

Limiting the market for information as a tool of governance: Evidence from Russia

Preliminary Draft, June 2017

Klaus Ackermann¹

Abstract

This paper presents a novel measure of subtle government intervention in the news market achieved by throttling the Internet. In countries where the news media is highly regulated and censored, the free distribution of information (including audio and any visual imagery) over the Internet is often seen as a threat to the legitimacy of the ruling regime. This study compares electoral outcomes at polling station level between the Russian presidential election at the beginning of March 2012 with the parliamentary election held three months earlier in December 2011. Electoral regions in two cases are compared: regions that experienced internet censorship at the presidential election but not the parliamentary election; versus regions that maintained a good internet connection without interference for both elections. Internet censorship is identified using randomised internet probing data in accuracies down to 15-minute intervals for up to a year before the election. Using a difference in difference design, an average effect of increased vote share of 3.2 percentage point for the government candidate is found due to internet throttling. Results are robust to different specifications and electoral controls are used to account for the possibility of vote rigging.

¹ackermann@uchicago.edu: Center for Data Science and Public Policy, University of Chicago; Department of Economics, Monash University

1 Introduction

Free distribution of information is generally believed to enrich the political process and improve participation during elections. An important source of information access for citizens is the Internet. In countries with poor democracy and low levels of “voice and accountability”² the Internet appears to be the only reliable source of information. As such, if government interventions hinder the free operation of the Internet, citizens can lose access to reliable information. Evidence suggests that some governments are tampering with the internet to limit access to information for their citizens; an apparently unobservable state intervention which is increasingly common in autocratically leaning jurisdictions. An example of this activity has been reported in Iran by the international organisation Reporters Without Borders (2012, p. 3):

“Shutting down the Internet is a drastic solution that can create problems for the authorities and can hurt the economy. Slowing the Internet connection speed right down is more subtle but also effective as it makes it impossible to send or receive photos or videos. Iran is past master at this. Syria’s censors also play with the Internet connection speed, fluctuations being a good indicator of the level of repression in a given region.”

Tampering with the Internet connection of a citizen is a human right violation. The United Nations Human Rights Council passed the resolution A/HRC/32/L.20 affirming that people deserve the same rights when online as are granted offline. Especially, the intentional disruption of information flow is a violation of international human rights law.

From an economic perspective information resides at the core of social choice (Sen, 2008). Each economic agent, be it a consumer or a voter, has an information set available as a source for decision making. In an election, citizens reveal their preferences as a collective society and institutions that arise from this should reflect their collective

²<http://info.worldbank.org/governance/wgi/pdf/va.pdf> “Voice and accountability captures perceptions of the extent to which a country’s citizens are able to participate in selecting their government, as well as freedom of expression, freedom of association, and a free media.”

intentions.

It is proposed that the market mechanism of revealed information that would allow an updating of institutions is especially distorted in the Russian Federation. Evidence surveyed in Enikolopov, Petrova, and Sonin (2013) report that parts of the Russian government are actively involved in Distributed Denial of Service (DDoS) attacks against blogs which reveal corruption in state-owned enterprises. Besides, the open net initiative³ reported that during election time subtle internet control measures (including speed tampering) were active.

There is documented evidence of electoral fraud in previous Russian elections. At the Duma elections of 2011 Enikolopov, Korovkin, Petrova, Sonin, and Zakharov (2012) conducted a field experiment randomly assigning independent voting observers to polling stations in Moscow. They found that the presence of an observer reduced the voting share for the incumbent government party. The treatment group had on average 11% lower results than the control group without observers, under the assumption that vote rigging was only conducted in favour of the government, and not opposition parties.

Besides relying on experimental data, post-electoral data revealed ballot box manipulations. A statistical election fraud detection was also conducted by Klimek, Yegorov, Hanel, and Thurner (2012). Using the assumption that a polling station within a country represents a sample taken out of the general population, they hypothesise that the large number of samples should follow an approximate normal distribution. One key identifying feature of election fraud that was documented in Russia where polling stations with above 90% voter turnout and a 90% vote share for the incumbent.

Meanwhile, the rise of new media is changing political activism. The Russian Federation has experienced an increase in social media internet use, in particular through a Russian only online platform called VK, and this greater capability to organise has lead to an increase in the probability of citizens taking part in protests as shown by Enikolopov, Makarin, and Petrova (2015).

³<https://opennet.net/research/profiles/russia>

It is important to distinguish between the different possible methods of how the sovereign might manipulate information distribution. The type of internet censorship studied in this paper relies on throttling the rate of data throughput rather than the blocking of specific websites. Blocking of content is not a standard tool for the Russian government in its exertion of influence over the internet as highlighted by Franke and Pallin (2012). It relies more on the creation of a political environment, which fosters self-censorship amongst internet service providers. Furthermore, if the Internet surprisingly becomes slow making all content less accessible (whether it is content within the country or possible blocked content from outside using virtual private networks) individual citizens are more likely to believe that internet overuse or poor infrastructure is at fault rather than the government in power.

This project started in 2013 with a hypothesis that contrary to the popular belief that the internet has always been a bastion of free information, autocratic states were managing to manipulate information flows on the internet in ways too subtle to be publicly criticised. Although government intervention into available internet bandwidth had been suspected since 2007 by internet watchdog organisations, at the time of hypothesis no data existed to prove or reject these claims.

It was proposed that active interference should be visible through examining global data sets of Internet Control Message Protocol (ICMP) probes. A single probe, usually referred to as ping, is a test packet sent from one device to another device over a network. Each probe returns a round trip time in milliseconds, measuring how much time is taken for information to physically traverse between devices. A fast connection has a lower response time, while slower connections take longer. Sending a series of those probes to a target IP over time provides a time-series of response measurements, which allows state speed throttling intervention to be detected by comparing distortion to expected values across a time-series.

The contribution of this paper is twofold. Firstly, it shows how automatically collected internet data can be used to look beyond the stated democratic principles of a sovereign,

and discover their real (revealed) actions. Secondly, an attempt has been made to estimate the causal effect benefits to a government in power on election day brought through these hidden actions

2 Literature & Background

Previous attempts have been made to explain the relationships between news sources and the voting patterns of citizens, as well as studies on the effects of internet availability and censorship by governments.

