Turning the Virtual Tables: Government Strategies for Addressing Online Opposition with an Application to Russia

Sergey Sanovich*

Denis Stukal

Duncan Penfold-Brown

Joshua Tucker

New York University

June 16, 2015

Preliminary Draft:

Please do not circulate without contacting authors for latest version.

Abstract

We introduce a novel classification of strategies employed by autocrats to combat hostile activity on the web and in social media in particular. Our classification looks at these options from the point of view of the end internet user and distinguishes both online from offline response and exerting control from engaging in opinion formation. For each of the three options – offline action, infrastructure regulation and online engagement – we provide a detailed account for the evolution of Russian government strategy since 2000. In addition, for online engagement option we construct the tools for detecting such activity on Twitter and test them on a large dataset of politically relevant Twitter data from

*Corresponding author: sanovich@nyu.edu.

Russia, gathered over the period of nine month in 2014. We make preliminary conclusions about the factors of internet policy choice in non-democracies.

1 Introduction

On December 3rd, 2014, the Russian news website *Meduza.io* reported that the 100th mirror of another Russian news website *Grani.ru* was banned by the Russian Federal Service for Supervision of Communications, Information Technology and Mass Media (Roskomnadzor). In a month since then this count reached 154 mirrors. *Grani.ru* was a popular news website with extensive coverage of opposition activity and alternative opinions. It was blocked in the Spring of 2014, at the height of Russian-Uranian conflict, using a technical system developed by Roskomnadzor to block content deemed as extremist, as allowed under Russian law. As of early 2015 none of the (supposedly easy) legal avenues to lift the ban – provided by the same law that was used to block the website – appeared to be working. *Meduza.io* itself is a new Russian news website, created by the former editorial staff of the most popular Russian news website, *Lenta.ru*, after its editor-in-chief Galina Timchenko was dismissed by the website owner over the coverage of the Russian-Ukrainian conflict. The new media source was created in the neighboring Baltic state of Latvia and most of the editorial staff moved there. Instead of using the Russian domain zone *.ru* it is located in the *.io* zone, which belongs to UK-governed British Indian Ocean Territory. Around the same time, the blog of the Russian opposition leader Alexey Navalny - one of the most popular political blogs in Russia - was also permanently banned on the "Livejournal" platform, and in early 2015 authorities began to crack down on its mirrors too.

Interestingly, the Russian government's response to unfriendly activity online (and on social networks in particular) has not always been through bans and legal action. As late as in 2010 the report of the *Internet in Russian Society* program at the Berkman Center for Internet and Society at Harvard University noted that "the political blogosphere appears to remain a free and open space for Russians of all political stripes to discuss politics, criticize or support

government, fight corrupt practices and officials, and to mobilize others around political and social causes" (Etling et al. 2010). Moreover, as recently as 2009 the newly elected Russian president Dmitry Medvedev opened his own blog at Livejournal and subsequently established a presence on Twitter and Facebook, as did many of his aides. Accordingly, the pro-government youth movements, which were created to confront possible "colored revolutions" on the streets of Moscow were charged with the duty of competing with oppositional voices in the cyberspace and promoting government-friendly content. In some cases they even directly engaged with leading oppositional bloggers on the pressing issues of the day.

Why were the changes in policy so quick and dramatic? What is the menu of options the government can choose from to respond to emerging online challenges? In this paper we address these questions by creating a new classification system for different forms of government response to online opposition, which we then apply to the Russian experience in the last 15 years. Our classification looks at these options from the "point of view of the user" and distinguishes both online from offline responses and exerting control from engaging in opinion formation. As a result we highlight three options: offline responses, which include legal action and market regulation; online infrastructure responses, which rely on digital tools to filter the information available to end users; and direct online engagement with users that aims to shape online conversations, usually through content generation.[1] Applying this form of classification, we analyze the evolution of internet policies in Russia during Putin's first term in office (2000-2008), the Medvedev interregnum (2008-2012), and the period of time since Putin's return to Kremlin after 2012. We investigate why the government almost completely ignored the internet when it was actively imposing its will on traditional media and why this policy changed after Putin left the Kremlin in 2008. We also look at why under Medvedev online engagement rather than imposing heavy restrictions was chosen as a premier strategy and why this choice was reversed when Putin and Medvedev switched offices in 2012.

[1] Another tactic is intimidation.

In addition, we provide initial findings from our attempts to employ tools of digital forensics to identify and describe bot activity in the Russian Twittersphere using a large dataset of tweets collected during a very turbulent nine month period of recent Russian history (from February to November of 2014). Using four different statistical methods of bot detection, we successfully identify a set of bots and perform semi-supervised of patterns and content of their posts.

Finally, we discuss what general conclusions regarding the internet policy choice in non-democracies we can take from the Russian experience under Putin. In particular, we emphasize domestic politics considerations, which are often overlooked in discussions about technical capacity and general repressiveness of the regime.

The paper is organized as follows. First, we review the existing classifications of government approaches to controlling internet, discuss their shortcomings and present our new classification of the options available to the government. In Section 3 we discuss each option at government's disposal in detail. In the next section we apply our classification to the evolution of internet control in Russia. Section 5 contains quantitative empirical analysis of bot activity, using Russian Twitter data from 2014, which includes our surprising finding that anti-government bots - at least in our sample - appear to be just as active as pro-government bots. Section 6 concludes with a discussion of basic factors affecting internet policy choice as they emerge from the Russian experience we describe.

2 Government response online: in search of a framework

Various forms of government online activities have been the subject of intensive research in recent years. This literature covers everything from legal regulation of the internet in world's most advanced democracies (Giacomello 2008; Travis 2013) to online censorship tools used by different types of autocratic governments around the world (Ko, Lee, and Jang 2009; MacKinnon 2011; King, Pan, and Roberts 2013; Nabi 2013) to establishing an active presence of governments on social media platforms (Barash and Kelly 2012; Pearce 2014). Freedom

House even produces an annual international index of internet freedom (Freedom House 2011, 2012, 2013, 2014). However, few studies attempt to provide a framework for the systematic analysis of government behavior on the web that would allow us to analyze why particular tools are chosen under any given circumstances, or even what makes up the choice set in the first place.

A notable exception is E. Morozov (2011), who distinguishes between technological and what he calls "sociopolitical" means of controlling online activity, the latter combining technology with online and offline actions by humans. Technological responses include internet-filtering, which spans from targeted bans of particular webs-sites and keywords to larger national-level schemes to block entire segments of the internet (China) or – in the extreme – any outside internet access outside the country (North Korea). Sociopolitical responses range even more widely from distributed denial-of-service (DDoS) attacks to employing both automated bots and paid trolls to destroy online communities' social capital to physical attacks on bloggers.

Morozov hypothesizes that if "liberation technologies", such as those promoted by Diamond (2010), were to succeed, embattled governments could turn to potentially more violent methods such as smearing campaigns, criminal prosecutions or even physical attacks on bloggers: "as technological methods lose efficacy, sociopolitical methods could simply overtake them: an authoritarian government might find it harder to censor blogs, but still rather easy to jail bloggers" (E. Morozov 2011, 63).

While this option is certainly not hypothetical, and Morozov's classification is useful for studying the dangers and promises of "liberation technologies" (what makes sociopolitical response different is exactly being beyond the reach of these technologies), it does not properly distinguish between government actions that restrict, or otherwise structure, online media environments and those where the government actively engages in shaping the formation of opinions online. This distinction is important for at least three reasons. First, while in *censor-*

ing online media the government could build (at least to some extent) on strategies from long before the internet was created, online *propaganda* in distributed networks is fundamentally different from a top-down broadcasting of the party line through hierarchical monopolies of traditional media. Second, this part of the government response is experienced differently by users: not as an outcome (e.g., an inaccessible web page), but as a point of interaction with the state (paid pro-government troll replying to your tweet). Last but not least, the study of government online activities will increasingly focus on social media, which is simultaneously the most abundant and versatile data source and the key point of contestation between the government and civil society (Deibert et al. 2011; Etling, Roberts, and Faris 2014; Pearce 2014; Lange 2014; Gunitsky 2015). Since social media data capture exactly the moment of interaction between the user and the state, it is important to understand the place of this type of government action (which leads to this interaction) in the wider menu of options available to the government.

Therefore, we propose a new classification of government options for responding to independent online activity. In addition to differentiating between "**offline**" and "**online**" tools, it distinguishes between online tools aimed merely at restricting the flow of information and those entailing active engagement with the users on behalf of the government. While the former operates largely through exerting control over the internet **infrastructure**, the latter typically involves some **content generation**. Since users experience each of these possible government actions differently, our classification, effectively, dissects government options *from the internet user's point of view*. In the next section we discuss each option in detail, providing examples and identifying the key resources needed to employ each of these options. This classification will then inform our analysis of the strategy pursued by the Russian government in Section 4 and the theoretical discussion in Section 6.

3 A new classification system for government responses to online opposition

In this section, we introduce our tripartite classification system for classifying government responses to online opposition. We begin with offline response, primarily focused on changes to a country's legal infrastructure regulating internet usage but also including attempts to change ownership structure. The second category focuses on attempts to control the infrastructure of online communications by the use of tools such as firewalls and online censorship; one way to think of these approaches are as attempts to *limit* online content. The final category also involves online interaction, but instead of placing limits on content, this tactic involves *engaging* online to attempt to effect the tone of the online convesation in a more pro-government direction.

3.1 Offline Responses

The first set of options at any government's disposal is based on digital age implications of traditional governing advantages: nodality ("network centrality"), organizational capacity, legal authority to enforce the law, a monopoly on the legitimate use of violence and the right to regulate human activity and, last but not least, the ability to expend large financial resources through taxation (see a detailed discussion in Ackland 2013, Chapter 8). The actions facilitated by these advantages could have a huge impact online, but take place offline; thus end-users either observe the consequences online after-the-fact (and, if necessary, adjust their behavior) or encounter these actions in person, but offline. The latter case, of course, applies to legal prosecution and violence against (usually, social media) users. In more favorable circumstances, users just observe the outcomes of government actions when their comments (and other user-generated content, or UGC) become legally designated as properly regulated media content (with corresponding obligations and risks on part of media outlets), and they find commenting functionality turned off at their favorite news web-sites.

Another option is to do something along the lines of steps recently taken by Russian Duma: forcing "popular" bloggers – in Russia's case defined as bloggers having more than 3000 readers a day[2] – to register as semi-media outlets, making each individual blogger responsible for her own content, on par with actual commercial media outlets (Macfarquhar 2014). As Ackland notes, such a legal designation – as well as other forms of regulation of UGC – could be attributed by the government to either genuine or fabricated popular demand stemming from concerns over public safety or morality (Ackland 2013, 143, 145-146).

Finally, the government could attempt to change the landscape of the digital media market and alter the choice of online platforms available to users. Relying on their authority to regulate commerce, autocrats around the world designate certain companies and industries, including in telecommunications, as "strategic", upon which they start to enforce various restrictions which "strategic industry" status implies, such as banning foreign ownership and/or investments, appointing state representative to the board, etc. For example, in late 2013 the publicly-owned – but heretofore relatively independent editorially – major Russian news agency *RIA Novosti* was beheaded, renamed, restructured and put under the leadership of a fervent regime supporter, Dmitry Kiselev (Sumlenny 2013). Nevertheless, to ensure the complete control of this already loyal news outlet (with a large, popular and respected social media presence), in early 2014 it was included in the list of "strategic enterprises", along with the second largest Russian news agency, *ITAR-TASS* (Tetrault-Farber 2014).[3]

In 2014 Russia changed its media law to ban (beginning January 1, 2016) foreign ownership (defined as more than 20%) of any media operating in Russia. In addition, any foreigner or Russian citizen with another citizenship was banned from serving as the founder or editor of any media (Gulyaeva et al. 2014; Birnbaum 2014). While this law primarily affected print media, many print publication have a significant presence on the web and are widely shared through social media (The Moscow Times 2014).

[2] See Section 4.3 for details.

[3] The list was first published as an Executive Order of the President of Russia and has been continuously updated ever since. Unlike the Strategic Sector Law (2008), the list extends beyond oil and military equipment production (Bam 2013).

If control over digital media is impossible to legislate or order, especially in the case of private companies, then governments can use other means, in particular purchasing power and extra-legal pressure, to assume control over important internet platforms. The so-called "Russian Google", Yandex, sold a golden share to state-owned *Sberbank* in 2009, allegedly after negotiations with Dmitry Medvedev and multiple proposals to designate companies such as Yandex as "strategic", which would have forced them to re-register in Russia[4] and severely diminish their appeal to international capital markets (Grishin 2012). In 2014 Yandex founder Arkady Volozh resigned as the company Russian CEO.[5] A similar attempt was made in the case of Vkontakte, known as the "Russian Facebook", which resulted in the hostile takeover of the company by business groups loyal to the Russian government (Kononov 2014; Yaffa 2013) and founder and former owner and CEO Pavel Durov (and his team) fleeing the country (Lunden 2014).

