustness of our empirical models to the inclusion of a measure of legal tradition. The Online Appendix also provides descriptive statistics for the variables that we use in our analyses.

Estimation Strategy

Our dependent variable is a count measure that exhibits overdispersion; countries that rarely make content removal requests coexist with states that use this content regulation mechanism regularly. Therefore, we model takedown requests using negative binomial regression (see e.g., Cameron and Trivedi 2005). We include random intercepts for countries in all tests, to model baseline variation in takedown requests across countries.¹⁶ While we would like to use country-specific fixed intercepts to control for cross-national variability in our models, such an approach is not possible here. Because our measures of institutions are largely static, we cannot include institutional factors and fixed unit effects in the same model. Also, given our necessarily short observation period, even our demand measures—patent filings and terrorist incidents—are slowly evolving, rendering fixed effects altogether impractical.¹⁷ We do include fixed effects for time period (half year) in all specifications, to account for differences in incentive to send requests over time.

PREDICTING GOOGLE TAKEDOWN REQUESTS

Table 2 shows the results of negative binomial regression models that predict content takedown requests. In general, the results are consistent with our theoretical expectations about the determinants of government censorship activity. Model 1 includes only demand factors, model 2 adds electoral institutions, while models 3 and 4 replicate models 1 and 2, but include controls. Model 4, therefore, tests all of our hypotheses while controlling for a variety of plausible determinants of censorship activity. Finally, model 5 tests the robustness of our domestic stability findings by substituting the WGI stability/violence measure for terrorist incidents. Figure 1 illustrates the substantive effects of key variables from model 4, predicting number of takedowns in a six-month period, across independent variables' observed ranges.¹⁸ The top panels of the figure display predicted takedowns while the bottom panels show the

¹⁶ Alternatively, one could use clustered standard errors to account for correlated residuals within countries. Models 6 and 7, in the Online Appendix, show that our results are robust to using this alternative approach.

¹⁷ We recognize that relying on random effects is a potential drawback of our analysis, but note that the accumulation of a long panel of internet censorship observations with significant variance in institutions, a necessity for conducting fixed effects analysis, will take many years. Thus, we believe that our approach represents a reasonable compromise given currently available data.

¹⁸ We fit all random effects models using the glmmADMB R package (Bolker et al. 2012), using Rubin's (1987) method to calculate point estimates and standard errors from models fit to multiply imputed data. We generated all predictions and estimates of marginal effects from model 4, and, unless otherwise noted, held all other variables at their means when making predictions. While we report multiple imputation-adjusted maximum likelihood fits in Table 2, we used Bayesian simulation techniques—using uninformative priors—to generate predictions and posterior probabilities. Posterior predictions incorporate information from all imputations.

TABLE 2 *Predicting Google Content Removal Requests*

	Model 1	Model 2	Model 3	Model 4	Model 5
Log terror incidents	0.28* (0.11)	0.27* (0.10)	0.31* (0.11)	0.30* (0.11)	
Log patents per capita	0.58* (0.18)	0.53* (0.18)	0.49* (0.23)	0.44* (0.22)	0.44* (0.22)
WGI stability/violence					-0.79* (0.39)
Log district magnitude		-0.36* (0.18)		-0.36* (0.18)	-0.42* (0.18)
Log GDP per capita			0.05 (0.49)	0.11 (0.48)	0.33 (0.52)
Time to start a business			0.02* (0.01)	0.02* (0.01)	0.02* (0.01)
Internet users per capita			2.10 (2.12)	1.84 (2.08)	2.06 (2.18)
Google search share			-0.01 (0.02)	-0.00 (0.02)	-0.00 (0.02)
Intercept	5.77* (1.67)	6.08* (1.58)	3.63 (5.10)	3.04 (4.98)	1.07 (5.30)

Note: The dependent variable is content requests issued. The observation level is country half-year. We estimated all models using negative binomial regression with random intercepts for country and dummies for half-year time period (not shown). We used chained multiple imputation for missing data; $N = 322$.

WGI = Worldwide Governance Indicators; GDP = gross domestic product.

* $p < 0.05$.