Classical print media historically enhanced the political diversity of democracies. Gentzkow, Shapiro, and Sinkinson (2011) examined the impact of newspaper and television broadcasting on election outcomes in the US. They found that an increase in available information increases the election turnout, whilst a lack of information availability has a dampening effect. The competition in the newspaper market, with new newspapers entering and others exiting, did not favour any specific party or candidate. Similarly, Gentzkow and Shapiro (2010) argued that newspapers respond to their readers' political attitude by adapting to consumers preferences in the US.

Evidence from Russia indicates a contrasting situation to observations made in the US. For instance, Lipman and McFaul (2005) report that the Russian presidents' agenda after the election in 2000 was to suppress or abandon media with content critical of the ruling government.

One sided media attention can significantly change the election as the elective representative does not necessary represent the efficient market outcome. The effect of unbalanced information access was studied by Enikolopov, Petrova, and Zhuravskaya (2011). They focused on availability across regions of the only national television channel independent from the state to measure the effects on voting habits of citizens exposed to alternative information. It was shown that exposure to the television channel increased the probability of voting for the opposition in Russia with the effect powerfully demon-

strated in the parliament elections of 1999. By 2003 the independent television channel had been bought by a company closely related to the state and no measurable persuasion effect remained.

The economic literature provides possible pathways for how autocratically leaning governments could respond to a change in possible information flows. A theoretical model of the impact on mass media on peoples preferences and political choices was done by Petrova (2008). In the empirical section of their study a relation between internet penetration and censorship in autocratic states was highlighted, showing that an increased uptake of uncontrollable media (such as the internet news sources) amongst the community could lead to a tightening of restrictions on the controllable media (e.g. print). The author relied on the World Development Indicator Database as their source of Internet penetration.

It is important to compare studies on the effect of internet penetration based on democratic development. In established democracies, such as Germany, the introduction of the Internet initially, lead to a decrease in voter turnout especially during the switching phase as modems were upgraded to broadband from 2004 through 2008. Falck, Gold, and Heblich (2014) related this phenomenon to entertainment consumption was increasing when people switched from television to internet use. To identify the effect they constructed an instrument based on telephone line availability to assess the effect of the faster internet on elections. At that stage, an election deciding influence due to an increase in available information access was not found.

A more nuanced picture of these effects was observed in a study of broadband introduction staged over a longer period in Italy. The study, conducted by Campante, Durante, and Sobbrino (2013), showed that at first, the introduction of broadband decreased voter turnout, but the effect was later reversed as the technology became more established. The authors relate the change to the rise of new political actors using the internet as a primary source of activism content. Like in the previous study, the causal estimates again relied on the distance of a region to a high-speed hub.

The impact of an established technology on the election was also confirmed in Brazil. Menezes (2015) discovered similar effects in the presidential election 2010, with candidates from smaller parties benefiting most. Those who previously did not have the resources to advertise to broad demographics had a significantly raised vote share in municipalities with a high internet penetration. Geographical variation and the cost of deploying infrastructure was used for identification purposes.

Differing results from internet expansion have been seen in more repressive regimes. An early great contribution is Miner (2015), which examined the spread of the Internet in Malaysia from 2004 to 2008 finding that the information increase benefited opposition parties through higher voter turnout and a higher vote share.

Careful consideration of the generalizability of these findings is warranted however. The author utilised possible Internet Protocol (IP) address space as a proxy for penetration, while the previously mentioned studies used reported physical line connections as a proxy. The possible IP space does not necessarily give an accurate reflection of possible actual use due to stale IP ranges as pointed out in (Ackermann & Angus, 2014). The author makes use of the *Maxmind* geographical location database for IP addresses to derive a measurement of internet penetration over time. The quality of this database is poor⁴, accordingly, Miner takes a few different approaches to try to reduce the error rate by combining datasets with the Asia-Pacific Network Information Center, which is responsible for delegating IP addresses to internet service providers. This methodology does not however account for both the technical procedure behind IP delegation as well as the market allocation issues discussed in Kuerbis, Asghari, and Mueller (2013). The reservation of an IP space for an Internet service provider does not necessary imply a representation of new customers who gained additional access to the internet. Furthermore, there is no guarantee that the new IP space is assigned homogeneously throughout a region.

There has been diverse literature from disciplines such as economics and geographical

⁴As of July 2016 only 30% of the data entries are correct as stated on their website in cities: https://www.maxmind.com/en/city_accuracy

sciences, which has focused on Internet diffusion and its impacts on society. Li and Shiu (2012) provide a summary in the appendix of their paper enumerating studies which have estimated internet effects at a national level. To explain the spread of the Internet within a country, these studies generally use indicators such as national telecommunication infrastructure or measures of competition within the telecommunication market. Kolko (2009) made use of official US Internet service provider registration forms (FCC form 447), which provide the number of customers per postcode to derive a measure of differentiated broadband availability. A few years later the same author used the dataset to estimate local growth based on broadband availability, using population growth and the slope of the local terrain as an instrument (Kolko, 2012). Another approach was taken by Mack and Rey (2014) who also used the FCC postcode data but created a spatially dependent structural model using lagged variables of broadband and population density as the instrument.

Research on existing literature on the role of media on election outcomes shows that any analysis which hopes to define the impact of the Internet needs to differentiate between democratic and more authoritarian government structures. A time lag appears in mature democracies between the internet being used as a distraction from politics and the new medium enriching the political participation in the same way as traditional print media. On the other hand, in less developed democracies use of the Internet by political opponents to the government is often quicker than traditional media thus seen as a threat to the ruling class and often censored heavily. The four mentioned papers focused on the introduction period of broadband technologies and used the possible connections as a proxy for Internet consumption. In contrast, this paper focuses on regions which have already had internet infrastructure in place for years yet experience active limitations.

Since 2007 the Open-Net Initiative identified that tampering with the speed of the internet service is an almost undetectable but widely suspected tool of media access control by nation states. Due to the metrics used in previous internet studies, comparisons of bandwidth availability have been unable to shed light on an issue this nuanced. The

majority of work thus far has focused on additional information gain from internet availability with data often limited to yearly intervals. This type of data can give a rough estimation of the quantity of internet connections being made, however, are unable to give any insight into the quality of connection available during the year, or before the month leading up to the elections. Research has mainly been limited to areas of political activism and the political engagement of the general population rather than any state targeted interventions.