Of course, instead of attempting to take control over existing public and private media and communications platforms, the government could try to increase its influence through artificially generated competition. In several countries, including Russia and Turkey, governments reportedly allocated generous funds to the creation of "national" search engines, "national" social networks or "national" email services.[6] A Russian national search engine has been discussed since at least 2008, when it was mentioned by then President Dmitry Medvedev after the so-called 7-days war with Georgia. Turkey began a similar project in 2009 (E. Morozov 2011). The results of government investments in creating artificial competition are less than impressive, though: the type of poor management by corrupt executives these countries could get away with in mineral resources extraction does not work nearly as well under intense national and international competition, especially given the problematic implementation of compulsory consumption online. In late 2014, the Turkish project was still in development,

[4] Yandex is incorporated in the Netherlands as Yandex N.V. – a fact that in 2014 was publicly condemned by Vladimir Putin at his meeting with People's Front for Russia (Brennan 2014).

[5] He kept the executive position in the international operations, though (Beard 2014).

[6] In the Russian case these pleas were especially suspicious given that Russia already had a "national" search engine, social network and email service, all created without any government aid. Yandex and Vkontakte have larger market shares in Russia than Google and Facebook, respectively, an impressive result in the absence of any protectionist measures against foreign competitors.

with a scheduled launch in 2016 [7]. In Russia, Rostelecom – the largest telecommunications company in Russia, with the government as its majority shareholder – released a beta-version of the search engine "Sputnik" in May of 2014, but it has not been able to acquire a noticeable market share as of yet. In the Russian case, though, successful hostile takeovers of existing platforms eliminated the need to build new ones from scratch.

3.2 Online: infrastructure

Offline means of controlling online activity are popular among autocrats around the world, not the least because they usually require zero IT competence or resources. However, rapid growth in internet penetration rates and the emergence of the internet as a principal source of information for increasing numbers of people creates challenges even for autocrats who are able to successfully employ offline tools of control. To begin with, information can be produced and distributed by foreign citizens and entities that are out of reach for the autocrat's security apparatus. Second, some local activists and/or journalists can use their digital proficiency to distribute information anonymously and therefore avoid offline prosecution. Finally, for various reasons autocrats could simply prefer putting flows of information under their control rather than going after its producers. If, for example, an autocrat wants to avoid taking responsibility for the government's actions, a DDoS attack on a popular oppositional blog can be blamed on "unidentified" hackers, while most types of offline response require at least some involvement of the state apparatus.[8]. Thus, governments facing a serious threat from the opposition or an insurgency may try to acquire tools to exert control over internet infrastructure and deploy such tools either routinely or in times of political instability.

Of course, there always exists the option to completely monopolize the telecommunication infrastructure inside the country and cut any connections with international networks.

[7] http://english.yenisafak.com/news/tubitak-to-develop-national-search-engine-2024960

[8] Even if the attack on a blogger is carried out by private citizens, and the government claims to have no relation with them, there is a strong expectation and laws on the books that require criminal investigation in the case of violence or even threat of violence. DDoS attacks on the other hand, and cybercrime in general,remain weakly (if at all) legally regulated in most countries, and usually neither the law nor public opinion demands any action on the part of the government in the case of a cyberattack.

North Korea did just that: it maintains *Kwangmyong*, a national intranet, and a national mobile phone service *Koryolink*. The North Korean authorities have complete control over the information accessible through *Kwangmyong* and tightly monitor personal communications through *Koryolink*. Communications with the outside world through both channels are prohibited (except for the ruling elite and foreign tourists). However, such a system imposes a heavy toll on the national economy and state capacity.

A step removed from this extreme approach – albeit still with non-trivial costs – is the highly sophisticated Chinese "Great Firewall", probably the best example of blocking sensitive information without fatally hurting either government communications or commercial activity (Economist 2013). Country experts anticipate that even North Korea will eventually take the Chinese and Cuban path of establishing 'such a 'Mosquito-Net" model of internet access, which facilitates the"use [of] the Internet as a propaganda machine in addition to taking advantage of it economically, [...] while keeping out information deemed threatening by the regime" (Chen, Ko, and Lee 2010; see also Ko, Lee, and Jang 2009).

The ability of autocratic governments to filter internet communications (including the access to social media) is, primarily, a function of three factors: control over the critical infrastructure; planning ahead and implementing a long-term, comprehensive, but not overly costly solution; and financial and human resources as well as the technical expertise necessary to build filtering tools. The first component is rarely an advantage of autocratic regimes: most of the world's internet infrastructure (ISPs, search and social media platforms, transaction services, etc.) is located in advanced democracies and therefore out of reach for autocrats seeking to control them. Consequently, the latter two factors – i.e., preparedness and money – determine ultimate success. According to Howard and Hussain (2013, 71-72) "sophisticated long-term investments in managing information infrastructure" made by countries such as Iran, Saudi Arabia and Bahrain have made these countries much less vulnerable to growing discontent than Egypt, Libya, and Yemen, who have had to rely on ad hoc solutions. The key risk of the latter approach is that an abrupt shutdown of internet access can end up hurt-

ing the government's ability to communicate with its own forces allies while simultaneously provoking further unrest by interfering with the non-political uses of internet.

Two primary technological options for regimes are filtering/blocking of particular websites or segments of the web and DDoS attacks. The former has the advantage of being permanent and customizable. China, for example, blocks only certain platforms and content (by keywords), while North Korea famously maintains its local web segment in complete isolation from the outside web. Both policies, though, share a common disadvantage of this approach: high transparency for local users and susceptibility to documentation by outsiders (including other governments, human rights organizations, etc.) (E. Morozov 2011).

DDoS attacks, on the other hand, are usually hardly traceable, relatively cheap, can be deployed during particularly sensitive political events such as elections or protests and can be more easily outsourced to loyal but independent groups, such as the Syrian Electronic Army (Noman 2011). On the other hand, their ability to break up online communications is limited in time and web space, i.e., a small set of web-sites at best. Moreover, the most popular platforms, such as Google and Twitter, are highly protected from DDoS attacks.

The most distinctive (and important for our discussion) feature of internet filtering is that it is observed by users as an end result and does not create any kind of interaction with the state (or its representatives) in course of user's *activity* online. If Twitter is blocked in your country (permanently, as in China, North Korea, Iran and several other countries, or temporarily, as in Venezuela and Turkey in 2014), you either cannot get there, or you can use one of the available tools (anonimizers, proxy-servers, etc.) to restore your access. In either case, users observe government actions as end results. The impact of such actions is either in successful breaking inter-personal communication or access to websites, or the lack thereof. In other words, the government cannot possibly shape the conversation through these means. To achieve this latter goal government has to directly engage with users online.

3.3 Online: content

Establishing a government presence on the web and using it to promote the government's agenda constitutes the third and final option at a government's disposal. This type of government response actually takes place online and users encounter it in course of their online – particularly social media – activity. Mainly, it includes the government actively producing content, either through artificial intelligence or real human effort. However, internet-censorship, if it is targeted, prompt and identifiable as such, could constitute its own form of communicating the government's point of view. Similarly, hacking and publishing bloggers' personal communications (such as emails, instant messages, etc.) could allow the government to expose and implicate the opposition and shape the conversation that way. A typical example of the latter is the case of Russian blogger and hacker *torquemada_hell*, a Russian-speaking person allegedly living in Germany. In 2010-11 he successfully hacked the email accounts of multiple Russian opposition politicians and released potentially damning information to the public (Sidorenko 2010).

Targeted internet-censorship is well documented by King, Pan, and Roberts (2013), who describe the immense Chinese system of monitoring and censoring of UGC across the country's dispersed social media platforms. King, Pan, and Roberts (2013) estimate the average amount of censorship at around 13%, find that most of the censorship happens within 24-hours from intial posting, and claim that in spite of obvious technical difficulties the level of censorship increases disproportionately during periods of especially voluminous discussions on social media after significant political (as well as non-political) events. Moreover, they mention that censors either leave notes such as "Sorry, the [ost you were looking for does not exist, has been deleted, or is being investigated" or post "pictures of Jingjing and Chacha, Internet police cartoon characters". Given that censors do not delete or block the entire threads which contained the deleted message, such interference could be easily observed by participants in the course of their conversation. While the discussion of censors themselves is also subject to censorship, the conversation could continue, affected or unaffected by the intrusion.

Still the most obvious – and increasingly popular – tool employed by governments to alter political conversations on the social media is using either "bots" or real people to advocate pro-government positions, turn conversation meaningless or prohibitively divisive, or distract users from sensitive political issues altogether.

Bots could perform two key functions: either cluttering conversations with "digital dust" or altering search results, internet rankings, top lists and other automated tools for sorting, sharing, discovering and consuming online content. As such, bots could be used to support real people. For instance, a ranking of the most popular Russian blog posts, maintained by Yandex, was closed in 2009, inundated by bots promoting (mostly pro-government) posts[9] (Sidorenko 2009).

The possible functions of humans acting on the government side are much more diverse. It is useful, therefore, to provide a basic classification of pro-government users. This classification does not look into users' honesty, consciousness and convictions. Instead, it is based on formal or informal ties with government (or lack of thereof).

To begin with, the government could hire students or other low-paid workers to submit rather simple messages, which would nevertheless pass the human intelligence tests integrated in many modern social media platforms. Their messages could be either identical or contain the same message worded differently. One particular example of this type of bloggers are the so-called Chinese 50-centers (Cook 2011). Russian pro-government youth movements, such as *Nashi* and *Young Guard of United Russia* were often accused of running a similar network of 11-rublers (Nossik 2013). Leaks released by the Russian arm of Anonymous in 2012 indicated that Nashi paid hundreds of thousand of dollars in fees for comments, statuses, Facebook likes, Youtube dislikes, etc. (Elder 2012b, 2012c).

[9]The key concern for Yandex, though, was government outrage in the cases when anti-government posts got traction in the ratings (Odynova 2009).

Cheap bloggers paid per comment are not the only group of friendly users that could be put on the government payroll. Bribing prominent and trusted bloggers, celebrities or journalists – although potentially much more expensive – could turn out to be a better investment in terms of persuading the public. The same leaks as noted in the previous paragraph revealed that along with paying small fees to thousands of low-skilled bloggers, Nashi also put aside tens of thousand of dollars to be paid to a small group of popular and heretofore considered independent bloggers for highly sophisticated positive publicity for the Russian leadership (the Putin and Medvedev tandem at that time) (Elder 2012a).

The next group consists of government supporters whose social media activity is not paid per se, but is facilitated through participation in various political projects or actual employment by the government. These sets of bloggers range from members of various youth political movements to the MPs from the ruling (or affiliated) parties to relatively prominent politicians (ministers, party leaders) who are encouraged to take on the challenge of representing the "government's point of view" in an often hostile social media environment.

Finally, the government could also try to mobilize genuine supporters with no ties – formal or informal – to the government or ruling party. Obviously, this group could include various types of people, but two groups deserve special attention. First, if famous people – particularly, celebrities and journalists – volunteer to support the government agenda, it could help the autocrat both directly and indirectly through endowing the ideas already promoted by the armies of bots and paid bloggers with the weight of fame, reputation and personal independence. Second, people in the opposition could occasionally be legitimately attracted by government policies and join the ranks of pro-government bloggers. Prominent example from recent Russian history is the brief, but notablem, excitement of previoulsy largely oppositional nationalists about the Russian annexation of Crimea and support for separatists in Eastern Ukraine (Balmforth 2014; Yudina 2014; Nechepurenko 2014; Goble 2014). In the spirit of working class vs. fellow travelers debates between communists in 1920s, we could call them *poputchiks* (Caute 1988). Obviously, this latter category rarely would end switch-

ing sides permanently. However, given their relative immunity to direct bribing and indirect manipulation, fluctuations in their support of the government could provide a useful baseline for studying other groups.

In the next section, we illustrate the usefulness of this taxonomy by applying it to a particular case: Russia over the past decade and a half, or, put another way, Russian in the age of Putin.

4 Russian government online: a long way from Putin to Putin

4.1 2000-2008: Internet vs. traditional media in Putin's Russia

Russian reforms of 1990s got mixed assessment from both outside observers (Shleifer and Treisman 2001) and reformers themselves (Gaidar 2003), and their reception among Russians remains ambiguous at best (Denisova, Eller, and Zhuravskaya 2010). However, a wide consensus holds that if anything worked during the painful transition from Communism, it was media freedom (Zasursky 2004). Diverse, influential and competitive news outlets emerged almost immediately and by the end of 1990s several powerful media conglomerates were operating alongside a large network of independent federal and regional media, usually free of any government control (Lipman 2009). This is proved by the enormous role media played in political fights throughout 1990s (Koltsova 2006) and, above all, during so-called war for Yeltsin's succession (Gelman, Travin, and Marganiya 2014, 104-108), when high-powered media was mobilized by both Putin and his opponents. Role of the media in Putin's rise to power is well documented by rigorous quantitative studies (Enikolopov, Petrova, and Zhuravskaya 2011), Putin's biographers (Gessen 2012, Chapter 2), Western observers (Judah 2013, Chapter 2) and Russian political memorialists alike (Tregubova 2003, Chapter 10). This experience allowed Putin to fully appreciate the power of the media to change public opinion and reverse political fortunes. Putin's media policy in the next 15 years demonstrates that he took this lesson extremely seriously and worked tirelessly to put media under his control (Lipman 2009; Burrett 2010). However, as we shall see, this policy was not universally applied to all types

of media. To the contrary, online media enjoyed the *laissez-faire* regime that was in many respects on par with most advanced democracies. To uncover the reason, why Putin had drawn such sharp line between traditional and online media, we will start by examining the Soviet experience of media control. Seemingly inconsistent strategy Putin adopted would become much less surprising if discussed in light of the strategy once devised by Putin's (former KGB lieutenant colonel) employers at the Central Committee and Lubyanka. Next, we will discuss main features of Putin's dual strategy and assess the changes in political and media environment, which ultimately rendered it unworkable.