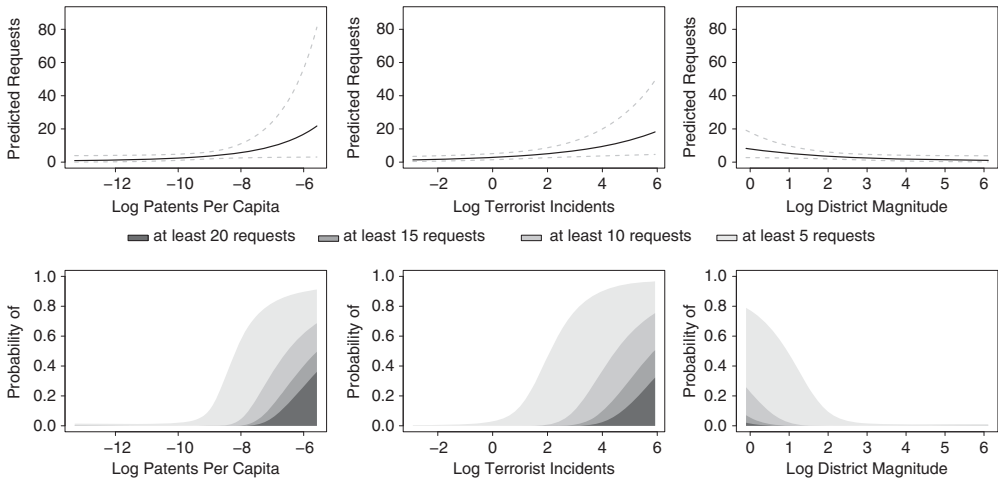


Fig. 1. Model predictions

predicted probabilities of observing at least 5, 10, 15, and 20 requests. Because observed request counts are skewed toward 0 (see the Online Appendix), the point estimates for predicted counts are often quite low, and given the functional form of the model, confidence intervals grow quickly as counts increase. Nonetheless, the predicted probabilities in the lower panels show that the substantive effects of the predictor variables are often quite substantial. For example, an average country with few terrorists incidents is extremely unlikely to lodge more than five takedown requests in any six-month period, while the most terror-plagued states in our data set

have around 97, 76, 51, and 32 percent chances of lodging at least 5, 10, 15, and 20 requests, respectively, in a single period.

All models support a demand-based explanation for digital censorship activity grounded in both regulatory and information-suppressive perspectives. We find a strong relationship between instability and political violence and takedown activity by governments. The coefficient for our measure of political violence—terrorist incidents—is statistically significant in all specifications. Countries experiencing more terrorist violence produce more takedown requests. The model predicts around four takedown requests per period for an average case. Such a case experiences around three terrorist incidents in a period. The model predicts that substantial increases in terrorism would generate commensurate jumps in takedown activity. In particular, if our average case were to experience a 2 SD increase in terrorism—about 47 events—the model would expect to see takedowns jump to around nine per period. The highest observed terrorism rate—373 events—corresponds to a predicted rate of 19 additional takedown requests per six-month interval. The second column of Figure 1 illustrates the substantive influence of terrorism more generally. Furthermore, model 5 indicates that our results are robust to our choice of measure of stability; the WGI stability index is negatively correlated with censorship.

The estimated effect of the private sector's demand for censorship, proxied by the log of patent applications per capita, is also consistently statistically significant and substantively meaningful. High patent countries send more requests. A comparison between South Korea—which produced the data set's highest number of patents per capita in the most recent time period—and Bosnia and Herzegovina—which in the same period produced the lowest number of patents per person—is instructive. During the final observation period South Korea lodged 33 takedown requests while Bosnia and Herzegovina sent only one. Model 3 generates predictions of 45 and one for these two cases, and both true values fall within 95 percent credible intervals for the predictions. Furthermore, the model allows us to explore counterfactual questions. In particular, the model predicts that Korea would have sent around four requests if it produced patents at the Bosnian rate while it predicts that Bosnia would have requested 25 takedowns if it contained knowledge industries as productive as those in South Korea. As Figure 1 illustrates, countries that produce numerous patents are most likely to also lodge many takedown requests. Thus, our results are consistent with the hypothesis that the presence of many IP producers in a state motivates governments to make intensive use of PPCs. Note that only a small percentage of the requests that we measure—see Table 1—directly address issues of copyright and trademark infringement. Thus, while our evidence is indirect, the strong relationship that we observe between IP production and government takedown requests implies that regulatory structures built to service IP producers may allow governments to influence digital information flows more broadly.

Next we examine our argument that government use of PPCs varies across electoral systems. Log of district magnitude taps underlying electoral system responsiveness to concentrated or minority interests and should be associated with increased censorship.¹⁹ Our theory suggests that regulatory and speech censorship makes a majority of the population worse off, while benefiting small groups of constituents (e.g., IP producers or aggrieved or defamed citizens), and that low district magnitude systems should be especially responsive to concentrated interests. The results appear to strongly support this supposition; the relationship between district magnitude and takedown requests is substantively and statistically significant across specifications. While the average effect of district magnitude depicted in Figure 1 is quite

¹⁹ Note that our results are robust to the use of a dummy for single member district systems instead, see models 8 and 9 in the Online Appendix.

modest, the figure—especially the bottom panel of the third column—illustrates that the tendency to make takedown requests drops off precipitously as district magnitude increases. Furthermore, because the model is non-linear, the role of district magnitude can be magnified in states that are otherwise predisposed toward extensive internet regulation. Take Israel, which has an especially large district magnitude (120) and asked Google to take down five or fewer items in each period in our sample. When we use our model to examine the counterfactual question of how many requests Israel would have lodged if it operated under a single member district system, it predicts requests counts ranging between 16 and 45 for the time periods in the sample. While this counterfactual exercise is admittedly speculative, it starkly illustrates the strength of the relationship between electoral system and governmental use of PPCs.