Given the impact that slower internet speed can have on the general population during the lead up to an election, it is important to be able to identify state intervention on this level. Slower internet speeds reduce the usability of the internet, especially by making videos and pictures (a strong component of social media). Speed throttling also reinforces and obscures obvious state censorship to VPN's and other bypasses (TOR, screen forwarding, etc), however not giving away why it is not working. In this way citizens have their information channels shifted to public media and news channels within the possible influence sphere of the government without realising they are being persuaded this way.

3 Identifying Censorship

3.1 Data processing

The dataset created in this paper is similar to the previous paper (Ackermann, Angus, & Raschky, 2016) with some key distinguishable properties. It utilized two types of internet scans carried out as described in Heidemann et al. (2008). An internet census, which scans every IP address of the IPv4 Internet, as well as a survey approach, where a representative IP address is selected out of every 254 address blocks and then scanned repeatedly every 11 minutes over a period of weeks. Representative IP addresses are selected based on their historical responsiveness in previous scans (Fan & Heidemann, 2010).

The scans are conducted from various locations globally, which provides a series of

measurements between devices worldwide. The source locations used in this paper are as follows: US Ft. Collins in Colorado, Arlington in Virginia and Marina del Rey in California. In addition, some operations were conducted from Fujisawa-shi, Kanagawa in Japan. The period of these census operations followed a random pattern, census scanning could be carried out at the same time from different locations, while survey or re-probing scans were only operated from one place at a time to maintain data integrity. These collaborated measurements allowed a detailed picture to be formed of internet activity in a region over time, and in this manner, it became possible to construct a signal for detecting censorship in an area of interest. Each location has a different base time offset (additional response time as the signal originates from Japan versus the US), a measurement limited by the physical condition of the local network infrastructure. Because ping measurements take the time between the source and destination, and not a measurement of each intermediate router, it is unknown where along the route any deviation to standard response time is occurring. In general, international routing protocols will send packages on the fastest known path available between two devices.

Each IP probe taken was geolocated to a particular region in addition to being associated with the relevant Internet Service Provider (ISP). Measurements within a region from the same ISP are assumed to experience similar speeds in a given location as the same physical infrastructure is used. More than 1.5 trillion raw observations were made, and these were aggregated down to 15-minute intervals⁵, logging latitude and longitude location, ISP, scanning origin and scan type (survey or full census). The approach taken to calculate average speeds was to aggregate all positive (online) ping responses by summing up response times, and counting the number of online probes within the 15-minute interval. This allows the calculation of average response times on different levels, as it leaves data raw enough to manipulate results on both geographical and time dimensions. After filtering, a total of 226,663,017,494 online responses and 1,052,728,499,987 offline

⁵This aggregation step was performed using Apache Hadoop on the computational facilities of the Australian Synchrotron. The successful aggregation step consumed roughly 50000 cpu hours, which would take a single computer 5.8 years to complete.

responses, globally, are present for further analysis.

3.2 Detecting of Internet Censorship

The first step to identifying internet censorship after identification of a country of interest is to calculate average round trip time rtt , with data aggregated by day t , country c , region i and scanning source s . To avoid false positives in the detecting of censoring, there needs to be a control taking into account world wide internet traffic events. For this reason, the worldwide average is calculated at the same time, utilising data from all countries using ISP in place of a region to derive an average measure. The rationale behind this distinction is that while regional political boundaries are important in detecting censorship by sovereign states, these regions are not necessarily representative of the existing network architecture. Furthermore, they are not needed to identify increased traffic consumption which could explain disruptions at the state level.

A few measures are taken to control for day specific country effects, which could distort the average round trip time into a country of interest. For each time-series created along those previously mentioned dimensions the first and last day of a consecutive period are dropped to avoid distortions due to possible unequal sampling. The first stage is the fixed effect regression

$$\log(rtt_{icst}) = \alpha_t + \beta_s + \gamma_c + \epsilon_{icst} \quad (1)$$

where α_t is a dummy for a day fixed effect, β_s a dummy for the scanning source, γ_c a country dummy and ϵ_{icst} the residual of unexplained variation by region and day in a logarithmic scale. The time series of residuals is converted into a stationary series by taking the first difference to eliminate the level difference between the scanning sources. $\Delta\epsilon_{icst}$ now becomes a signal of average response time by day and geographical location of interest. When a sovereign state tampers with the connection speed within a region, this would cause a distortion to the connection speed. It is proposed that Internet censoring

occurs on a day which experiences a shift greater than three standard deviations from the average $\Delta\epsilon_{icst}$ over a time window.

The detection mechanism described is dependent on periods where no internet censoring occurs in order to document a baseline expected value. The calculation of the standard deviation is based on the location of the scanning origin as well as the time frame over which the scan was taken. In comparison, figure 1 rests on a collection of data with an order of magnitude more observations than that shown in figure 2, which creates a more noisy signal. The advantage of this extra dimension is the presence of scans taken of the same location at the same time from different sources. This allows an integration of the signal using majority rule of agreement to determine whether a day is defined as censored or not, reducing the possibility of false positive detection.

3.3 Validation of the detection mechanism in Tehran, Iran

Disrupting information flow of the internet is a phenomenon prevalent in Iran. Previous research by Anderson (2013) used a dataset focusing on file-sharing data to detect speed throttling as a form of censorship. This speed throttling was later confirmed by Iran's minister for communications and information technology as having taken place during the Iranian election in 2013 (Esfandiari, 2013). This provided an ideal event to test the censorship detecting capabilities of the data set constructed for this paper, as well as the methodology proposed in the last section.

Anderson (2013) used test data collected through the file-sharing platform Bit-torrent to identify periods of heavy internet throttling in Iran. The data was collected in the background when internet users downloaded data of any kind. On average, the number of clients varied between 500 and 1000 users per week during a period from 2010 to 2012. The author defined thresholds during times of non-expected censorship and compared those measurements against times of expected government intervention.