Soviet Union, as a relatively long-lived 20th century dictatorship[10], survived several waves of technological advancement. Already in 1917 Bolsheviks famously recognized the role of modern technologies and communications by seizing (along with the Winter Palace, railway stations, bridges and army headquarters) Petrograd's Telegraph and Telephone Exchanges. This recognition entailed two different strategies – for mass and personal communications – which were implemented fairly persistently throughout Soviet history. The government maintained the complete monopoly over the mass communication and fiercely prosecuted those who tried to challenge this monopoly. Most famous examples include jamming of Western radio broadcasters (Radio Liberty, Voice of America, BBC, etc.) and strict regulations over using printers, plotters and photocopier after they were installed at various Soviet administrative departments and institutions[11] (Komaromi 2004).

Personal communications was a different matter. Instead of monopolizing their usage, government allowed Soviet citizens to use them promiscuously and then used it to identify and prosecute those who were disloyal. Phones, for example, were spreading rapidly in the USSR, approaching roughly 37 million, or about 13 per 100 inhabitants in 1990 (Banks and Wilson 2013). It was still six times as small as the U. S. at the time, but the difference was due to technological and economic reasons, not political restrictions. However, government

[10] According to Przeworski et al. (2000) average dictatorship which was overthrown between 1950 and 1990 lasted 27.4 years and average dictatorship still in course in 1990 was 26.2 years old.

[11] The last among these regulations, which was not enforced anymore, was struck down by the Russian Supreme Court only in 2009.

used every opportunity to spy over its citizens using wiretapping. Under Stalin serious efforts were put in the research on speaker identification, which was famously depicted by Aleksandr Solzhenitsyn in the autobiographical novel *In the First Circle*. Later the elaborate system of surveillance and spread of personal, but publicly registered home phones eliminated the need in voice identification. Similarly, typewriters were allowed for personal use, however their printout had to be handed to the *First Department* (local KGB affiliates at any Soviet enterprize or institution) and could be cross-verified to identify the exact typewriter used in printing "inappropriate" materials.

After assuming power in 1999, Pitin gradually implemented similar strategy of complete monopolization of mass media and liberal policy on personal communications. While the latter was virtually left free from interference (but not from surveillance[12]), national television networks were returned to government ownership and Soviet-style management with weekly instructions delivered at Kremlin to the news executives (Gehlbach 2010). Press and radio remained more diverse, with some pro-government and some relatively independent outlets competing with each other (Lipman 2009; Womack 2014). The process of putting traditional media under government control in the early 2000s included such colorful episodes as imprisonment of media mogul and oligarch Vladimir Gusinsky (in order to be released he signed a secret protocol with Russian Minister of Communications Mikhail Lesin and handed over his media assets to state natural gas monopoly Gazprom) and stunning reversal of fortunes for the architect of both Yeltsin's electoral victory in 1996 and Putin's one in 2000 Boris Berezovsky (who lost control of main Russian TV channel and went to exile already by 2001) (Becker 2004). This and other episodes and their consequences for the media landscape are well documented in literature (see academic stydies in Burrett 2010; Dunn 2008; Burrett 2008; for more informal discussion see Gessen 2012, Chapter 7; and an illuminating personal memoir by Tregubova 2003, Chapters 11-13). Putin's approach to the internet, on other hand, was rarely assessed. Observers simply noted that "the Russian blogosphere is a space that

[12]Sophisticated "deep pocket" web surveillance system known as SORM (-2,-3) was installed by the Russian government no later than in the late 1990s and has been being updated constantly ever since (see Soldatov and Borogan 2013).

appears to be largely free of government control" (Etling et al. 2010, 33); "the absence of Internet filtering is notable. Based on tests run through the OpenNet Initiative, we continue to find no evidence of significant technical filtering of the Russian Internet" (Alexanyan et al. 2012, 10-11), etc. A recent account by one of the leading Russian internet news producers Anton Nossik suggests that it was no accident. Instead, already in 1999, still prime-minister Vladimir Putin had a clear preference for non-interference in the internet space:

> ... in December 1999, three days before he became acting president of Russia, Vladimir Putin [...] summoned all the heads of Russias nascent Internet industry for a meeting [...]. In his brief but passionate speech that day, Putin made special mention of Chinese and Vietnamese models of Internet regulation, stating that he viewed them as unacceptable. "Whenever we'll have to choose between excessive regulation and protection of online freedom, we'll definitely opt for freedom," he concluded [...]. (Nossik 2014)

Under the auspices of such a benevolent government policy, Russian online media flourished, creating a vibrant sector of the economy and a reliable source of information for millions of Russians. Russia is one of the few countries where Google is not the most popular search engine and Facebook is not the most popular social network. Remarkably, both occurred without restrictions on American competitors. Unlike Baidu and Weibo, Yandex, Odnoklassniki and Vkontakte won virtually fair competition with their American counterparts[13]. Successful development of local services did not mean that foreign ones were not actively used by Russian bloggers and readers. Livejournal, the most popular Russian social network in 2001-2011, while being originally American and predominantly English-speaking, developed Russian community so large that it was eventually overtaken by a Russian media holding and became dominated by Russian users (Greenall 2012). And as of April, 2014 Facebook has 24 million users from Russia [14] and Twitter has more than 8 million[15], which makes Russian one

[13] Still, Vkontakte (but not Yandex or Odnoklassniki) had significantly benefited from the lax enforcement of the property rights. However, this doesn't make comparison with China less impressive, given that China is also famous for wide-spread piracy.

[14] http://ria.ru/technology/20121004/766127348.html

[15] http://digit.ru/internet/20131031/407481403.html

of 10 most popular languages on Twitter[16].

Again, in stark contrast with most other countries, Russian most popular news web-sites do not represent traditional media such as newspapers, radio and TV broadcasters (Nossik 2014). Instead, *Gazeta.Ru*, *Lenta.Ru*, *NewsRu.com*, *Polit.ru* and alike were built from scratch and became major news outlets in their own right (i.e. their staff does original reporting, often as an eyewitness, rather than just digitalize other's content).

As a result, Russia developed a strong, powerful and independent internet media sphere, which was a remarkable achievement for any non-democratic country, but especially for one, where traditional media are so tightly controlled. As Alexanyan et al. (2012) note "Russia is unusual in the degree of freedom found online compared to offline media and political restrictions". Such imbalance, however, proved to be unsustainable. In the late 2000s internet media increasingly supplemented and eventually supplanted TV as the main news source at least for the educated Russians (Clover 2011). One of the leading Russian TV anchors Leonid Parfenov, who has been banned from air since 2004, aptly summarized this process in 2010 speech, which went viral on YouTube (Remnick 2011):

> These evergreen tricks are known to everyone who has witnessed the Central Television of the USSR. Reports are replaced by protocol shootings like "Meeting at the Kremlin"; reporter's intonations support the officials in the picture; broadcasting models are implemented to show "the leader receiving a minister or a governor", "the leader campaigns among the masses ", "the leader holding a summit with his foreign colleague", etc. These are not news; this is old record that repeats the already established patterns of broadcasting. Even a news hook isn't a must. In the emasculated media environment any small fry will pass for a big shot just because of getting some airtime.
>
> [...]
>
> It hurts twice as much to speak about television journalism, given the evident suc-

[16] http://bits.blogs.nytimes.com/2014/03/09/the-languages-of-twitter-users/

> cess of the large-scale TV shows and Russian school of television series. Russian TV is getting more and more sophisticated in exciting, fascinating, entertaining and amusing people, but it hardly could be called civic social and political institution. I am convinced it is the reason for the dramatic decline in TV viewership among the most active part of the population. People of our type say: "Why bother turning on the box? It's not intended for us".

However, as Nossik (2014) notes, this dual strategy – tight control of the traditional media and almost complete nonintervention in the web – was devised when Russian internet penetration was almost negligible. Even three years after Putin came to power, in 2002, Russia has 2.1 million people (2% of the adult population) who used internet daily [17]. By 2008 this share increased to 14 million (16% of the adult population), and by 2013 to 52.2 million people (46% of the adult population)[18]. Needless to say, the quality of access changed dramatically after wide access to broadband connection replaced slow dialup. This circumstances diminished the value of monopoly in TV broadcasting and strong influence in other traditional media which Kremlin enjoyed (Oates and Lokot 2013) and simultaneously made the online communities sufficiently large and well-structured to become politically significant. Dual nature social media, which is simultaneously mass and personal communication, presented a particular challenge for the government.

These changes coincided with the constitutionally required transition of power from Putin (who served two consecutive terms) to Medvedev in 2008. While Putin was appointed Prime Minister of Russia immediately after elections and Medevdev was widely considered as a weak leader, who never freed himself from Putin's oversight, Medvedev had his own agenda and probably nowhere else it was more visible than in his approach towards information technologies and internet in particular.

[17] http://bd.fom.ru/report/map/projects/internet/internet1133/vesna2011
[18] http://fom.ru/SMI-i-internet/11417

4.2 2008-2012: "Blogger-in-Chief" and his followers

Dmitry Medvedev's approach towards internet was integral part of his general agenda. Laid out in an article "Russia, Forward!", which was published in liberal (and online-only) newspaper *Gazeta.ru*, his *modernization* plan aimed to preserve the basic parameters of the political system built by Putin, but make it more efficient and friendlier towards businesses and citizens (Sakwa 2014, Chapters 3-5). This included, for example, establishing Moscow as an international financial center, police reform, boosting higher education international competitiveness and creation of a functioning e-government (see and overview of Medvedev's reforms in Black 2014; Black and Johns 2013). Medvedev signature project was Skolkovo, a publicly funded, but semi-independently managed high-tech incubator near Moscow. Obviously, the success of these projects was dependent on creative class in major population centers, and IT professionals in particular. Thus establishing a communications channels with these people, who were largely ignored by the blatant Soviet-style TV propaganda, was the first order of business for Medvedev. And unlike in many other areas, he did not hesitate to break with Putin legacy, and put the traditionally solemn and unquestioned presidential speech in the caustic domain of the social networks.

Less then a year after assuming office, in early 2009 Medvedev started a video blog which quickly moved to Livejournal – then Russian main social network and blogging platform. In 2010 he visited Silicon Valley, met Steve Jobs and opened a twitter account at Twitter headquarters in San Francisco. Notably, his account began to follow (in addition to foreign heads of states and Russian officials) several bloggers known for their criticism of the government and newsfeed from radio station *Echo of Moscow* – perhaps the most critical of government among major media outlets in Russia. Finally, in 2011 he opened his Facebook page, which he occasionally used to communicate with its readers on the matters not covered or ill-covered by the official media (such as 2011 protests) using a different, more frank tone. In all social networks he build a large readership, which is typical for heads of states, but still notable since the environment was completely different from the general media environment Medvedev

was used to: here he could not get his message through simply by eliminating competition and controlling the platform and the agenda (Yagodin 2012). In addition, in a rare occasion in 2011 he visited small private TV channel *Rain*, which at the time was mainly accessible online. As a result, Medvedev got permanently associated with blogging and social networks, and even called both in Russia and abroad "Blogger-in-Chief" (see for example West 2010), which simultaneously gave him credit for being up-to-date with the internet age and suggested that his rhetoric translates in little action.

Medvedev was not embarking on social media platforms alone. While it still remained an exception for high-level public officials at the time, several of his aids established significant presence on the social media. In particular, his close aid and economics adviser Arkady Dvorkovich maintains one of the most popular Russian twitter accounts with close to half a million followers; he also has a Facebook page, as does Medvedev's press-secretary Natalya Timakova (who as a former journalist is Facebook friend of many prominent liberal reporters). However, probably even more important was the establishment of large-scale and permanent operation to push pro-government agenda on the web and in social media in particular. Following the long-standing Russian tradition, government action came late, but quickly. Pro-Kremlin youth movements, created to combat color revolution on Moscow streets and squares (Hale 2006) were partially repurposed to push pro-government agenda online. Its leaders (in case of *Nashi* they were called *commissars*) became active bloggers, but they never relied on the persuasion capacity of their messages. Instead they gradually created a network of online support. Until then Russian government presence on social media was very limited. A report by the Berkman Center for Internet and Society at Harvard University, which was published in late 2010 and covered Russian blogosphere – concentrated in Livjournal at the time – in May 2009 – September 2010, found that "pro-government bloggers are not especially prominent and do not constitute their own cluster" (Etling et al. 2010, 3). Moreover, those affiliated with the government "are not central nodes in any of the political or social clusters [...] investigated" (33).