The statistically significant coefficients for time to start a business in Table 2 provide further evidence that institutions matter for internet censorship request activity. How heavily the hand of regulation falls on the brick and mortar economy correlates the with extent of internet regulation. Countries with more red tape censor the internet more actively.

Finally, the coefficients for our remaining control variables in Table 2 are statistically insignificant. Once we account for unit effects with random intercepts and control for key demand factors, countries with higher GDP per capita tend to send takedown requests at about the same rate as poorer countries. Google's search market share also has no clear relationship to requests, with a coefficient close to 0. Google is the dominant search engine across our sample—its average market share is 93 percent, and while the minimum share is 33 and the SD is 9, the first quartile is 92—so this non-result may be attributable to a lack of meaningful variation on this dimension. Finally, while there is a large positive relationship between internet usage and content removal requests, this effect is not statistically significant in our models.

MECHANISM TRACING: THE TURKISH CASE

The previous section corroborates our arguments, but cannot conclusively establish the causal mechanisms that underly our theory. Takedown requests are a recent phenomenon and the empirical record of the use of PPCs is limited. While we can present results on how electoral institutions correlate with digital censorship, we observe no within-country variation in electoral district magnitude during our observation period—we can only establish that states with majoritarian electoral institutions have used takedown requests more aggressively than their counterparts, after accounting for a series of plausible drivers of digital regulatory activity. Similarly, while there is some variation in patent production within countries across the waves of our panel, these shifts are progressive rather than revolutionary, and, even where we observe sizeable changes in patent production, we would not expect politicians, bureaucrats, and legal actors to routinely translate evolving demands from firms into policy in the time frame that we examine here. Nonetheless, these results provide value, both because they describe an empirical landscape that has yet to be explored (see Gerring 2012 for an argument about the importance of “mere description”) and because they serve as baseline tests of the theoretical framework that we develop in the second section, providing a road map for future research on this topic. Yet, one of our hypothesized determinants of internet censorship—internal unrest—does exhibit within-unit variability in our sample. In particular, Turkey represents an exceptional example of recent volatility in terrorism, protests, and other indicators of political instability. We therefore present a short study of this case to help establish the plausibility of a causal relationship between internal unrest and digital speech suppression through PPCs.

Turkey is a state where our theoretical framework indicates strong potential for PPC use. First, while Turkish holds closed-list proportional elections, they have a strong tradition of

personalism, average district magnitude is quite low (seven), and high electoral thresholds amplify majoritarian tendencies. And while internet penetration is relatively low, at 41 percent, Turkey's IP-producing sector is also sufficiently developed—Turkish patent generation is within 1 SD of the sample mean—to generate private-sector demand for digital content policing. Finally, the Turkish state has long-faced challenges to internal stability from the separatist Kurdistan Workers' Party (PKK), which has waged an armed insurgency—punctuated by numerous cease-fires—since 1984.

Turning to recent events, the PKK ordered a cease-fire in April of 2009 (Jenkins 2009), which held until May of 2010 (*PKK Announces Ceasefire in Turkey* 2010). Internal conflict with the PKK remained relatively low during this period, primarily induced by the government's promises of reforms and political and cultural opening (US Department of State 2012, 85). PKK attacks, however, increased in the second half of 2010 and there was a “spike” in attacks and kidnapping in the run-up to, and aftermath of, national elections in June of 2011 (National Counterterrorism Center 2012, 9). Indeed, the Global Terrorism Database lists less than five terrorist events in Turkey in the second half of 2009 and the first half of 2010, but 18 in the second half of 2010, 21 early in 2011, and 32 in the second half of that year.