The Iranian governments' attitude towards freedom of expression over the internet has historically been revealed in the lead up to expected protests. The 14th of February 2012




Figure 1: Survey: Average ping times and world traffic corrected difference residuals in Tehran and a mobile phone ISP pre and post the anniversary of the green movement. Based on 6,741,895 survey probes.

marked the one year anniversary of protest of the green movement, at which time the Coordination Council for the Green Path of Hope had called for a march of silence (*Iran: Widespread arrests reported on day of protest*, 2012). Figure 1 illustrates the average ping times by scanning location and $\Delta\epsilon_{icst}$ corrected by world traffic. A vertical line marks the Green path of hope anniversary. In the days before the suspected protest against the government in power, international news media, such as the Washington Post, reported limits to the availability of internet bandwidth (Flock, 2012). Specifically, a blog post marked the 9th of February as a day of substantial limitation (Petrossian, 2012). This event coincided with a massive increase in average response times and corresponded to the orange peak shown in figure 1. The same shift is also visible when observing census probes in isolation as displayed in figure 2.

In comparison, Anderson (2013) was required to use week and month averages for the entire country due to limitations in his available data. A quirk of his method of



Figure 2: Census: Average ping times and world traffic corrected difference residuals in Tehran and a mobile phone ISP pre and post the anniversary of the green movement. Based on 596,733 census probes.

measurement meant that upon the advent of internet throttling, file-sharing users were often motivated to postpone their activities until speeds were reasonable again. This limited the number of available clients available for data collection during suspected throttling events. The method described in this paper is advantaged by use of actively sent probes to IP addresses, allowing an almost instantaneous resolution down to the city and ISP level. The investigations using probing data confirmed Anderson’s suspicion that censorship in Iran occurred on an ISP specific level. Smaller Internet Service Providers, as well as companies large enough to act as their own ISP, appeared not to be impacted. Worth noting in the graph is that scanning data does not cover all days, and as such if an event of interest does not fall within a period of scanning, no conclusive finding can be reached. This method is optimised towards only detecting true positive and avoiding false positives.

4 Data and Estimation

A critical time in any political actors calendar is in the lead up to an election. Politicians and political parties try to disseminate information about their views, past achievements and future plans to as wide an audience as possible. For an autocratic leaning country interference with the electoral process remains a convenient method of legitimising claims to power. While censorship of the mainstream media within a country is often assumed, censoring of the internet has been perceived to be harder to carry out due to the inherent levels anonymity and the vast diversity of potential information sources. It is here that data throttling can play a vital role for a dictatorship.

Electoral data is available at different geographical levels in the Russian Federation. Freely available from www.vybory.izbirkom.ru, data is available right down to polling station outcomes, or precinct electoral commission information. The data used for this paper was collected from the Duma election in 2011 and then three months later from the Presidential election 2012. Each polling station was then manually matched to the global administrative units⁶ on sub-region level 2 (adm2). The collected internet data described in the previous sections was also aggregated at the adm2 level in order to allow a coherent comparison. The unit of analysis was the polling station level, with treatment being defined at the sub-region level.

A difference-in-difference (DID) approach was used to compare electoral outcomes at precinct electoral level (between sub-regions) which experienced an active intervention into their internet information flow at the presidential election versus regions that had unrestricted access to the internet. For each adm2 region to be included in the selection for either treatment or control, it was a requirement that the state (adm1) that the sub-region is located in had possessed consistent IP address space since 2006 as defined in Ackermann, Angus, and Raschky (2017).

The internet use leading up to the election was defined as the time frame of interest, with a selected baseline for the calculation being the start of 2011 right up until a day

⁶www.gadm.org

before the presidential election on the 4th of March 2012. The parliament election day on the 3rd of December 2011 has been excluded.

The time range for detecting possible intervention was defined as from the start of November for the Duma election and the time period between the Duma vote cast up until the day before the polling for the presidential elections. The census scan was combined using a majority rule, requiring there be at least 2 scanning locations are observing on any given day. If there was a distortion of three standard deviations from the average speed by 2 out of 2 or 3 scanning locations, then a region was defined as having suffered from internet censorship. The survey scan data, which suffers less from noise and overlapping, requires one outlier to contribute to the treatment assignment. There is the possibility that bad internet connection is prevalent in the control groups, and to avoid false positives sub-regions which crossed the two standard deviations border more often than twice from the average were dropped. The rationale behind this was that if the speed fluctuated more than once but never broke the three standard deviation border then it was likely that the region suffered from poor quality internet infrastructure. Figure 4 and 3 present the average ping time and the corrected first difference in Orenburg, Russia.

As outcome variables the voting share of a state-supported party candidate was used. Of interest were voting shares of candidates whose parties obtained a seat in the 2011 parliamentary election. It was assumed that minimal average swings would be experienced within such a short period of time between elections. The three-month difference was controlled for by time and sub-region variation. The difference between sub-regions pre and post treatment of internet censorship relative to the control group was defined by:

$$Y_{irt}^j = \alpha + \beta_1 \cdot censorship_r + \beta_2 \cdot post_t + \beta_3(censorship_r \cdot post_t) + \epsilon_{irt}^j \quad (2)$$

where Y_{irt}^j is the voting share of the party or candidate j , the precinct electoral commission i , $censorship_r$ denotes if a sub-region experience censorship, $post_t$ a time trend which is 0 for the parliament election and 1 for the presidential election and $(censorship_r \cdot post_t)$ is




Figure 3: Survey: Average ping times and world traffic corrected difference residuals in Orenburg. Based on 4,382,231 survey probes.

the interaction of interest. The coefficient β_3 is the average treatment effect of censorship at the polling station. Standard errors are clustered at sub-region level, to account for serial correlation and to mitigate possible overstating of the statistical significance as emphasised in Bertrand, Duflo, and Mullainathan (2004). To account for other possible hidden government interventions designed to sway voters, the second specification, is estimated with controls that vary between sub-regions and time periods:

$$Y_{irt}^j = \alpha + \beta_1 \cdot censorhip_r + \beta_2 \cdot post_t + \beta_3(censhorhip_r \cdot post_t) + \beta_4 \cdot X_{irt} + \epsilon_{irt}^j \quad (3)$$

where X_{irt} are the electoral controls, precinct election turnout and eligible voters by