A network of support started with artificial intelligence rather than human effort. Networks of bots got frequently employed first to flood opposition blogs with meaningless or assaultive content. Later they began to push alternative, pro-government messages to top charts and help pro-government bloggers to attract new followers. A report by the Berkman Center noted that "there is a concentration of bloggers affiliated with pro-government youth groups among the Instrumental bloggers [i.e. bots]" (Etling et al. 2010, 3). However, real bloggers soon followed. In less than a year – which also witnessed the transition of the discussion core of the Russian blogosphere from Livejournal to Twitter – pro-government bloggers emerged as a distinct, and indeed, one of the largest clusters on Russian political Twitter (Kelly et al. 2012, 11). This result holds even after filtering out bots and other instrumental accounts, which remained numerous in the pro-government segment.

Continuous monitoring of the Russian blogosphere, undertaken by "Internet in Russian society" program at the Berkman Center for Internet and Society at Harvard University in 2010 – 2014 reveals several distinctive characteristics of pro-government segment in Russian social networks, as compared both to oppositional and "uncommitted" users. First, due to the general weakness and high fragmentation of the Russian opposition, "many active Russian bloggers [...] engage on political topics without 'choosing a team'. [...] most Russian bloggers prefer to declare an independent intellectual posture, and eschew group affiliations" (Etling et al. 2010, 19). In contrast, pro-government bloggers tend to declare their political preferences and affiliation. Moreover, usage of predominantly pro-government hashtags in Twitter was highly concentrated among pro-government users, at least compared to predominately oppositional hashtags, which were more widely used in different clusters. Finally, while pro-government users demonstrate high commitment in terms of the number of hashtag mentions (after the first one), they usually did it in a short time period, producing sharply peaked distribution of hashtag popularity (Barash and Kelly 2012).

As blogosphere remain the most ideologically diverse media environment in Russia, pro-government users experience pressures absent in other media. A comparative study of Russian

blogosphere and TV in the year before the Duma elections of 2011 reveals that such competitive environment forces pro-government bloggers to engage with their adversaries in cases when TV and even newspapers could largely ignore oppositional activity. Etling, Roberts, and Faris (2014) give an example of the oppositional youth retreat in the outskirts of Moscow, which was intended to countervail large government-sponsored youth camp "Seliger". Largely overlooked by the traditional media, it became the subject of the heated discussion between leading oppositional and pro-government bloggers on Twitter.

Online response to hostile (or perceived as such) internet activity through direct engagement with users remained the "weapon of choice" during Medvedev presidency, but certainly it wasn't the only one. Both offline response and attempts to go through the online infrastructure did take place, but the latter were relatively rare and quite limited in their scope and the former was not a part of any systematic internet policy, and as such could not (and wasn't intended to) change the digital media landscape.

Up until the end of Medvedev's presidential term the only type of internet infrastructure infringement known in Russia were relatively brief (lasting up to several days) DDoS attacks on particular web resources (Agora 2011; Freedom House 2011). The first major attack was launched on August 6, 2009 – the first anniversary of the Russia-Georgia 7-days war. The target was pro-Georgian blogger *cyxymu*. The attack was strong enough to significantly disrupt Facebook and completely shut down Twitter and Livejournal (Mills 2009). The series of smaller attacks on various Livejournal blogs and independent media culminated on the weekend of the Russian Duma elections of 2011, when two dozens of the most prominent independent media (including *The New Times*, *Kommersant*, *Echo of Moscow*, *Novaya Gazeta*, *Slon*, etc.), blogs (including the entire Livejournal platform) and, most crucially, election monitors' coordinating portals (including the largest one, GOLOS) were shut down for hours (Roberts and Etling 2011). Later many of the very same resources were attacked during oppositional rallies after elections and in the early 2012.

Importantly, DDoS attacks, unlike filtering (and offline response), could be used not only be the government, but also by the opposition. In early 2012 Russian branch of international cyber activist group *Anonymous* blocked web sites of the Russian government, Kremlin and several major state media, such as *Vesti* and *RIA Novosti*[19]. These attacks, however, did last only several hours (compared to several days in case of Livejournal), and obviously, could impede state response to demonstrations.

Finally, offline response by Russian government to unfriendly internet activity was not yet separated from the general anti-opposition activity and was not legally or organizationally institutionalized. Market regulation and government entrepreneurship was still targeted at traditional media: for example, in 2011 newspaper *Moskovskiye Novosti* was relaunched by state news agency *RIA Novosti*. As it was widely assumed, the project was aimed to provide moderate competition to privately owned (and quite critical) *Vedomosti*, simultaneously being more friendly to Medvedev than most state media, loyal to Putin (Meyer 2011; Barykova and Zotova 2011; A. Morozov 2011). Later that year Medvedev announced the establishment of the Public Television of Russia, which faced no private competition, but shared the second goal with *Moskovskiye Novosti*[20].

Violence and legal action against bloggers were relatively rare and mostly took place in the North Caucuses. Legal restrictions, if any, were imposed under the auspices of the general anti-terrorist laws and orders, mostly having to do with combating Chechen and Dagestani insurgencies. While anti-terrorist rational was often abused for the sake of winning over political enemies in the respective republics, these cases were rarely consequential at the federal level (Simons 2013). In few cases outside the Caucuses prosecutions were largely a regional matter or the result of local security apparatus initiatives rather than implementation of any national strategy. Prominent cases from that time included blogger Savva Terentyev from Komi Republic, who in 2008 was convicted of defamation of the "social group 'law

[19] http://habrahabr.ru/post/143501/; http://lenta.ru/news/2012/05/10/attack/

[20] As mentioned, execution of government projects in media suffers from general government inefficiency and the TV channel went on air only in 2013, long after Medvedev switched offices with Putin.

enforcement personnel'" and sentenced to one year of imprisonment with a probation period of one year after an anti-police comment at Liverjournal. Another prominent case took place in 2009 in the Republic of Tatarstan, where a former government official turned opposition blogger posted false rumor that the governor of the republic has died. He was convicted of libel and defamation of the "social group 'government officials'" and sentenced for 2 years in prison (Yudina 2012).

Institutionalization of offline response, as well as means of control over the online infrastructure happened only after Dmitry Medvedev handed his office back to Vladimir Putin in 2012. However, the process was so quick that already by 2014 the relative importance of different types of government response was reversed: sheer force of offline response and establishment of a comprehensive system of internet filtering rendered the online engagement with users, created by Medvedev, almost irrelevant.

4.3 2012 – : Cracking down and giving up

Compared to transition from Putin to Medvedev in 2008, the reverse transition in 2012 was much less smooth. Announced on Septermber 24, 2011 and immediately nicknamed as "castling", it was met with resentment by both Medvedev supporters and those in opposition to both Medvedev and Putin (Judah 2013). This resentment has transformed in large-scale street protests after parliamentary elections in December, 2011, which were widely considered as rigged[21]. As we mentioned above, close relationship between Putin and Medvedev (culminated in "castling"), did not mean that Medvedev lacked his own agenda. In this case too his response was a program of moderate, but significant political reforms, announced in the Address to the Federal Assembly (Russian equivalent of the State of the Union) in late December of 2011, three weeks after Duma elections, and just after major protests had started. This program included, most importantly, reinstatement of popular elections of Russian governors and elections of MPs in districts (switching back from pure proportional to mixed electoral

[21] Post-election analysis revealed that suspicion was well-grounded(Kobak, Shpilkin, and Pshenichnikov 2012; Enikolopov et al. 2013).

system)(Sakwa 2014, 129-132). These reforms, however, were either striped of any substance (like change in party registration rules) or explicitly reversed (like decriminalization of libel) (Chapters 7-8). Protest activity, on the other hand, was severely restricted after on May 6, 2012 (one day before Putin's inauguration) an opposition rally was dispersed by force (hundreds of people and several dozens of them were subsequently prosecuted for inciting riots and assaulting police).

It is in this context, when the freedom of Russian internet from filtering came to an end (Freedom House 2012, 2013). Already in July of 2012, despite vocal protests, including Russian Wikipedia temporary voluntary shut down, Russian State Duma adopted (and Vladimir Putin signed into law) so-called Internet Restriction Bill (Federal law of Russian Federation no. 139-FZ), which created a continuously updated Russian Internet Blacklist[22]. The list, maintained by the Russian Federal Service for Supervision of Communications, Information Technology and Mass Media (Roskomnadzor), contains domain names which any Russian ISP has to permanently block on the grounds of containing pornography, copyright infringement or "extremist content". Initially, items were to be included in the list per a court order and only if the hosting website fails to remove the content in 24 hours after receiving the notification. However, in December of 2013 new amendments to the Law on Information, Information Technology, and Information Protection provided the Office the Prosecutor General with the authority to block websites without any court order. Moreover, the procedure was changed, so the web page were to be blocked first, and allowed to be accessible again only after it removes the content deemed as "calling for mass disorders, extremist activity, and participation in mass public events, which fail to follow appropriate regulations"[23] (Human Rights Watch 2014).

However, when at the height of Russian-Ukrainian conflict in March of 2014 several oppositional news web sites were blocked, even these loose rules were not followed. On March 13, 2014 *Grani.ru*, *Kasparov.ru* and *EJ.ru*, as well as popular opposition politician (in 2013 he

[22] http://eais.rkn.gov.ru/
[23] Text of the law: http://www.rg.ru/2013/12/30/extrem-site-dok.html.

ran for Moscow mayor and came second) Alexey Navalny's Livejournal blog, were blocked by all ISPs per government order. Since then several suits were brought to courts demanding the reason for the blocking. Journalists and Alexey Navalny asked authorities to identify specific materials on these websites that triggered the blocking, so that the materials could be removed and access reestablished. Throughout 2014 authorities repeatedly denied that they are under any obligation to provide such information and courts repeatedly dismissed the cases. Early in 2015, all three websites and Navalny's Livejournal page remain completely blocked in Russia.

Still, Russian government incomplete control over the online infrastructure significantly impedes its ability to crack down on opposition activity simply by blocking web pages. The greatest thereat, of course, are largest social media platforms – Facebook and Twitter. First, unlike most other web resources, Facebook's and Twitter's individual pages (say, a particular post or user profile) could not be blocked by the filtering software currently available to the Russian authorities (Sivkova 2014). Blocking the entire platforms, on the other hand, is still considered undesirable: it would further hurt Putin's regime reputation abroad and simultaneously hurt and potentially antagonize a large number of politically indifferent (and regime-friendly) users in Russia. In a rare event, a public official, Roskomnadzor deputy head Maxim Ksenzov, who speculated over such possibility, was publicly disproved by Prime Minister Dmitry Medvedev in a Facebook post[24] and later was formally reprimanded. Of course, instead of blocking them, Russian government could ask them to police themselves and remove access to certain pages at least for users inside Russia. However, unlike Vkontakte, foreign social networks could easily ignore such orders. For example, in December of 2014 authorities requested Facebook and Vkontakte to block access to pages, allowing supporters of Alexey Navalny to register for a rally protesting his looming criminal conviction and receive updates about the place and time of the event. Vkontakte blocked the page and all subsequent attempts to create a copy, posting a warning that "This page is blocked upon receiving a Roskomndazor notification of restricting access to information, which contains

[24] https://www.facebook.com/Dmitry.Medvedev/posts/10152047885946851

calls to participate in mass public events, which fail to follow appropriate regulations, as per the request of the Office of the Prosecutor General of Russia." [25]. Facebook also blocked access to a similar page inside Russia[26], but after a huge outcry in Western media, refused to block any other pages. Moreover, some Russian media outlets, which were afraid to report the scheduling of the event itself, covered the Roskomnadzor order and social networks response. As a result, more people learned about the event and the new event page opened on Facebook attracted even more people[27].

Given that second page attracted more than 33 thousands people, who stated that they are "going to the rally" (plus almost 6 thousands, who stated that they are "likely going"), it's not surprising that the authorities resorted to offline response: they simply changed the data of the return proceedings to two weeks earlier. The new date was the day before the largest Russian holiday (The New Year's Eve) and Navalny was informed less than 24 hours in advance. While the third event also attracted considerable number of supporters, combination of suddenness, cold weather and pre-holidays preparation likely reduced the turnout.