As the period of relative quiet ended, the Justice and Development Party (AKP) consolidated its control of government, winning its third straight general election and almost 60 percent of the seats in the National Assembly, forming a single party majority government. The empowered Turkish government responded to renewed PKK activity, in part, by ratcheting up its censorship activity, altering media content, and imprisoning journalists using “overly broad and aggressively applied anti-terrorism laws, combined with a judicial system that too often sees its role as protecting the state, rather than the individual” (Corke et al. 2014, 14). This speech curtailment included attempts to regulate internet content through PPCs. Indeed, Google's transparency report includes vignettes from this period that reflect both the AKP's strict enforcement of laws prohibiting “criticism of [Mustafa Kemal] Atatürk, the government or national identity or values,” a broad tool for quashing politically subversive speech, and several more specific examples of attempts to suppress information related to Kurdish activism and independence, including “two requests from a government agency to remove a blog that contains information about the Kurdish Party and Kurdish activities as well as a Google+ picture showing a map of Kurdistan” and requests “to remove blogs for discussing minority independence” (Google Incorporated 2015).

Currently available Google transparency reports begin in July 2009 and extend through the end of 2013, covering nine half-year periods.²⁰ We use these reports to trace the above-described process quantitatively, and use synthetic case-control methods (Abadie, Diamond and Hainmueller 2010) to determine if post-election takedown requests by the Turkish regime are indicative of a policy intervention, or simply an artifact of ongoing cross-national trends in digital content regulation. We use tools described in the study by Abadie, Diamond and Hainmueller (2011) to construct a synthetic control case for Turkey—a weighted amalgam of the other cases in our data set designed to match Turkey as closely as possible, during the pre-intervention period, with respect to a set of covariates that should predict takedown requests.²¹

Figure 2 presents the results of our synthetic control study. The left-hand panel compares pre- and post-election²² takedown requests in Turkey to a weighted average of requests

²⁰ Our previous panel analysis extends only through 2012 because few of our independent variables are available for 2013.

²¹ The Online Appendix provides additional details about the synthetic control study.

²² The election took place toward the end of period 4.

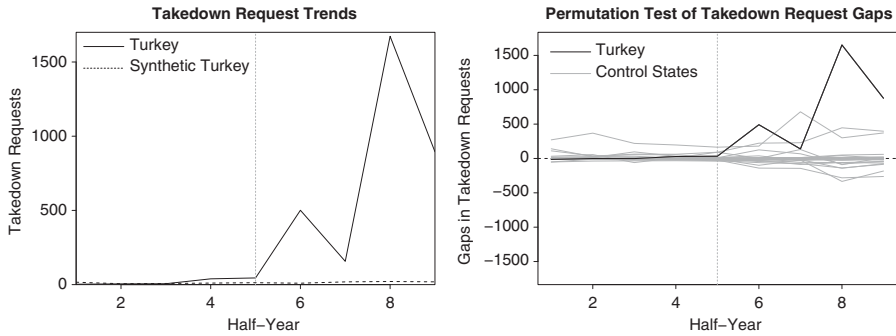


Fig. 2. Synthetic control study results

Note: The left-hand panel compares Turkey's takedown request trend to that of a matched synthetic control case; the right-hand panel plots gaps between cases and matched synthetic controls for Turkey and each country in the potential control sample.

produced by the cases that contribute to the synthetic control. Turkish requests jump substantially in the post-election period before skyrocketing in 2013, when the AKP vastly expanded its censorship activities, likely in light of the Gezi Park protests, but also because of the intensifying of conflict with the PKK associated with the onset of the Syrian civil war (Dombey and Fielding-Smith 2012). The conflict only abated in March 2013 (the end of half-year 8), when the group declared a cease-fire. The synthetic control case maintains a trajectory of slow growth—the scale of the Turkish expansion of takedown requests obscures this modest upward trend—in both the pre- and post-election periods, implying that the jump in Turkish requests reflects a case-specific policy intervention. Corroborating this interpretation, the right-hand panel of Figure 2 displays the result of a permutation test where we iteratively constructed a synthetic control for each country in our sample, and then plotted the gap between each case's takedown requests and their matched synthetic control, across the observation period. The permutation test indicates that the likelihood of seeing an intervention effect of the size that we observe in Turkey is small. Indeed, only two cases show gaps that approach the magnitude of the Turkish case, and one of these cases is poorly matched to its synthetic control in the pre-intervention period. While this synthetic control study cannot definitely establish a causal relationship between internal unrest—initially sparked by the end of the PKK cease-fire, but then fueled by popular protest in response to perceived AKP overreach and a corruption scandal—and takedown requests,²³ it helps to rule out the possibility that the expansion of the use of PPCs in Turkey was simply the result of wider trends in the use of such regulatory tools. This quantitative exercise also lends credence to our interpretation of the qualitative evidence. More broadly, the mechanism tracing that we do here buttresses the large- N analysis that we presented in the previous section, at least with respect to the relationship between internal unrest and takedown requests.