Figure 4: Census: Average ping times and world traffic corrected difference residuals Orenburg. Based on 509,908 census probes.

square kilometres, as well as economic indicators, all in logarithmic scale. A higher density implies a tighter social network compared to less dense areas, which would allow the government to interfere with different channels without the need for election tampering. As pointed out by Klimek et al. (2012), suspiciously high voter turnouts are a signal for ballot stuffing; thus measures are included to account for this possibility. In maintaining consistency with existing economic literature regarding elections in Russia (Enikolopov et al., 2011), income per household, and active physicians per thousand people were used to reflect the strength of public infrastructure. As the estimation strategy relies on a difference in difference design, the controls needed to be time varying. The economic data from the OECD Regional Accounts was available at adm1 level by year and has

been lagged. To capture the effect of differences in internet infrastructure by region, lagged IP per capita was also included. A robustness check was added, dropping all sub-regions containing both high 90% turnouts as well as above 90% voting share for the incumbent in addition to dropping the list of regions mentioned in Kobak and Shpilkin (2016), whose historical pattern of reported election results contained irregularities. This step was to account for possible election manipulation outside of the scope of the control variables.

4.1 Balance test

To assess the comparability between the control group of 82 sub-regions and treatment group of 32 sub-regions, a balancing test on the outcomes and covariates was performed. All parties that won seats in the parliamentary election in 2011 and nominated a candidate for the presidential election were included. The incumbent party was “United Russia (UR)”, with Vladimir Putin as its head, and he was the candidate who ultimately triumphed. The Opposition parties include the “Communist Party of the Russian Federation (KPRF)” led by Gennady Zyuganov, the party “A Just Russia (JR)” headed by Sergey Mironov and the “Liberal Democratic Party of Russia (LDPR)” with Vladimir Zhirinovskiy as its leader. Mikhail Prokhorov was another candidate on the ballot sheet for the presidential election, but he ran as an independent. Table 1 presents the balance test at baseline with standard errors clustered at sub-region level. The voter turnout was significantly different at the 10% level indicating that this control variable, when applied to all regions, was not balanced. Conversely, in the robustness sample with areas with suspicious election results dropped (60 versus 26 sub-regions), the difference was no longer statistically significant as shown in table 4. Furthermore, the internet infrastructure variable regarding IP per capita differed across regions. Concerning outcome variables, only the vote share of “A Just Russia” differed a statistically significant amount between the control and treatment group.

Table 1: Balance test

Variable	Control	Treatment	Diff	p-value
Vote Share UR	0.496 (0.203)	0.450 (0.196)	-0.046 (0.039)	0.238
Vote Share KPRF	0.200 (0.095)	0.192 (0.087)	-0.008 (0.017)	0.631
Vote Share JR	0.132 (0.073)	0.162 (0.084)	0.029 (0.015)	0.056
Vote Share LDPR	0.129 (0.069)	0.143 (0.058)	0.013 (0.009)	0.137
Election Turnout at PEC	0.619 (0.174)	0.575 (0.147)	-0.045 (0.025)	0.082
Election Turnout at PEC (log)	-0.518 (0.278)	-0.583 (0.237)	-0.065 (0.041)	0.109
Eligible voters by km2	4.522 (2.115)	4.170 (2.347)	-0.352 (0.558)	0.530
IP per capita	-4.431 (0.528)	-4.182 (0.429)	0.249 (0.114)	0.032
Active Physicians Rate by 1000	1.510 (0.188)	1.491 (0.207)	-0.019 (0.051)	0.708
Primary Income Private Households per head	8.557 (0.339)	8.618 (0.318)	0.061 (0.074)	0.415

Note: Standard error of the mean and difference is clustered at sub-region level.

4.2 Empirical Estimation

The regression results suggest that repression of information flow by the government has a significant effect on election outcomes. Equation two provides the primary estimate, which represents the voting share for "United Russia" or Vladimir Putin in table 4, indicating that on average the censorship of internet connections in a region increased vote share of the incumbent by around 3.2 percentage points. In the full sample, without accounting for electoral fraud as in table 3 the effect is around 2.8 percentage points. The results comparing the coefficient concerning voter turnout across the full sample and the robust specification, it appears that dropping areas suspected of dubious voting results acted to reduce the estimated influence of turnout. The result difference between the standard and controlled difference estimation, is similar to the robust sample, indicating that this is an outcome of the balance between treatment and control as mentioned in (Abramitzky & Lavy, 2014). Whilst the communist parties voting base, shown in equation

Table 2: Balance test restricted Sample to account for vote ridding

Variable	Control	Treatment	Diff	p-value
Vote Share UR	0.428 (0.138)	0.391 (0.123)	-0.037 (0.029)	0.206
Vote Share KPRF	0.229 (0.087)	0.208 (0.068)	-0.021 (0.017)	0.230
Vote Share JR	0.152 (0.063)	0.187 (0.074)	0.035 (0.016)	0.039
Vote Shae LDPR	0.148 (0.060)	0.155 (0.049)	0.007 (0.009)	0.411
Election Turnout at PEC	0.562 (0.136)	0.544 (0.114)	-0.019 (0.019)	0.335
Election Turnout at PEC (log)	-0.603 (0.235)	-0.630 (0.201)	-0.027 (0.035)	0.444
Eligible voters by km2	3.954 (1.962)	3.998 (2.081)	0.045 (0.646)	0.945
IP per capita	-4.350 (0.534)	-4.010 (0.387)	0.341 (0.140)	0.018
Active Physicians Rate by 1000	1.577 (0.160)	1.502 (0.210)	-0.075 (0.060)	0.211
Primary Income of Private Households per head	8.528 (0.333)	8.571 (0.259)	0.043 (0.082)	0.603

Note: Standard error of the mean and difference is clustered at sub-region level.

4, appears to have been unaffected by Internet censorship, an adverse impact is shown for “A Just Russia”, however, this could be a result of baseline imbalance, requiring further investigation.