Offline response was certainly not limited to ad hoc solutions just described. Instead, government complete control over law enforcement apparatus and law making was actively used to augment its limited ability to censor social media platforms. Criminalization of online activity was first implemented trough targeted amendments to existing criminal law, but was soon institutionalized in a dedicated law. Using media to spread information deemed extremist was always an aggravating circumstance in Russian criminal law. Laws missing such provision were sooner or later corrected: for instance, when in 2011 punishment for Article 280 of the Criminal Code was severed, using mass media for "extremism propaganda" became an aggravating circumstance. However, when just two years later, in 2013, a new extremism crime appeared in the Criminal Code (Article 280.1, Public Appeals to the Violation of the

[25] https://vk.com/blank.php?rkn=32274605; It should be noted that according to those "appropriate regulations" authorities could not be notified about the upcoming rally earlier than 15 days in advance. The page was blocked 26 days before the event it announced was scheduled to take place.
[26] https://www.facebook.com/events/417200101767938
[27] https://www.facebook.com/events/406603402849183

Territorial Integrity of the Russian Federation), using "mass media, including telecommunication networks (including 'Internet')" was added as the aggravating circumstance (Agora 2012, 2013).

In May, 2014 Vladimir Putin signed into law a requirement for any blogger with the daily readership in excess of 3000 people to register with the government and reveal her true identity and email address [28]. In addition, bloggers will be held accountable for failure to verify the information they "spread", have to keep archives of their postings and follow laws which regulate news production during electoral campaigns. However, institutionalized regulations expectedly are much less effective then targeted actions: in half a year after the law came into force just 369 people got registered with Roskomnadzor (Rothrock 2015) and the only known real consequence is the shut down of Intel's forum for developers - hardly a platform of political significance, which was closed by Intel voluntarily out of precaution (Lunden 2015). Among the reasons is unclear definition of "readership": Roskomnadzor guidelines on the subject[29] call to use rigorous "page views" count (rather than hits, number of friends or followers or any other metric), but not all platforms generate such statistics, and it is especially hard to do in social networks.

Using loyal business groups to restructure the online media market proved much more reliable tool to ensure that at least Russian major platforms are under control. Hitherto mostly focused on traditional media (TV and press), power brokers in the Presidential Administration, Ministry of Communications and largest media conglomerates have been increasingly preoccupied with online news outlets and platforms. The methods they used were not much different. Two cases are particularly revealing. In 2014 billionaire Alexander Mamut fired the editor-in-chief of the most popular Russian online news portal *Lenta.ru*, allegedly on the grounds of insufficiently "pro-Russian" coverage of the Ukrainian revolution of 2013-2014. Complete lock out of the entire editorial staff was strikingly similar to the one at the NTV

[28] According to the survey reported by Alexanyan et al. (2012), even without any legal requirement Russian bloggers rarely conceal their identity. They do use pseudonyms (following internet tradition), but usually alongside, not instead of their real names. This is particularly true for politically-engaged bloggers.

[29] http://rkn.gov.ru/docs/prikaz_Roskomnadzora_ot_09.07.2014_N_99.pdf

channel in 2001 and countless others since then. However – and here comes the difference between TV and website – this fired team of journalists was able to relaunch their media. The insurance of their independence and security from outside pressure was physical relocation of most of the editorial staff to neighboring Baltic country of Latvia and opening website in *.io* domain zone, which belongs to British Indian Ocean Territory and is administered by a UK company. New media name *Meduza* (Russian for jellyfish) matches the geographical location of its domain.

Hostile takeover of Vkontakte in 2013-2014 by Kremlin-affiliated businessmen also followed the approach which earlier successfully secured the loyalty of various media outlets (such as *Izvestia* and *Kommersant*): involuntary ownership transfer, usually, compensated at the market rate (dependent on the cooperation of the former owner). This transfer usually came after former owners and/or managers refuse to cooperate in politically sensitive matters for too long. Vkontakte received requests from FSB similar to the ones described (and followed) in case of pro-Navalny rally in late 2014 for years. Specifically, requests to remove pro-Navalny groups came first in the wake of large-scale protests after Duma elections in 2011 (Razumovskaya 2011), but Vkontakte owner and CEO, libertarian internet-guru Pavel Durov refused to comply. However, when in early 2014 Vkontakte was served with request to disclosure personal data of administrators of Euromaidan-related pages in Vkontakte [30], government did not take no for an answer. Durov had to sell what was left of his share, resigned and left the country (Ries 2014; Kononov 2014; Lunden 2014).

Lesson of Vkontakte was taken seriously not only by Russian media managers and owners, who wanted to keep their positions and businesses. Foreign companies which wanted to be able to refuse involuntary cooperation with Russian government had to assess if they had any vulnerable assets in Russia. For instance, Facebook ability to change its response and refuse to block any more groups in late 2014 episode of pro-Navalny rally was secured by company's retreat from Russia. Its development office was closed and the entire engineering

[30] https://vk.com/wall1_45621.

team was invited (and accepted the offer) to move to Google offices in Europe and elsewhere. While reasons for this move were not disclosed, observers assumed company concerns with the potential access to Google code to the Russian government and even coercive methods to get this access through pressure on individual engineers (Bershidsky 2014).

What in this context happened with the online response? Did response offline and through the online infrastructure supplanted any serious engagement with users on behalf of the government? Yes, but it was not the only factor. Another key change was in the target audience of government online effort. If Medvedev was trying to build a coalition around values of modernization and reformist policy agenda, "Putin Redux" entailed rather dramatic change in regime ideology, not just compared to Medvedev years, but also compared to Putin's first two terms (Sakwa 2014). Reorientation towards conservative, even traditionalist values in domestic policy was paired with expansionist, revanchist foreign policy (Smyth and Sobolева 2014; Snyder 2014). This change had implications of truly historical scale, but one of less known consequences was reorientation of online propaganda machine from winning over neutral or even already oppositionally-inclined users towards protecting wider public, those receiving most of news news from TV but starting to use internet for entertainment or consumption, from the dangerous influence of the opposition voices. Typical example from the same episode with pro-Navalny rally in late 2014 was Youtube videos with prominent Navalny supporters, who were showing on air the web address of the supposedly pro-Navalny website with information about the rally. In reality they were fooled and the address was leading to page full of anti-Navalny videos. These videos and apparent endorsement of them by famous artists and journalist were then promoted on social media.

Such provocations obviously could not build a reputation that Medvedev, or his economic advisor Arkady Dvorkovich, or Perm governor in 2004 - 2012 Oleg Chirkunov were seeking to build online. Their goal is not to engage, agitate and invite for discussion; it is to disorganize, discourage and mute opposition. And this goal is much better served by filtering technologies and targeted prosecution of influential bloggers. Extensive online debates be-

tween oppositional politicians and pro-Putin "Nashi" youth movement, which occasionally happened before 2012, are no longer possible. With the gradual, but persistent political retreat of Medvedev's team, and government officials of liberal and pro-Western inclination in general (which in many cases includes leaving public service or even the country for good), the government presence online would not vanish. However the government officials and their speakers, both paid and volunteers, would speak with themselves and their most loyal supporters. Government would be online, but it would not be responding to anybody online, much less waiting for response from anyone.

5 Online engagement: preliminary analysis

5.1 Statistical methods for bot detection: a useful starting point

Online engagement is a complex phenomenon, ranging from completely automated bots that produce large volumes of gibberish to flood popular communications platforms to high-profile paid bloggers with independent reputations, who send nuanced, targeted messages to different groups of the public. While ultimately, we intend to study all of these phenomena, we start with bots.

There are three main reasons for this choice. First, bots produce by far the largest volume of content and without the tools to remove that, it would be almost impossible to study the human-generated content. Second, the only practical way to identify bots is by using automated algorithms; starting the empirical part of our research in this manner has the advantage of therefore creating an objective – and replicable – approach that can be employed in future analysis. There is, however, a third, less methodolically inspired reason to start with bots, which is that they are both important and interesting objects to study. While social media provide citizens and politicians with new and powerful tools for expressing their political beliefs and preferences, affecting the political agenda, mobilizing supporters, and organizing political actions, they also bring the challenge of differentiating between real political communication and interaction with computer programs that imitate human activity. Even

though the latter are not necessarily malicious and may even help citizens hold politicians accountable (by sending out news about recent corruption scandals, for instance), they are potentially dangerous for a number of reasons. First, they may disseminate the wrong information (or deliberate misinformation) in a systematic way, affecting the routine functioning of local governments, as it happened in Louisiana on September 11, 2014 (Chen 2015). Another possibility is that the wrong information could misinform voters about politicians' activities, which could be potentially consequential in election periods. Second, bots may misrepresent the voters' preferences and their popularity, thus providing politicians with wrong cues regarding citizen preferences. Finally, bots may be used as a new weapon in cyberwars useful for preventing enemies from spreading their messages by swamping them with spam.

Bot detection is a relatively new topic in political science emerging from the burgeoning research of political communication on social media. Scant literature that exists on the subject mostly borrows methods developed in computer science to detect spam, i.e. undesired automated emails using primarily certain features of their text. Bot detection in social media is a different, but closely related, task (Chu et al. 2012) that can be solved with a wider range of techniques making use of both text and non-textual account (also referred to as "metadata") information. Nevertheless, bot detection in social media is still regarded as a challenging task within the computer science community, making Boshmaf et al. (2011) claim that 80% of bots are undetectable.

In this paper, we both adopt some of the existing bot-detection techniques to identify bots on Twitter that are engaged in writing tweets about Russian politics, and use expert coding to validate the results. For that purpose, we use a large set of tweets collected by the Social Media and Political Participation (SMaPP) lab at NYU by searching for predefined key-words and extracting from Twitter API both identified tweets and their authors' account information.

5.2 Methods of bot detection

The first – and arguably most crucial - step in any analysis of the behavior of bots is finding those bots in the purpose. We employ three types of bot-detection techniques – two of which have been used in other bot detection exercise outside of politics and one of which we believe is novel – to find bots in our collection of tweets: looking at how regular the time intervals are between tweets; examinig patterns in the numbers of friend and followers of accounts; and focusing on the propensity to post identical tweets to other users. In the remainder of this section, we desribe each of these in turn.

Entropy of inter-tweeting time intervals. Our first technique, which we label as our "entropy measure" is predicated on the idea that bots – contrary to humans – show a much higher regularity in their activity on Twitter. In the most simple case, bots may be programmed to send tweets every k seconds or using some other, more sophisticated, schedule. On the contrary, humans' tweeting activity is much less regular. These differences in predictability may be captured by entropy which is a measure of uncertainty popular in computer science and information theory.

In order to compute the entropy, we created a list of all accounts that appear at least once in our collection of tweets. Then, for every account found, we ordered all collected tweets in time, computed the length of time intervals between consecutive tweets, and used those time intervals to compute averaged entropy as follows:

$$entropy_i = \frac{1}{J_i} \sum_{j=1}^{J_i} p_j^{(i)} \times log_2(p_j^{(i)}),$$

where p_j is the probability of jth interval for the ith account; J_i is the total number of time intervals for the ith account.

Followers/friends ratio. Although entropy proved to be the most informative bot-detection technique in Chu et al. (2012), we did not restrict ourselves to this measure due to a special

type of data we analyze. Instead of having all the tweets from a given set of accounts, we have all the tweets mentioning predefined keywords. Thus, it is likely that we do not have many tweets sent by most of the accounts in our collection. Hence, the entropy measure might be noisy in our case, although to the extent that we are looking for bots designed to produce *political content*, this might not be too much of a problem, (i.e., if we think these bots are only tweeting about politics, we should get all of their tweets).[31]

Our second method of detection, however, does not rely on the posting of tweets, and thus can not be affected by whether or not we have all of our users' tweets in the dataset. Here, we rely on the the idea that bots' Twitter profile looks different than a valid user's, first, because bots would have much less followers. Indeed, humans would probably refuse to follow a bot that does not show signs of normal human online activity. At the same time, bots would tend to follow lots of users (i.e. have many friends, using Twitter terminology) in the hope that some of them will accidentally follow them back. Thus, we expect that some bots will tend to have a very small followers/friends ratio defined simply as

$$ratio_i = \frac{card(followers_i)}{card(friends_i)},$$

where $card(followers_i)$ denotes the number of accounts that follow ith account, and $card(friends_i)$ stands for the number of accounts that are followed by ith account.

At the same time, an interesting case is a different type of bots involving the ones that do not follow anybody. These bots may be just news seeds and restrict their activity to spreading news on Twitter. Obviously, the followers/friends ratio is undefined for these bots due to division by zero, and we code them with a special system-missing value.

Identical tweets. Our final bot-detection technique involves identifying accounts sending the same tweet. This may be the same account that repeatedly sends out the same tweet, or

[31] This might, however, cause us to miss bots in our "identical text" test that we describe below; again, however, if all of the identical text in question contains our political keywords, then this may not be a problem either.

different accounts. The motivation for this bot-detection technique is rather straightforward, since bots are programs and have some pre-programmed data-generating process for their tweets. Hence, the least sophisticated bots would just repeated some of their tweets over and over again. [32]

In order to identify candidate bots using this approach, we created a dictionary of tweets in our collection and associated with every tweet both its frequency (i.e. the number of times this tweet shows up in the collection) and a list of Twitter accounts that sent out that tweet.

Even though there is a non-zero probability that some tweet may appear multiple times by pure chance, given the maximum tweet length is 140 characters, the probability that a reasonably long tweet appears hundreds or even dozens of times in a data set like ours by chance alone is negligibly small.