Finally, it is worth noting that, while many current accounts of strong Turkish government pressure on PPCs focus exclusively on the government's handling of the Gezi Park protests in May 2013, “the tools used to pressure and control media outlets and individual journalists existed before the AK Party came to power. But the party, with its extraordinary political dominance, has used them unapologetically and with increasing frequency and force”

²³ Concerns about endogeneity remain. Furthermore, our short observation period limits the viability of the synthetic control method and we must, therefore, be cautious not to over-interpret these results.

(Corke et al. 2014, 8). The case of Turkey demonstrates the relatively fungibility of legal frameworks and institutions that facilitate censorship. In particular, much of the legal framework that would eventually be applied to terrorist organizations in 2011 and then to urban protesters in 2013 have their basis in a legacy of the Turkish military's control and influence on the media, which dates back to before the AKP took power (Corke et al. 2014, 6–7). The events of 2011 represent the confluence of a renewal of hostilities with internal PKK forces and a consolidation of power for the AKP in the aftermath of highly successful elections, but they also reflect the role that pre-existing mechanisms for control play in facilitating censorship by democratically elected regimes.

POLITICS ON THE INTERNET

The primary contribution of this paper is to provide an initial answer to an outstanding, substantively important, question: how do democracies censor the internet, and why (Breindl 2013, 41)? We provide a theoretical framework that points to a pair of factors—political instability and violence, and IP production—that generate demand for digital censorship, and political institutions—namely electoral system design—that translate those demands into government activity. We argue that the answer to the question of when and why democracies censor the internet is, at least partially, political. Furthermore, while previous work has argued that PPCs represent a key content management tool for democracies, this paper provides, to our knowledge, the first large-sample comparative empirical study of what factors drive, and which institutions modulate, this form of digital censorship, showing that patterns in government censorship activity vary systematically with political demand factors and institutions. Finally, we overcome a key obstacle to the comparative study of digital censorship—the lack of cross-nationally comparable measures of this activity—by focusing on Google content removal requests. In sum, this study can provide a theoretical and empirical foundation for subsequent work on this topic. Nonetheless, we do not wish to overstate the strength of our evidence. In particular, data availability limits our ability to subject our causal arguments to strong tests. Our analysis is observational and we currently only have a short panel to work with. We argue only that our theory is consistent with the empirical record to date.

Our study has important implications for facilitating the free flow of information in modern democracies. One key finding is, ironically, that countries that are most invested in the information economy—those with large knowledge-producing sectors—are also those that most actively restrict their citizens' access to information. While this activity helps incentivize IP production, spurring both economic growth and knowledge accumulation, it is important to realize that only a small percentage of the government censorship that we measure in this study actually pertains to IP protection (see Table 1). In fact, the bulk of content removal requests in our data set fall under the rubrics of defamation and privacy and security. Indeed, Google has a system in place to field IP infringement requests directly from private actors, heading off many such IP challenges before the government gets involved. The fact that the marginal relationship between IP production and digital censorship through PPCs is so strong implies that regulatory structures built to satisfy economically motivated constituencies are being leveraged for other purposes. Of course, this paper does not directly trace the mechanisms underlying this argument. Yet, our study highlights an empirical regularity that raises a question that warrants further investigation: do IP-protecting institutions that were designed to promote business interests and to spur innovation allow states to more widely interfere in the free flow of digital information? Digital rights activists, including the Electronic Frontier Foundation (2010), have

long argued that IP-oriented legislation like the DMCA can have unintended consequences. If, as others have argued (e.g., Adler 2011), censorship through PPCs is subject to less oversight and accountability than traditional censorship methods, then citizens of information-rich societies should find our analysis disturbing, especially given that we find that governments also use PPCs more as political instability and violence increases.

Our institutional results contribute to an ongoing debate in political economy about how electoral institutions affect the balance of power between diffuse and concentrated interests. Furthermore, while our empirical focus is on digital content management, our findings imply that electoral institutions may influence the degree of media freedom in society more broadly. We find evidence to support the claim that, because electoral rules modulate politicians' incentives—both to cater to focused interests and to protect their own reputations—we should see more government interference in digital information transmission in low district magnitude systems. But this argument is not tethered to the details of internet regulation and may be applicable beyond the digital domain.