In order to complete a robustness check, it is necessary to include additional estimations. The tampering effect could differ across sub-regions due to unequal numbers of eligible voters. To account for the implications of this, weighted regression of eligible voters is performed using block bootstrapped standard errors on the robust samples in table 5. This resulted in a slight increase in the estimates. The weighted regression procedure does not have an option for clustered standard errors, which led to the use of block bootstrapped errors in order to minimise possible overconfidence in the results following procedures described in Bertrand et al. (2004)

Table 3: Full Sample

VARIABLES	Party share							
	(1) UR	(2) UR	(3) KPRF	(4) KPRF	(5) JR	(6) JR	(7) LDPR	(8) LDPR
Difference in Difference Estimator	0.038* (0.019)	0.028* (0.015)	-0.008 (0.008)	-0.004 (0.006)	-0.024* (0.012)	-0.021* (0.011)	-0.012** (0.006)	-0.009* (0.005)
Treatment	-0.046 (0.039)	-0.010 (0.024)	-0.008 (0.017)	-0.021* (0.012)	0.029* (0.015)	0.021* (0.012)	0.013 (0.009)	0.004 (0.007)
Timetrend	0.158*** (0.010)	0.147*** (0.011)	-0.022*** (0.005)	-0.012** (0.006)	-0.097*** (0.005)	-0.094*** (0.006)	-0.061*** (0.003)	-0.060*** (0.005)
Election Turnout at PEC (log)		0.403*** (0.034)		-0.162*** (0.016)		-0.099*** (0.009)		-0.091*** (0.010)
Eligible voters by km2		-0.018*** (0.004)		0.007*** (0.002)		0.005*** (0.001)		-0.001 (0.001)
IP per capita		-0.046*** (0.014)		0.015* (0.008)		0.012 (0.009)		0.013** (0.005)
Primary Income of Private Households per head		-0.068*** (0.023)		0.006 (0.010)		0.013 (0.012)		0.006 (0.007)
Active Physicians Rate by 1000		0.066 (0.042)		-0.043* (0.022)		-0.005 (0.016)		0.020* (0.011)
Observations	13,834	13,834	13,834	13,834	13,834	13,834	13,834	13,834
R ²	0.204	0.578	0.024	0.313	0.469	0.601	0.268	0.418
Number of sub regions	108	108	108	108	108	108	108	108
F-statistics	125.9	169.2	15.27	26.62	146.8	104.8	192.3	164.2

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Results based on 2935544223 observations sent from beginning of 2011 to 4th of March 2012

Table 4: Robust sample free of suspected electoral fraud related to counting of the votes

VARIABLES	Party share							
	(1) UR	(2) UR	(3) KPRF	(4) KPRF	(5) JR	(6) JR	(7) LDPR	(8) LDPR
Difference in Difference Estimator	0.032* (0.018)	0.032** (0.016)	-0.002 (0.009)	-0.002 (0.009)	-0.028** (0.014)	-0.028** (0.013)	-0.011* (0.006)	-0.011* (0.006)
Treatment	-0.037 (0.029)	-0.029 (0.022)	-0.021 (0.017)	-0.022 (0.016)	0.035** (0.016)	0.035** (0.014)	0.007 (0.009)	0.004 (0.008)
Timetrend	0.190*** (0.009)	0.181*** (0.010)	-0.029*** (0.006)	-0.019* (0.010)	-0.114*** (0.007)	-0.106*** (0.009)	-0.070*** (0.004)	-0.071*** (0.005)
Election Turnout at PEC (log)		0.167*** (0.028)		-0.065*** (0.018)		-0.051*** (0.011)		-0.037*** (0.012)
Eligible voters by km2		-0.019*** (0.004)		0.009** (0.004)		0.002 (0.002)		-0.002 (0.002)
IP per capita		-0.008 (0.019)		-0.000 (0.013)		-0.011 (0.011)		0.012 (0.009)
Primary Income of Private Households per head		-0.049** (0.021)		-0.012 (0.016)		0.019 (0.013)		0.001 (0.013)
Active Physicians Rate by 1000		-0.023 (0.042)		0.012 (0.026)		-0.007 (0.027)		0.016 (0.012)
Observations	6,007	6,007	6,007	6,007	6,007	6,007	6,007	6,007
R ²	0.442	0.549	0.058	0.146	0.622	0.648	0.395	0.430
Number of sub regions	60	60	60	60	60	60	60	60
F-statistics	216.9	158.1	15.34	9.586	145	80.52	236.9	128.9

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Results based on 2935544223 observations sent from beginning of 2011 to 4th of March 2012

Table 5: Robust sample free of suspected electoral fraud, weighted regression results by eligible voters

VARIABLES	Party share							
	(1) UR	(2) UR	(3) KPRF	(4) KPRF	(5) JR	(6) JR	(7) LDPR	(8) LDPR
Difference in Difference Estimator	0.032* (0.017)	0.033** (0.016)	-0.001 (0.009)	-0.002 (0.009)	-0.027** (0.014)	-0.027** (0.013)	-0.011* (0.006)	-0.011* (0.006)
Treatment	-0.040 (0.030)	-0.030 (0.023)	-0.020 (0.018)	-0.021 (0.016)	0.034** (0.016)	0.035** (0.015)	0.007 (0.009)	0.004 (0.008)
Timetrend	0.192*** (0.009)	0.182*** (0.011)	-0.030*** (0.006)	-0.019** (0.010)	-0.114*** (0.006)	-0.106*** (0.008)	-0.070*** (0.004)	-0.071*** (0.005)
Election Turnout at PEC (log)		0.165*** (0.027)		-0.065*** (0.017)		-0.052*** (0.012)		-0.037*** (0.012)
Eligible voters by km2		-0.019*** (0.004)		0.009** (0.004)		0.002 (0.002)		-0.002 (0.002)
IP per capita		-0.009 (0.024)		-0.000 (0.016)		-0.011 (0.012)		0.011 (0.011)
Primary Income Private Households per head		-0.049* (0.029)		-0.012 (0.021)		0.018 (0.015)		0.001 (0.016)
Active Physicians Rate by 1000		-0.021 (0.048)		0.011 (0.030)		-0.007 (0.029)		0.016 (0.014)
Observations	6,007	6,007	6,007	6,007	6,007	6,007	6,007	6,007
R ²	0.443	0.550	0.058	0.148	0.619	0.645	0.394	0.429
Number of sub regions	60	60	60	60	60	60	60	60

Standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Results based on 2,935,544,223 observations sent from beginning of 2011 to 4th of March 2012

Table 6: Robust sample free of suspected electoral fraud, extend by 2 previous elections

VARIABLES	Party share							
	(1) UR	(2) UR	(3) KPRF	(4) KPRF	(5) JR	(6) JR	(7) LDPR	(8) LDPR
Difference in Difference Estimator	0.040** (0.018)	0.039** (0.018)	-0.004 (0.013)	-0.003 (0.012)	-0.036** (0.018)	-0.035** (0.017)	-0.010 (0.007)	-0.008 (0.007)
Election Turnout at PEC (log)		0.120*** (0.016)		-0.062*** (0.012)		-0.029*** (0.006)		-0.026*** (0.006)
Eligible voters by km2		0.538** (0.220)		-0.459*** (0.092)		0.009 (0.186)		-0.116 (0.076)
IP per capita		-0.001 (0.021)		-0.008 (0.010)		0.009 (0.014)		0.010 (0.007)
Primary Income of Private Households per head		0.064 (0.208)		0.001 (0.089)		-0.008 (0.092)		0.046 (0.053)
Active Physicians Rate by 1000		-0.168 (0.174)		0.168 (0.102)		0.193 (0.153)		0.012 (0.073)
Observations	9,993	9,993	9,993	9,993	9,993	9,993	9,993	9,993
R^2	0.688	0.708	0.544	0.563	0.809	0.815	0.492	0.501
Sub region fixed effect	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year fixed effect	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Sub region time trend	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Number of sub regions	51	51	51	51	51	51	51	51

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Results based on 2935544223 observations sent from beginning of 2011 to 4th of March 2012

The validity of the difference in difference design for causal estimates is reliant on the parallel trend assumption. Ideally, it would have been preferable to test the parallel trend assumption with an additional election that year. The next best scenario involved adding election results from the 2007 Duma and the 2008 presidential elections to the pre-treatment period. Using Angrist and Pischke (2009) as a guide to creating a robustness test, sub-region specific time trends and sub-region fixed effects were added to equation 3 in order to allow different time trends for each sub-region. Due to voting anomalies at previous elections, nine sub-regions were excluded. The results presented in table 6 show a slight increase in the difference in difference estimator in equation two, which strengthens the evidence for the effects of censorship in the Russian Federation on electoral outcomes.

5 Discussion

It has been shown that limiting bandwidth is a tool used by autocratic leaning governments to censor the Internet. The question that arises is whether results of this censorship are generalizable. Difference in Difference estimates are generalizable when the selection into treatment or control follow random assignment (Varian, 2016). In the setting described, an active Internet connection was a requirement in addition to the technical possibility for the government to interfere with it. Although results showed an increase of 3.2 percentage points to the government in areas which suffered internet censorship, this estimate cannot be generalised as the method used by the government for selection is unknown. Compared to Enikolopov et al. (2011), who showed that citizen access to a private television station, supportive of the opposition, led to a 1.55 percentage point decrease in support for the government, limitation of the market for information appears to have more effect. These results contrast starkly to a market over saturated with information options. In the United States DellaVigna and Kaplan (2007) found that the one-sided information broadcast by Fox News biased towards the Republican party had a favourable election impact of between 0.4 and 0.7 percentage points. It should be noted

that the detection of government censorship may coincide with other governmental interventions and thus act as a signal for other activities intended to subvert democracy. Future research will investigate this avenue.

The estimated effect of internet tampering was not substantive enough to have changed electoral outcomes, however, it was just one of the methods employed by the Russian government. Opposition candidates for other parties, such as Yabloko, were excluded from ballot papers - a ruling from the central electoral commission asserted that signature requirements for the candidate were impermissible due to not being of 'high enough quality' (Nichol, 2012).

Russian interference with the distribution of news is not only limited to media markets within the Country. The German newspaper *Sueddeutsche* analysed a leaked dataset detailing the operations of a company of 600 employees, tasked to manipulating social media entries and comments on news websites in foreign countries (Hans, 2014). Following Russia's actions in the Crimea, a stark divergence was observed between public opinion in serious surveys and comments found on German news articles on the internet.

6 Conclusion

This research paper presented a unique data set and a novel method of detecting previously disguised government censorship. Data showing this intervention was then used to estimate the impacts on the electoral process in Russia for the 2012 presidential election using a difference in difference design. On average, a treatment effect of 3.2 percentage points benefited the government supported candidate.

The results suggested that the rigging of votes on election day is not the only dimension when it comes to electoral fraud. Future research is needed to investigate the effects of internet throttling in other countries.

Although the findings of government intervention into the market for information in Russia are instructive and provide tools to identify internet censorship in other coun-

tries, publication of these activities are unlikely to have an impact within the Russian Federation. An insight into current president Vladimir Putin's contempt for freedom of information is illustrated by his quote:

“Contrary to a common perception, mass media is an instrument, rather than an institution.”⁷

⁷Recorded in Enikolopov et al. (2011)

References

- Abramitzky, R., & Lavy, V. (2014). How responsive is investment in schooling to changes in redistributive policies and in returns? *Econometrica*, *82*(4), 1241–1272.
- Ackermann, K., & Angus, S. D. (2014). A Resource Efficient Big Data Analysis Method for the Social Sciences: The Case of Global IP Activity. *Procedia Computer Science*, *29*, 2360–2369.
- Ackermann, K., Angus, S. D., & Raschky, P. (2016, August). *High granularity internet activity traces for 1647 major cities worldwide 2006-2012* (Working Paper). Monash University.
- Ackermann, K., Angus, S. D., & Raschky, P. A. (2017). The internet as quantitative social science platform: Insights from a trillion observations. *arXiv preprint arXiv:1701.05632*.
- Anderson, C. (2013). Dimming the Internet: Detecting throttling as a mechanism of censorship in Iran. *arXiv.org*.
- Angrist, J. D., & Pischke, J. S. (2009). Mostly harmless econometrics: An empiricist's companion. *Princeton Univ Pr*.
- Bertrand, M., Duflo, E., & Mullainathan, S. (2004). How Much Should We Trust Differences-In-Differences Estimates? *The Quarterly Journal of Economics*, *119*(1), 249–275.
- Campante, F. R., Durante, R., & Sobbrío, F. (2013, May). *Politics 2.0: The multifaceted effect of broadband internet on political participation* (Working Paper No. 19029). National Bureau of Economic Research.
- DellaVigna, S., & Kaplan, E. (2007). The Fox News Effect: Media Bias and Voting. *The Quarterly Journal of Economics*, *122*(3), 1187–1234.
- Enikolopov, R., Korovkin, V., Petrova, M., Sonin, K., & Zakharov, A. (2012). Field experiment estimate of electoral fraud in Russian parliamentary elections. *Proceedings of the National Academy of Sciences*, *110*(2), 448–452.