5.3 Data

Our dataset includes more than 18 million tweets, posted by around 3.8 million Twitter users between February 6, 2014 and November 5, 2014. There is a large variation in the number of tweets written by different users in our collection, ranging from 1 to almost 73 000. The data were collected by scraping all tweets that contain a predefined list of keywords and hashtags with Twitter API. These words were chosen to represent a broad set of politically relevant terms, including major politicians' names (Putin, Medvedev, Navalny, Khodorkovsky, Udaltsov, etc.), events (Sochi olympics, opposition rallies at Bolotnaya square in Moscow, "Direct Line with Vladimir Putin", etc) and slogans ("Party of thieves and crooks", "Sobyanin is our mayor", "Stop feeding the Caucasus", etc.). Hashtags and keywords were drawn from all over the political spectrum, including Putin and United Russian, loyal and radical opposition, Russian nationalists and others, up to and including Pussy Riot.

[32] Some of these primitive have even attacked one of the co-authors of this paper on occasion following political tweets related to Russian politics!

We use three methods outlined above to identify in the dataset a list of suspicious accounts that we regard as candidates for bots. Quantitative measures described above require selecting a threshold that would separate legitimate human users from bots. As any other case of threshold selection, this is not a straightforward problem since no theory has so far been developed to justify the choice. Given this lack of theory-driven knowledge, we followed an empirical approach having two main considerations in mind. First, the threshold should be small enough so that we could consider a large set of accounts. Second, the threshold should not be too small so that qualitative validation of the suspicious accounts would be feasible.

Given the distribution the ratio measure (see Figure 1), we regard as suspicious those accounts that have a followers/friends ratio between 0 and 0.001. The zero value would mean that nobody follows that account, which is really suspicious. The value 0.001 implies that an account follows 1000 times more accounts than follow it. That is a large enough disproportion in the number of followers and friends to raise doubts in the "humanity" of the user. Application of 0.001 as a threshold value for the ratio measure gives us 208 515 suspicious accounts, which is 5% of the total number of accounts found in our collection.

Another set of suspicious accounts comes from the users for whom the ratio is undefined (not shown on Figure 1), since they have no friends. This set contains 202 705 accounts which is also 5% of the collection.

In order to make the verification task feasible, we restricted our attention to a subset of suspicious accounts that had reasonably high activity on Twitter. The distribution of the number of tweets that the suspicious accounts produced is extremely skewed to the right, as it follows from Table 1, so we restricted our attention to those accounts that sent out not less than 300 tweets. Thus, we ended up with 254 accounts with a low ratio and 194 accounts with an undefined ratio.

As one can see, the chosen threshold of 300 tweets provides us with a reasonably large set of candidates for bots that is, though, feasible to verify qualitatively by expert evaluation.

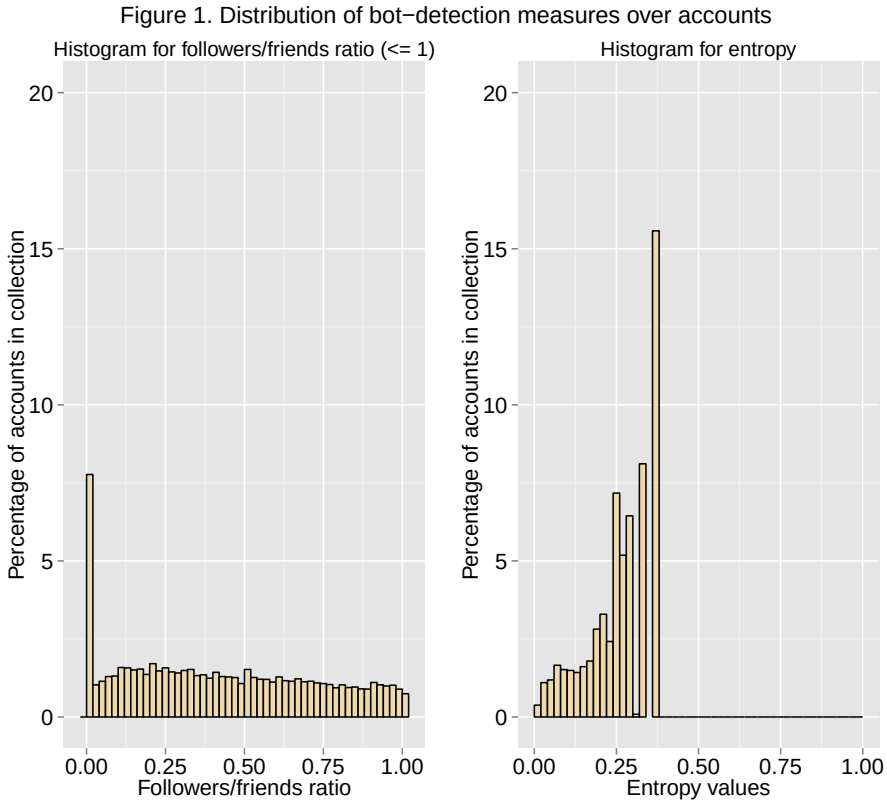

Figure 1: Distribution of Twitter accounts in the dataset by the followers/friends ratio (left panel) and timing of activity entropy (right panel)

Table 1: Tweet counts from suspicious accounts

Descriptive Statistics	Accounts with ratio <= 0.001	Accounts with no friends
min	1	1
$q_{0.25}$	1	1
median	1	1
$q_{0.75}$	3	2
$q_{0.99}$	55	82
$q_{0.999}$	330	295
max	16 581	72 990

Note: Summary statistics refer to the number of tweets written by suspicious accounts and found in our collection.

Moving that threshold down, to 200, would substantially increase the number of suspicious

accounts to be verified[33] and become much less feasible. At the same time, there is no resource gain in moving the threshold up.

One more set of accounts to be considered is produced by the entropy measure. Here, suspicious are accounts with the smallest values of the measure, since their online activity shows high regularity uncommon for humans. In order to ensure comparability of the accounts selected using the entropy measure with those we chose using the ratio measure, we also restricted our attention to those account that produced at least 300 tweets throughout the period under study. From those, we selected for further analysis 300 accounts with the lowest entropy values.

The final source of candidates for bots is the set of accounts engaged in writing repeating tweets. As we explained above, this means an account either writes some tweet multiple times, or that it writes the same tweet as some other account. In order to make the verification step feasible, we restricted our attention to 183 accounts that produced at least 500 repeating tweets.

Thus, we end up with four sets of suspicious Twitter accounts that we identified using different bot-detection techniques. How similar are these sets? Table 2 presents a cross-table with entries showing the number of accounts that are common for two sets. Since bots are programs, their online behavior depends on the creativity of their creators and specific purposes they were created for. Thus, unsurprisingly, different bot-detection techniques catch different types of bots, as one can see from Table 2. The fact that four suspicious sets do not have large proportions of accounts in common (in total, there are only 113 duplicates among 931 accounts we identified) implies that we probably managed to identify different types of bots that may be used for different purposes.

[33] Moving the tweet count threshold to 200 increases the number of suspicious accounts with a low followers/friends ratio up to 583 accounts. The number of accounts with no friends goes up to 581.

Table 2: Numbers of accounts that are common for four sets

	No friends	Low ratio	Entropy	Repeating tweets	**Totals**
No friends	—	0[a]	13	16	**194**
Low ratio		—	24	14	**254**
Entropy			—	56	**300**
Repeating tweets				—	**183**
Totals	**194**	**254**	**300**	**183**	

Note: Entries are numbers of Twitter accounts that are common to two corresponding sets.
[a] These two sets are exclusive, since an accounts belongs to no friends set if the ratio is undefined.

5.4 Verification and assessment

In order to assess the accuracy of our bot detection algorithms, we looked at each of 818 unique "bot candidates" we identified. Expert verification of the suspicious accounts revealed high accuracy of the methods employed. 570 accounts were classified as certainly bots and additional 67 as "likely bots". The latter category includes accounts with activity very similar to clear bots, but with some, quantitatively minor, element of human activity, such as seemingly original replies to other users. Only 28 accounts were not classified as bots, and even those exhibit behavior – such as posting tens or hundreds of thousands of twits around the clock on a diverse set of topics – that suggests that they are (in most cases) not personal, but institutional accounts. Indeed, several of them were official accounts of large news media (including NTV, NEWSRU and RBK), while several others were unofficial media accounts.

Moreover, by 2015 126 accounts from our 2014 dataset were suspended (additional 27 were deleted by users themselves). Given that accounts are usually suspended for spam activity, it further confirms that our methods accurately captured bot activity. Table 3 highlights the high accuracy of the quantitative results.

Table 3 also reveals high individual reliability of three of the methods we developed: accounts with no friends or low followers to friends ratio, as well as those with highly repetitive content, almost never belong to humans. Low entropy of posting timing is the least reliable predictor of expert classification as bot, but most of the accounts in "not bots" category

belonged to media outlets and low entropy in their case is simply explained by the constant flow of tweets posted one after another during the working hours or around the clock.

Table 3: Results of suspicious accounts verification

	No friends	Low ratio	Entropy	Repeating tweets
Certainly bots	160	172	169	173
Probably bots	3	6	60	0
Not bots	4	0	21	1
Suspended	25	68	34	8
Deleted	2	8	16	1
Totals	**194**	**254**	**300**	**183**

Note: Entries are frequencies.

From now on, we proceed analyzing the accounts that were classified as certainly bots, thus ending up with 570 accounts.

These verified bots could be further classified by type of content they feature (see Table 4). The lion share of the them are primitive newsfeeds which simply post news titles (or, in much more rare case, excerpts from news stories) in bulk. Such accounts could be clearly identified as bots at the first glimpse. Their handles and names are usually meaningless and often allow to identify a series of identical bots simply by comparing their names, which are similarly spelled or contain consecutive numbers. These accounts typically lack the description and their userpics appear to be lifted from the web image search. They post large volume of tweets, mostly don't follow other accounts and rarely have substantial number of followers. Importantly for our purposes, these bots, while not necessarily impartial, by their very nature dilute any ideological message of their source. First, newsfeeds (even if supplied by a newspaper, not professional news agency) are usually less biased and ideologically charged than feature stories. Whatever bias they might have is further diluted by reducing the news piece to the title or one-line excerpt. Finally, and most crucially, the sheer volume of information and diversity if topics covered prevernts these bots from communicating any coherent message.

Table 4: Verified bots by type

	No friends	Low ratio	Entropy	Repeating tweets	**Totals**[a]
Newsfeeds	119	71	87	158	**344**
Rich content, political	4	59	60	1	**121**
Non-political	36	41	23	14	**105**

Note: Entries are frequencies. [a] Total values are not equal to row sums because items belonging to several categories were counted only once.

Bots deemed in Table 4 as "rich content, political" could not be more different from the newsfeeds. The difference starts from how they look like: these bots are visibly much more interesting, which reflects more effort and resources put in their maintenance. They usually feature a high-quality, often custom-made userpic and background picture, witty name and handle, informative description. Many of these bots contain large number (up to tens of thousands) of embedded pictures and videos and often use multiple hashtags. As a rule, such bots retweet other users, and it seems that they retweet both other bots and real people. This makes their content to look richer than often text-only (or, at best, hyperlink-augmented) newsfeed accounts: retweets always contain a link to and userpic of another user plus date and time of the original tweet. As mentioned above, retweeting allows not only to spread particular message, but also add to the visibility and authority of the source account, thus making bots automated support staff for paid and genuine Twitter propagandists. These retweets and other content appear to be well-curated to cover a certain topic, region, political figure, etc. This enhances these bots' ability to send a coherent message to readers. In addition, these bots usually mimic the human behavior: they follow other users and in many cases are able to accumulate large number of followers, including real users (surprisingly, some of them are followed by experienced opinion leaders on Twitter). Thus in their activity these bots could rely on the network effects of Twitter.

The third type of bots we identified are both newsfeeds and rich-content bots that are non-political and thus largely irrelevant for our analysis. In our dataset the majority of these bots were related to Sochi Olympic Games, but instead of covering their political dimension, were preoccupied with sports events.

5.5 Sentiment analysis of verified political bots

Next step in our inquiry is qualitative sentiment analysis of content posted by verified bots. While studying newsfeeds and comparing political and non-political bots are interesting avenues for future analysis, for now we will focus on 121 well-networked verified bots containing politically relevant, rich content.

These most interesting accounts in our dataset we classified into three groups: pro-Kremlin, pro-opposition, pro-Ukrainian. The motivation behind the choice of these sentiment groups is as follows. Pro-Kremlin bots are the most well-known bots in the Russian segment of social media, widely discussed (along with trolls) in newspapers around the world. However, there is anecdotal evidence that anti-Kremlin groups may also used that tool in social media for propaganda and mobilization purposes. Since the timing of our data collection coincided with the political crisis in Ukraine and further Russian involvement, we split the anti-Kremlin bots into two categories: pro-opposition, which encompasses bots tweeting mainly about the intra-Russian political agenda, and pro-Ukrainian, whose Twitter activity is related to the developments in Ukraine. The results are shown in Table 5.