From a policy-oriented perspective, our results imply that citizens in lower magnitude electoral systems face an uphill battle when it comes to protecting their digital rights and thus must work hard to organize to protect consumer interests in free information access. Politicians in low-magnitude systems have incentives to make it easy to remove damaging information from the internet and knowledge-producing firms are ideally situated to obtain protection when elections are low magnitude. It is instructive that one of the biggest wins for consumer advocates of digital rights in the United States—the defeat of Stop Online Piracy Act/PROTECT IP Act (SOPA/PIPA)—was largely organized by content providers who, while substantial IP holders themselves, were concerned that protections that would benefit other knowledge producers would hurt them. Similarly, civil society organizations partnered with content providers to play a critical role in organizing support for Brazil's Civil Rights Framework for the internet. Thus, because organizing average consumers to effectively lobby government is notoriously difficult, one of the most effective strategies that consumer groups in low-magnitude states may have to limit digital censorship is to take advantage of fault lines across knowledge industries. When no fault lines exist, consumer protections in low-magnitude systems are likely to suffer.

A number of intriguing questions for future research emerge from our theory. One of the most timely is: can patterns in content removal requests tell us anything about level of democracy or democratic survival? Could spikes in takedowns be an early warning sign of future autocratic tendencies? On one hand, the answer would appear to be no—many consolidated democracies use these tools extensively. On the other hand, our findings in Turkey provide an intriguing counter example. As we describe, prior to the events of 2013 in Gezi Park and its subsequent crackdown, when many outside observers identified authoritarian trends in the AKP, Turkey progressively moved from a low of between one and ten requests in the second half of 2009, to an incredibly high figure of 501 contacts in the first half of 2012. Perhaps major changes in the extent to which states leverage PPCs are an early warning sign for the erosion of democracy?

Finally, this analysis focuses solely on censorship requests, ignoring compliance. Google provides information on its compliance rate that we hope to investigate in future work. Moreover, other firms—notably Twitter and Microsoft—recently begin releasing transparency reports of their own. Eventually, takedown data from multiple firms will provide a powerful tool for examining the robustness of our findings and cross-firm variations in compliance should help to shed light on when governments are best able to compel private actors to do their digital content regulation for them.

REFERENCES

- Abadie, Alberto, Alexis Diamond, and Jens Hainmueller. 2010. 'Synthetic Control Methods for Comparative Case Studies: Estimating the Effect of California's Tobacco Control Program'. *Journal of the American Statistical Association* 105(490):493–505.
- Abadie, Alberto, Alexis Diamond, and Jens Hainmueller. 2011. 'Synth: An R Package for Synthetic Control Methods in Comparative Case Studies'. *Journal of Statistical Software* 42(13):1–17.
- Adler, Julie. 2011. 'The Public's Burden in the Digital Age: Pressures on Intermediaries and the Privatization of Internet Censorship'. *Journal of Law and Policy* 20(1):231–66.
- Baldwin, Peter. 2014. *The Copyright Wars: Three Centuries of Trans-Atlantic Battle*. Princeton, NJ: Princeton University Press.
- Barlow, John Perry. 1996. 'A Declaration of the Independence of Cyberspace'. Available at [https://homes.eff.org/\\$/sim\\$barlow/Declaration-Final.html](https://homes.eff.org/$/sim$barlow/Declaration-Final.html), accessed 5 April 2012.
- Bawn, Kathleen, and Michael F. Thies. 2003. 'A Comparative Theory of Electoral Incentives: Representing the Unorganized Under PR, Plurality and Mixed-Member Electoral Systems'. *Journal of Theoretical Politics* 15(1):5–32.
- Beck, Torsten, George Clarke, Alberto Groff, Philip Keefer, and Patrick Walsh. 2001. 'New Tools in Comparative Political Economy'. *World Bank Economic Review* 15(1):165–76.
- Bolker, Ben, Hans Skaug, Arni Magnusson, and Anders Nielsen. 2012. 'Getting Started with the glmmADMB Package'. Technical report. Available at <http://glmmadmb.r-forge.r-project.org/glmmADMB.pdf>, accessed 2 January 2012.
- Bormann, Nils-Christian, and Matt Golder. 2013. 'Democratic Electoral Systems Around the World, 1946–2011'. *Electoral Studies* 32:360–69.
- Boyle, James. 1997. 'Foucault in Cyberspace: Surveillance, Sovereignty, and Hard-Censors'. *University of Cincinnati Law Review* 66:177–205.
- Breindl, Yana. 2013. 'Internet Content Regulation in Liberal Democracies: A Literature Review'. Available at http://www.gcdh.de/files/1113/6549/2342/YBreindl_Literature_Review_Mar2013_final.pdf, accessed 9 November 2015.
- Cameron, A. Colin, and Pravin K. Trivedi. 2005. *Microeconometrics: Methods and Applications*. Cambridge: Cambridge University Press.
- Carey, John M., and Mathew Soberg Shugart. 1995. 'Incentives to Cultivate a Personal Vote: A Rank Ordering of Electoral Formulas'. *Electoral Studies* 14(4):417–39.
- Castells, Manuel. 1996. *The Rise of the Network Society*. Cambridge: Blackwell.
- Chang, Eric C. C., Mark Andreas Kayser, Drew A. Linzer, and Ronald Rogowski. 2010. *Electoral Systems and the Balance of Consumer-Producer Power*. Cambridge: Cambridge University Press.
- Corke, Susan, Andrew Finkel, David J. Kramer, Carla Anne Robbins, and Nate Schenckan. 2014. 'Democracy in Crisis: Corruption, Media and Power in Turkey'. Freedom House Special Report. Available at https://freedomhouse.org/report/special-reports/democracy-crisis-corruption-media-and-power-turkey#_VXYCDc9VhBc, accessed 9 November 2015.
- Deibert, Ronald, John Palfrey, Rafal Rohozinski, and Jonathan Zittrain. 2011. *Access Contested: Security, Identity, and Resistance in Asian Cyberspace*. Cambridge, MA: MIT Press.
- Deibert, Ronald, John Palfrey, Rafal Rohozinski, and Jonathan Zittrain. 2008. *Access Denied: The Practice and Policy of Global Internet Filtering*. Cambridge, MA: MIT Press.
- DeNardis, L. 2009. *Protocol Politics: The Globalization of Internet Governance*. Cambridge, MA: MIT Press.
- DeNardis, Laura. 2012. 'Hidden Levers of Internet Control'. *Information, Communication and Society* 15(5):720–38.
- Dombey, Daniel, and Abigail Fielding-Smith. 2012. 'Analysts Link PKK Upsurge to Syrian War'. *Financial Times*. Available at <http://www.ft.com/intl/cms/s/0/dd2c9d8c-ec74-11e1-8e4a-00144feab49a.html#axzz3cflwlnP>, accessed 10 June 2015.
- Electronic Frontier Foundation. 2010. 'Unintended Consequences: Twelve Years Under the DMCA'. Available at <https://www.eff.org/wp/unintended-consequences-under-dmca>, accessed 14 November 2014.