- Enikolopov, R., Makarin, A., & Petrova, M. (2015). Social Media and Protest Participation: Evidence from Russia. *SSRN Electronic Journal*.
- Enikolopov, R., Petrova, M., & Sonin, K. (2013). Social Media and Corruption. *SSRN Electronic Journal*.
- Enikolopov, R., Petrova, M., & Zhuravskaya, E. (2011, December). Media and Political Persuasion: Evidence from Russia. *American economic review*, 101(7), 3253–3285.
- Esfandiari, G. (2013). Iran Admits Throttling Internet To 'Preserve Calm' During Election. *RFERL*. Retrieved 2016-07-07, from <http://www.rferl.org/content/iran-internet-disruptions-election/25028696.html>
- Falck, O., Gold, R., & Heblich, S. (2014, July). E-elections: Voting Behavior and the Internet †. *American economic review*, 104(7), 2238–2265.
- Fan, X., & Heidemann, J. (2010). Selecting representative IP addresses for Internet topology studies. In *Proceedings of the 10th acm sigcomm ...* (p. 411). New York, New York, USA: ACM Press.
- Flock, E. (2012, February). *Iran gets back e-mail access, but other sites remain blacked out ahead of protest*. Retrieved 2016-07-05, from https://www.washingtonpost.com/blogs/blogpost/post/iran-gets-back-e-mail-access-but-other-sites-remain-blacked-out-ahead-of-protest/2012/02/13/gIQAgxz5AR_blog.html
- Franke, U., & Pallin, C. V. (2012). *Russian Politics and the Internet in 2012*. Routledge.
- Gentzkow, M., & Shapiro, J. M. (2010). What drives media slant? Evidence from US daily newspapers. *Econometrica*, 78(1), 35–71.
- Gentzkow, M., Shapiro, J. M., & Sinkinson, M. (2011, December). The Effect of Newspaper Entry and Exit on Electoral Politics. *American economic review*, 101(7), 2980–3018.
- Hans, J. (2014). *Propaganda aus russland - putins trolle*. Retrieved 2016-07-19, from <http://www.sueddeutsche.de/politik/propaganda-aus-russland-putins-trolle-1.1997470>
- Heidemann, J., Pradkin, Y., Govindan, R., Papadopoulos, C., Bartlett, G., & Bannister,

- J. (2008, October). Census and survey of the visible internet. In *Proceedings of the acm internet measurement conference* (pp. 169–182). Vouliagmeni, Greece: ACM.
- Iran: *Widespread arrests reported on day of protest*. (2012, February). Retrieved 2016-07-05, from <http://www.payvand.com/news/12/feb/1147.html>
- Klimek, P., Yegorov, Y., Hanel, R., & Thurner, S. (2012, October). Statistical detection of systematic election irregularities. *Proceedings of the National Academy of Sciences*, *109*(41), 16469–16473.
- Kobak, D., & Shpilkin, S. (2016). Integer percentages as electoral falsification fingerprints. *The Annals of Applied . . .*, *10*(1), 54–73.
- Kolko, J. (2009, December). A new measure of US residential broadband availability. *Telecommunications Policy*, *34*(3), 132–143.
- Kolko, J. (2012, January). Broadband and local growth. *Journal of Urban Economics*, *71*(1), 100–113.
- Kuerbis, B., Asghari, H., & Mueller, M. (2013, March). *In the Eye of the Beholder: The Role of Needs-Based Assessment in IP Address Market Transfers* (Tech. Rep.). Rochester, NY.
- Li, R., & Shiu, A. (2012, November). Internet diffusion in China: A dynamic panel data analysis. *Telecommunications Policy*, *36*(10-11), 872–887.
- Lipman, M., & McFaul, M. (2005). Putin and the Media. *Putin's Russia: Past Imperfect, Future Uncertain*, ed. Dale R. Herspring. Lanham, MD: Rowman and Littlefield.
- Mack, E. A., & Rey, S. J. (2014, February). An econometric approach for evaluating the linkages between broadband and knowledge intensive firms. *Telecommunications Policy*, *38*(1), 105–118.
- Menezes, A. (2015). Internet Availability, Political Information, and Voting: Evidence from Brazil.
- Miner, L. (2015). The unintended consequences of Internet diffusion: Evidence from Malaysia. *Journal of Public Economics*, *132*, 66–78.
- Nichol, J. (2012). Russia's march 2012 presidential election: Outcome and implications.

- Current Politics and Economics of Russia, Eastern and Central Europe*, 27(3/4), 357.
- Petrossian, F. (2012, February). *Iran: Internet Blackout Ahead of Protest Day*. *Global Voices*. Retrieved 2016-07-05, from <https://globalvoices.org/2012/02/13/iran-internet-blackout-ahead-of-protest-day/>
- Petrova, M. (2008, February). Inequality and media capture. *Journal of Public Economics*, 92(1-2), 183–212.
- Reporters Without Borders. (2012, March). Internet Enemies Report 2012. , 1–72.
- Sen, A. (2008). social choice. In *The new palgrave dictionary of economics*. Basingstoke: Palgrave Macmillan.
- Varian, H. R. (2016, July). Causal inference in economics and marketing. *Proceedings of the National Academy of Sciences*, 113(27), 7310–7315.