We should note that while for now this classification remains provisional, it is highly reliable thanks to the nature of the content that bots generate. In none of these 121 accounts there was a small degree of doubt regarding their political orientation. Bots spread the most blunt message possible and pick up the most radical, viral and vicious content that serves their agenda.

Table 5: Sentiment analysis of verified political bots

	No friends	Low ratio	Entropy	Repeating tweets	**Totals**[a]
pro-Kremlin	0	22	15	1	**36**
pro-opposition	0	8	29	0	**37**
pro-Ukrainian	4	29	16	0	**48**

Note: Entries are frequencies. [a] Total values are not equal to row sums because items belonging to several categories were counted only once.

Table 5 highlights interesting and unexpected findings. First, anti-Kremlin bots, comprising pro-opposition and pro-Ukrainian bots, are twice as numerous as pro-Kremlin. This result may seem unexpected given the mass media's clamor about Kremlin's social media propaganda campaigns. At the same time, this result might imply that Kremlin prefers more sophisticated and expensive online propaganda techniques like payed trolls, while the Russian opposition and pro-Ukrainian users may so far lack resources to employ those techniques.

Figure 2. Tweeting activity of political bots

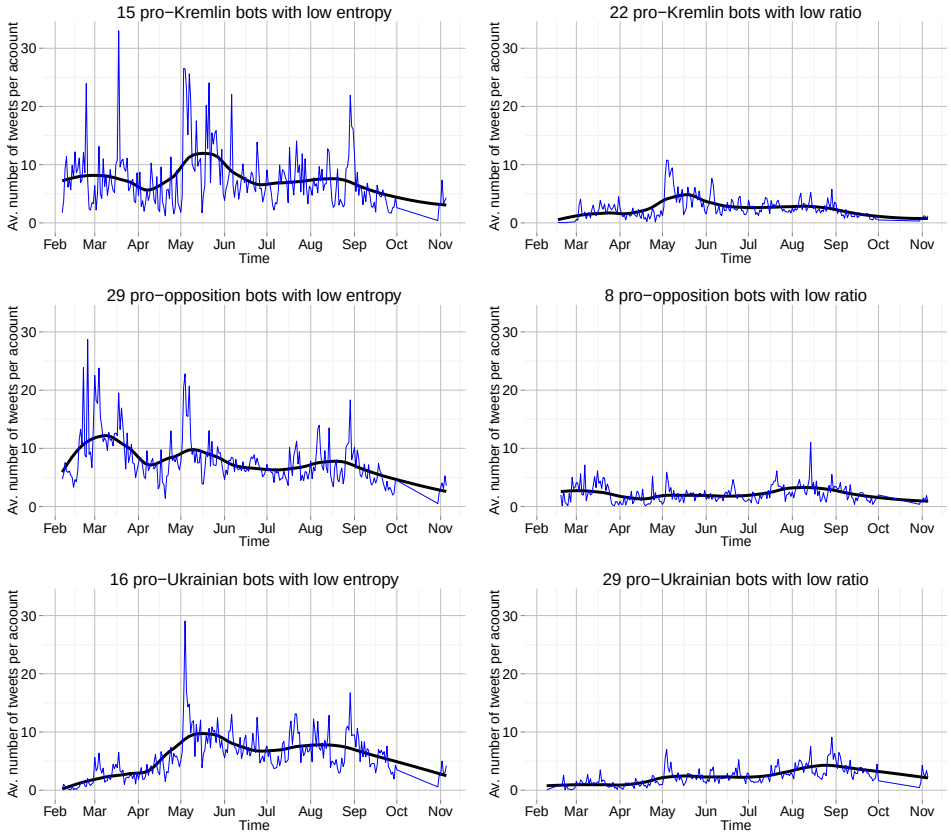

Figure 2 presents political bots' average Twitter activity per day over time. One can see from the graph that, overall, bots that are pre-programmed to tweet in specific moments of time are much more active on Twitter that other types of bots. This is true for all sentiment groups and might reveal a special purpose of these bots, that is to produce content on a regular basis. However, there is still a lot of variation in the average number of tweets over time, meaning that those time pre-programmed bots are actively regulated.

Even though the time-series do not look very similar at first glance, the correlation analysis reveals important patterns (see Table 6). One can see that two types of pro-Kremlin bots show a substantially high correlation of their tweeting activity. The same is true about pro-Ukrainian bots, but is not true about pro-opposition bots. Furthermore, there is a high correlation in the tweeting activity of the pro-Kremlin bots on the one hand, and pro-Ukrainian bots with low entropy. This suggests that pro-Kremlin and pro-Ukrainian bots' activity might have been coordinated in a centralized manner, and the bots tended to react either to the same developments off- or online. At the same time, Russian pro-opposition bots do not show any noticeable coordination.

Table 6: Correlation table for tweeting activity time-series

	pro-Kremlin entropy bots	pro-Kremlin low ratio bots	pro-opposition entropy bots	pro-opposition low ratio bots	pro-Ukrainian entropy bots	pro-Ukrainian low ratio bots
pro-Kremlin entropy bots	1.00	0.73	0.59	0.35	0.60	0.38
pro-Kremlin low ratio bots		1.00	0.43	0.31	0.75	0.41
pro-opposition entropy bots			1.00	0.48	0.34	0.19
pro-opposition low ratio bots				1.00	0.42	0.47
pro-Ukrainian entropy bots					1.00	0.69
pro-Ukrainian low ratio bots						1.00

Note: Entries are product-moment correlation coefficients. All coefficients are significant at $p < 0.01$.

At the same time, the focus of tweets written by the pro-Kremlin bots differs significantly both from other bots' and among themselves. Table 7 shows Jaccard similarities between sets of 20 most popular hashtags used in every group of bots. As we can see, pro-Kremlin bot identified with the entropy measure are very dissimilar than all other groups of bots, even the other pro-Kremlin group. Meanwhile, pro-opposition and pro-Ukrainian bots cluster all together and focus on the same topics, as revealed by the hashtags used. The list of the hashtags is given in the Appendix and suggests that pro-Kremlin bots with the low entropy are

mainly focused on the international perception of Russia. These bots frequently use hashtags about the U.S., sanctions and Olympic games, along with more common hashtags related to Ukraine. The other group of pro-Kremlin bots, according to the hashtags, is almost exclusively focused on the dramatic developments in Eastern Ukraine and the situation around Crimea. Unsurprisingly, pro-Ukrainian bots use a lot of hashtags about Euromaidan protests, Crimea and Ukraine, while pro-opposition bots split their attention and talk also about internal Russian issues like Chechnya and protests in Russian cities. Nevertheless, there is one hashtag that is common to all groups of bots – it is the presence of hashtags related to Putin.

Table 7: Jaccard similarities of 20 most popular hashtags across groups of bots

	pro-Kremlin entropy bots	pro-Kremlin low ratio bots	pro-opposition entropy bots	pro-opposition low ratio bots	pro-Ukrainian entropy bots	pro-Ukrainian low ratio bots
pro-Kremlin entropy bots	1.00	0.18	0.08	0.14	0.11	0.11
pro-Kremlin low ratio bots		1.00	0.21	0.21	0.33	0.25
pro-opposition entropy bots			1.00	0.25	0.43	0.67
pro-opposition low ratio bots				1.00	0.43	0.38
pro-Ukrainian entropy bots					1.00	0.67
pro-Ukrainian low ratio bots						1.00

Note: Entries are Jaccard similarities between sets of 20 most popular hashtags in every group of bots. Jaccard similarities range from 0 to 1, where 0 refers to a pair of sets that have no common element, while 1 refers to a pair of sets that are identical.

6 Lessons from the Russian experience: A Discussion

Creating a comprehensive theory of government internet policy choice is not the subject of this paper, but we want to briefly outline the key factors of this choice as they appear in the evolution of Russian government policy towards internet. Of course, as any individual case, it emphasizes some factors at the expense of others. For instance, some important parameters remain constant, while others change rapidly and thus provide more latitude for theoretical speculation. We list factors in no particular order as their relative significance remains to be

explored.

Probably the most obvious factor is internet penetration in the country. Until it reaches certain threshold it would be considered politically irrelevant, as it was throughout Putin's first two terms. Here the regional variation could be also vitally important. If the vast majority of ruling party electorate is living in provincial towns and villages, high internet penetration in capital city creates less of a concern, unless it poses an immediate threat of a rebellion or mass protests (as it happened in 2011-2012 in Moscow).

In addition to electoral significance (or lack of thereof), internet could be important for economic development and for the smooth functioning of government institutions. Relatively cautious treatment of Yandex, as compared to Vkontakte, probably comes not from just somewhat greater threat of online coordination (vs. mere providing information through news aggregation, which is done by Yandex), but also from an important role Yandex plays at Russian IT market. The latter is of the essence especially in those numerous non-democracies which are at the mercy of volatile commodity prices. Similarly in these countries bureaucracies, rotten by corruption, often fail to execute even the will of the autocrat. Thus monitoring capacity provided by internet is even more important than in democracies, where old mechanisms of representation ensure quality public services. If IT industry would be hit by it, repressive internet policy could become self-hurting. This dilemma exemplifies the balance autocratic regimes should struck between employing modern technologies for their own advantage and maintaining tight control over the information flows.

Next set of factors has to do with resources available for the regime to pursue the option(s) it chooses. Non-democracies are rarely restricted in offline options they can pursue, as full force of law enforcement apparatus (as well as extra-legal enforcement tools) are readily available to them. Technical capabilities vary much more across countries. Russia is almost perfect environment for the autocrat in this respect. First, skilled IT specialists are abundant to create high-quality filtering tools. Similarly, PR specialists, many with strong journalistic

background, can devise and execute a plan of online engagement with any audience. Therefore we can conclude that access to such resources was not important factor of policy choices in the Russian case. Russia was also comparatively well positioned in terms of major web platforms it could potentially control. This "control capacity" would probably be best estimated by the share local online media and platforms have on the countries online market. While in Russia it is lower than in China, it is higher than in most other non-democracies in the world. Still, as we have seen, while control over Vkontakte was easily established through the combination of relatively benign use of law enforcement and market intervention by loyal businesses, threats posed by opposition communities in Facebook and Livejournal and on Twitter were not completely neutralized even by combination of filtering and much more gross manipulation of the law.

Thirdly, policy towards online media and social networks is intimately related with media policy and politics in general. The direction of this relationship is unclear though. On the one hand, regime can make no distinction between traditional and digital environments and increase pressure online as it tightens the screws offline. Indeed, as similarities between the expropriation of Vladimir Gusinsky's media in 2001 and *Lenta.ru*'s lockout in 2014 reveal, experience of repression in traditional media could be easily transferred at least to the digital media if not social networks. On the other hand, internet could be this alternative space, where few disillusioned citizens would (as autocrats hope) harmlessly vent their frustration as their access to mainstream media rapidly evaporates. In addition, internet could be valuable source of information about public opinion and subordinates' conduct, when all other channels broadcast the party line. These considerations seems to explain virtual lack of filtering in Runet until as late as 2010. Finally, free internet in the absence of free offline media could work as window dressing for the regime: until he began to implement similar internet policies after 2012, Putin often publicly and privately bragged about the absence of Chinese style internet censorship in Russia. However, these arguments do not trump the security of the regime concerns, which in Russia took the leading role after 2011-2012 wave of protests.

Internet penetration, electoral significance of the "net citizens", differences (and similarities) between offline and online media in non-democracies as well as internet role in economic development and government functioning are all important factors of internet policy choice and are often discussed in the literature on the subject. However, Russian experience as presented here calls for taking several other, rarely considered factors into the account. These factors have to do with the domestic politics in non-democracies.

First dimension of the domestic politics important for the internet policy choice is the power dynamics and competition inside the ruling group. If the media environment is monopolized, parts of selectorate with disadvantaged access would seek alternative avenues to build support and exert influence at home and abroad. Medvedev's era serves as a perfect illustration. As opposed to economic controls – traditionally the responsibility of the government (headed by Putin at the time) – nominally Medvdev had all the keys to media management. Nevertheless, in reality the most powerful media – national TV – remained loyal to Putin throughout Medvedev's term in office. Medvedev then turned to online media, from digital TV Rain to social networks to online-only news outlets. His goal there was to build a reputation and hence a community of supporters, which determined his virtual abstention from restrictive policies. As the power was consolidated in Putin's hands again, the balance between online threats and opportunities had changed, moving policy into a different direction.