- Elmer-De Witt, Philip, David S. Jackson, and Wendy King. December 6, 1993. 'The First Nation in Cyberspace'. *TIME Magazine*. <http://content.time.com/time/magazine/article/0,9171,979768,00.html>, accessed 11 November 2015.
- Espinel, Victoria A. 2012. '2011 U.S. Intellectual Property Enforcement Coordinator Annual Report on Intellectual Property Enforcement'. Available at http://www.whitehouse.gov/sites/default/files/omb/IPEC/ipec_annual_2011_report.pdf, accessed 9 November 2015.
- Farrell, Henry. 2006. 'REGULATING INFORMATION FLOWS: States, Private Actors, and E-Commerce'. *Annual Review of Political Science* 9(1):353–74.
- Geist, Michael. 2003. 'Cyberlaw 2.0'. *Boston College Law Review* 44:323–58.
- Gerring, John. 2012. 'Mere Description'. *British Journal of Political Science* 42(4):721–46.
- Goldsmith, Jack, and Tim Wu. 2006. *Who Controls the Internet? Illusions of a Borderless World*. Oxford: Oxford University Press.
- Google Incorporated. 2013. 'Transparency Report'. Available at <http://www.google.com/transparencyreport/>, accessed 5 April 2013.
- Google Incorporated. 2015. 'Transparency Report: Explore Turkey Requests'. Available at <http://www.google.com/transparencyreport/removals/government/notes/?hl=en#authority=TR>, accessed 9 June 2015.
- Google Incorporated. 2016. 'Transparency Report'. Available at <https://www.google.com/transparencyreport/removals/government/data/?hl=en>, accessed 10 February 2016.
- Grossman, Gene M., and Elhanan Helpman. 2005. 'A Protectionist Bias in Majoritarian Politics'. *The Quarterly Journal of Economics* 120(4):1239–282.
- Haufler, Virginia. 2001. *Public Role for the Private Sector: Industry Self-Regulation in a Global Economy*. Washington, DC: Carnegie Endowment for International Peace.
- Jenkins, Gareth. 2009. 'PKK: Breaking the Stalemate'. ISN Security Watch. Available at <http://www.isn.ethz.ch/Digital-Library/Articles/Detail/?lang=en&id=100613>, accessed 10 June 2015.
- Johson, Joel W., and Jessica S. Wallack. 2010. 'Electoral Systems and the Personal Vote'. Available at [http://hdl.handle.net/1902.1/17901V1\[Version\]](http://hdl.handle.net/1902.1/17901V1[Version]), accessed 9 February 2014.
- Kaufmann, Daniel, Aart Kraay, and Massimo Mastruzzi. 2010. 'The Worldwide Governance Indicators: Methodology and Analytical Issues'. World Bank Policy Research Working Paper No. 5430. Available at <http://ssrn.com/abstract=1682130>.
- Kaufmann, Daniel, Aart Kraay, and Massimo Mastruzzi. 2012. 'World Governance Indicators'. Available at <http://info.worldbank.org/governance/wgi/index.aspx#home>, accessed 27 January 2013.
- 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(2):326–43.
- Lessig, Lawrence. 1999. *Code and Other Laws of Cyberspace*. New York: Basic Books.
- MacKinnon, Rebecca. 2012. *Consent of the Networked: The Worldwide Struggle for Internet Freedom*. New York: Basic Books.
- Magee, Stephen P., William A. Brock, and Leslie Young. 1989. *Black Hole Tarrifs and Endogenous Policy Theory*. Cambridge: Cambridge University Press.
- Marsden, Christopher T. 2011. *Internet Co-Regulation: European Law, Regulatory Governance, and Legitimacy in Cyberspace*. Cambridge: Cambridge University Press.
- Merriam-Webster. 2014. 'Merriam-Webster Online Dictionary'. Available at <http://www.merriam-webster.com/dictionary/censoring>, accessed 14 November, 2014.
- National Consortium for the Study of Terrorism and Reponses to Terrorism. 2012. 'Global Terrorism Database'. Available at <http://www.start.umd.edu/gtd>, accessed 6 February 2014.
- National Counterterrorism Center. 2012. 'Report on Terrorism 2011'. Available at <https://fas.org/irp/threat/nctc2011.pdf>, accessed 8 June 2015.
- Newman, Abraham, and John Zysman. 2004. 'Transforming Politics in a Digital Era'. BRIE Working Paper No. 165. Berkeley, CA. Available at <http://brie.berkeley.edu/publications/wp165.pdf>, accessed 9 November 2015.
- Persson, Torsten, and Guido E. Tabellini. 2005. *The Economic Effects of Constitutions*. Cambridge, MA: The MIT Press.

- PKK Announces Ceasefire in Turkey. 2010. 'Al Jazeera'. Available at <http://www.aljazeera.com/news/europe/2010/08/2010814131932869254.html>, accessed 10 June 2015.
- Post, David, and David Johnson. 1996. 'Law and Borders: The Rise of Law in Cyberspace'. *Stanford Law Review* 48:1367.
- Reidenberg, Joel. 1998. 'Lex Informatica: The Formulation of Information Policy Through Technology'. *Texas Law Review* 76(3):553–84.
- Rickard, Stephanie J. 2012. 'Electoral Systems, Voters Interests and Geographic Dispersion'. *British Journal of Political Science* 42(4):855–77.
- Rogowski, Ronald. 1989. *Commerce and Coalitions: How Trade Affects Domestic Political Alignments*. Princeton, NJ: Princeton University Press.
- Rogowski, Ronald, and Mark Andreas Kayser. 2002. 'Majoritarian Electoral Systems and Consumer Power: Price-Level Evidence from the OECD Countries'. *American Journal of Political Science* 46:526–39.
- Rosenau, James, and J. P. Singh. 2002. *Information Technologies and Global Politics*. Albany, NY: SUNY Press.
- Rubin, Donald B. 1987. *Multiple Imputation for Nonresponse in Surveys*. New York: J. Wiley and Sons.
- Samuels, David. 2001. 'Incumbents and Challengers on a Level Playing Field: Assessing the Impact of Campaign Finance in Brazil'. *The Journal of Politics* 63(2):569–84.
- Sell, Susan K. 1998. *Power and Ideas: North-South Politics of Intellectual Property Antitrust*. Albany, NY: SUNY Press.
- StatCounter. 2013. 'StatCounter GlobalStats'. Available at <http://gs.statcounter.com/>, accessed 9 November 2015.
- Tonnelson, Alan. 2000. *The Race to the Bottom*. Boulder, CO: Westview Press.
- US Department of State. 2012. 'Country Reports on Terrorism 2011: Turkey'. Available at <http://www.state.gov/documents/organization/195768.pdf>, accessed 8 June 2015.
- The World Bank. 2013. 'World Development Indicators'. Available at <http://data.worldbank.org/data-catalog/world-development-indicators>, accessed 9 November 2015.
- World Intellectual Property Organization. 2013. 'IP Statistics Data Center'. Available at <http://ipstatsdb.wipo.org/ipstatv2/ipstats/patentsSearch>, accessed 9 November 2015.
- Zittrain, Jonathan. 2003. 'Internet Points of Control'. *Boston College Law Review* 44:653–88.