Probably even more important is the target audience of the government propaganda. Obviously, people who use internet, and especially engage in political discussions there, are not representative of the population in general: they tend to be younger, more globally-oriented, more skeptical about the authorities, and at least in the Russian case, more liberal ideologically. Thus changes into the policies towards the online media platforms usually reveal changes in the structure of popular pro-government coalitions. In particular, the difference between Putin's policy before 2008 and after 2012 is not entirely due to the higher threat posed by online coordination between protesters in 2011-2012, but also the result of a change in the pro-Putin coalition itself. Contrary to the popular belief, in early 2000s it mainly

consisted of relatively well-to-do pro-Western urbanites, with more traditionalist and state-dependent electorate left to the unreformed Communist Party of Russia. As Putin's majority was rapidly changing, continuing Medvedev's engagement with online community became less and less politically profitable for Putin. Moreover, members of this community became the ideologically convenient and, importantly, easily identifiable targets for online restrictions and offline repressions. Optic internet cable became the sharpest line between loyal government supporters and opponents of the regime.

References

Ackland, Robert. 2013. *Web Social Science: Concepts, Data and Tools for Social Scientists in the Digital Age.* SAGE Publications. [7, 8].

Agora. 2011. *Internet Freedom in Russia 2011.* Kazan, Russia: Association of Human Rights Organizations "Agora". http://openinform.ru/fs/j_photos/openinform_353.pdf. [25].

———. 2012. *Internet Freedom in Russia 2012.* Kazan, Russia: Association of Human Rights Organizations "Agora". http://www.hro.org/node/15685. [31].

———. 2013. *Internet Freedom in Russia 2013.* Kazan, Russia: Association of Human Rights Organizations "Agora". http://eliberator.ru/files/Internet_2013.pdf. [31].

Alexanyan, Karina, Vladimir Barash, Bruce Etling, Robert Faris, Urs Gasser, John Kelly, John G. Palfrey, and Hal Roberts. 2012. "Exploring Russian Cyberspace: Digitally-Mediated Collective Action and the Networked Public Sphere." SSRN Scholarly Paper. Accessed April 30, 2014. http://papers.ssrn.com/abstract=2014998. [19, 20, 31].

Balmforth, Tom. 2014. "From The Fringes Toward Mainstream: Russian Nationalist Broadsheet Basks In Ukraine Conflict." RADIOFREEEUROPE/RadioLiberty. August 17. Accessed December 28, 2014. http://www.rferl.org/content/feature/26534846.html. [15].

Bam, Joshua. 2013. "Strategic sector legislation in Russia: critique and proposal for change." *Matters of Russian and International Law,* no. 2. [8].

Banks, Arthur S., and Kenneth A. Wilson. 2013. *Cross-National Time-Series Data Archive.* Jerusalem, Israel.: Databanks International. [17].

Barash, Vladimir, and John Kelly. 2012. "Salience vs. Commitment: Dynamics of Political Hashtags in Russian Twitter." SSRN Scholarly Paper. Accessed April 30, 2014. http://papers.ssrn.com/abstract=2034506. [4, 24].

Barykova, Olga, and Maria Zotova. 2011. "Ushlo li vremya "Moskovskih Novostey"?" BBC Russian. April 1. Accessed January 7, 2015. http://www.bbc.co.uk/russian/russia/2011/04/110331_moscow_news_scandal.shtml. [26].

Beard, Nadia. 2014. "Founder and CEO of Yandex, Arkady Volozh, resigns." Calvert Journal. August 26. Accessed December 26, 2014. http://calvertjournal.com/news/show/3035/founder-of-yandex-resigns-amid-controversy-arkady-volozh. [9].

Becker, Jonathan. 2004. "Lessons from Russia A Neo-Authoritarian Media System." *European Journal of Communication* 19 (2): 139–163. [18].

Bershidsky, Leonid. 2014. "Google's Retreat From Moscow." BLOOMBERGVIEW.com. December 12. Accessed January 12, 2015. http://www.bloombergview.com/articles/2014-12-12/googles-retreat-from-moscow. [33].

Birnbaum, Michael. 2014. "Russia′s Putin signs law extending Kremlin′s grip over media." The Washington Post. October 15. Accessed December 26, 2014. http://www.washingtonpost.com/world/europe/russias-putin-signs-law-extending-kremlins-grip-over-media/2014/10/15/6d9e8b2c-546b-11e4-809b-8cc0a295c773_story.html. [8].

Black, J. L. 2014. *The Russian Presidency of Dimitri Medvedev, 2008-2012: The Next Step Forward Or Merely a Time Out?* Routledge. [22].

Black, J. L., and Michael Johns. 2013. *Russia after 2012: From Putin to Medvedev to Putin – Continuity, Change, or Revolution?* Routledge. [22].

Boshmaf, Yazan, Ildar Muslukhov, Konstantin Beznosov, and Matei Ripeanu. 2011. "The Socialbot Network: When Bots Socialize for Fame and Money." In *Proceedings of the 27th Annual Computer Security Applications Conference,* 93–102. ACSAC '11. New York, NY, USA: ACM. Accessed June 8, 2015. doi:10.1145/2076732.2076746. [35].

Brennan, Christopher. 2014. "Putin Says CIA Created the Internet, Cites Foreign Influence at Yandex." The Moscow Times. April 24. Accessed May 17, 2014. http://www.themoscowtimes.com/news/article/putin-says-cia-created-the-internet-cites-foreign-influence-at-yandex/498903.html. [9].

Burrett, Tina. 2008. "The end of independent television? Elite conflict and the reconstructing the Russian television landscape." In *The Post-Soviet Russian Media: Conflicting Signals,* edited by Birgit Beumers, Stephen Hutchings, and Natalia Rulyova, 71–86. Routledge. [18].

———. 2010. *Television and Presidential Power in Putin's Russia.* Routledge. [16, 18].

Caute, David. 1988. *The Fellow-travellers: Intellectual Friends of Communism.* Yale University Press. [15].

Chen, Adrian. 2015. "The Agency." The New York Times. June 2. Accessed June 7, 2015. http://www.nytimes.com/2015/06/07/magazine/the-agency.html. [35].

Chen, Cheng, Kyungmin Ko, and Ji-Yong Lee. 2010. "North Korea's Internet strategy and its political implications." *The Pacific Review* 23 (5): 649–670. [11].

Chu, Zi, S. Gianvecchio, Haining Wang, and S. Jajodia. 2012. "Detecting Automation of Twitter Accounts: Are You a Human, Bot, or Cyborg?" *IEEE Transactions on Dependable and Secure Computing* 9 (6): 811–824. [35, 36].

Clover, Charles. 2011. "Internet subverts Russian TV′s message." The Financial Times. December 1. Accessed May 26, 2014. http://www.nytimes.com/2014/03/11/opinion/the-kremlins-social-media-takeover.html. [20].

Cook, Sarah. 2011. "China's growing army of paid internet commentators." Freedom At Issue Blog. October 11. Accessed May 17, 2014. http://www.freedomhouse.org/blog/china%C3%A2%C2%80%C2%99s-growing-army-paid-internet-commentators. [14].

Deibert, Ronald, John Palfrey, Rafal Rohozinski, and Jonathan Zittrain, eds. 2011. *Access Contested: Security, Identity, and Resistance in Asian Cyberspace.* Information revolution and global politics. MIT Press. [6].

Denisova, Irina, Markus Eller, and Ekaterina Zhuravskaya. 2010. "What do Russians think about transition?1." *Economics of Transition* 18 (2): 249–280. [16].

Diamond, Larry. 2010. "Liberation Technology." ¡p¿Volume 21, Number 3, July 2010¡/p¿, *Journal of Democracy* 21 (3): 69–83. [5].

Dunn, John A. 2008. "Where did it all go wrong? Russian television in the Putin era." In *The Post-Soviet Russian Media: Conflicting Signals,* edited by Birgit Beumers, Stephen Hutchings, and Natalia Rulyova, 42–55. Routledge. [18].

Economist. 2013. "The art of concealment." The Economist. April 4. Accessed May 17, 2014. http://www.economist.com/news/special-report/21574631-chinese-screening-online-material-abroad-becoming-ever-more-sophisticated. [11].

Elder, Miriam. 2012a. "Emails give insight into Kremlin youth group's priorities, means and concerns." The Guardian. February 7. Accessed May 17, 2014. http://www.theguardian.com/world/2012/feb/07/nashi-emails-insight-kremlin-groups-priorities. [15].

Elder, Miriam. 2012b. "Hacked emails allege Russian youth group Nashi paying bloggers." The Guardian. February 7. Accessed May 17, 2014. http://www.theguardian.com/world/2012/feb/07/hacked-emails-nashi-putin-bloggers. [14].

———. 2012c. "Polishing Putin: hacked emails suggest dirty tricks by Russian youth group." The Guardian. February 7. Accessed May 17, 2014. http://www.theguardian.com/world/2012/feb/07/putin-hacked-emails-russian-nashi. [14].

Enikolopov, Ruben, Vasily Korovkin, Maria Petrova, Konstantin Sonin, and Alexei Zakharov. 2013. "Field experiment estimate of electoral fraud in Russian parliamentary elections." *Proceedings of the National Academy of Sciences* 110 (2): 448–452. [27].

Enikolopov, Ruben, Maria Petrova, and Ekaterina Zhuravskaya. 2011. "Media and Political Persuasion: Evidence from Russia." *American Economic Review* 101 (7): 3253–3285. [16].

Etling, Bruce, Karina Alexanyan, John Kelly, Robert Faris, John G. Palfrey, and Urs Gasser. 2010. "Public Discourse in the Russian Blogosphere: Mapping RuNet Politics and Mobilization." SSRN Scholarly Paper. Accessed April 29, 2014. http://papers.ssrn.com/abstract=1698344. [3, 19, 23, 24].

Etling, Bruce, Hal Roberts, and Robert Faris. 2014. "Blogs as an Alternative Public Sphere: The Role of Blogs, Mainstream Media, and TV in Russia's Media Ecology." SSRN Scholarly Paper. Accessed April 29, 2014. http://papers.ssrn.com/abstract=2430786. [6, 25].

Freedom House. 2011. *Freedom on the Net 2011.* Washington, DC: Freedom House. http://www.freedomhouse.org/report/freedom-net/freedom-net-2011. [5, 25].

———. 2012. *Freedom on the Net 2012.* Washington, DC: Freedom House. http://www.freedomhouse.org/report/freedom-net/freedom-net-2012. [5, 28].

Freedom House. 2013. *Freedom on the Net 2013.* Washington, DC: Freedom House. http://www.freedomhouse.org/report/freedom-net/freedom-net-2013. [5, 28].

———. 2014. *Freedom on the Net 2014.* Washington, DC: Freedom House. https://freedomhouse.org/report/freedom-net/freedom-net-2014. [5].

Gaidar, Egor. 2003. *The Economics of Transition.* MIT Press. [16].

Gehlbach, Scott. 2010. "Reflections on Putin and the Media." *Post-Soviet Affairs* 26 (1): 77–87. [18].

Gelman, Vladimir, Dmitry Travin, and Otar Marganiya. 2014. *Reexamining Economic and Political Reforms in Russia, 1985–2000: Generations, Ideas, and Changes.* Lexington Books. [16].

Gessen, Masha. 2012. *The Man Without a Face: The Unlikely Rise of Vladimir Putin.* Penguin. [16, 18].

Giacomello, Giampiero. 2008. *National Governments and Control of the Internet: A Digital Challenge.* Routledge. [4].

Goble, Paul. 2014. "Ukrainian Events Have Deeply Split Russian Nationalists." The Interpreter. July 20. Accessed December 28, 2014. http://www.interpretermag.com/ukrainian-events-have-deeply-split-russian-nationalists/. [15].

Greenall, Robert. 2012. "LiveJournal: Russia's unlikely internet giant." BBC News. February 29. Accessed May 26, 2014. http://www.bbc.co.uk/news/magazine-17177053. [19].

Grishin, Nikolay. 2012. "Yandexed Everything." Kommersant - Trade Secret. March 12. Accessed May 17, 2014. http://www.kommersant.ru/doc/2065978. [9].

Gulyaeva, Natalia, Maria Baeva, Oxana Balayan, and Maria Sedykh. 2014. "Russia tightens foreign ownership restrictions in media." Hogan Lovells Global Media and Communications Watch. October 20. Accessed December 26, 2014. http://www.hlmediacomms.com/2014/10/20/law-restricting-foreign-ownership-in-media-business-in-russia/. [8].

Gunitsky, Seva. 2015. "Corrupting the Cyber-Commons: Social Media as a Tool of Autocratic Stability." *Perspectives on Politics* 13 (1). [6].

Hale, Henry E. 2006. "Democracy or autocracy on the march? The colored revolutions as normal dynamics of patronal presidentialism." *Communist and Post-Communist Studies,* DEMOCRATIC REVOLUTIONS IN POST-COMMUNIST STATES, 39 (3): 305–329. [23].

Howard, Philip N., and Muzammil M. Hussain. 2013. *Democracy's Fourth Wave?: Digital Media and the Arab Spring.* Oxford University Press. [11].

Human Rights Watch. 2014. "Russia: Halt Orders to Block Online Media — Human Rights Watch." March 24. Accessed May 27, 2014. http://www.hrw.org/news/2014/03/23/russia-halt-orders-block-online-media. [28].

Judah, Ben. 2013. *Fragile Empire: How Russia Fell In and Out of Love with Vladimir Putin.* Yale University Press. [16, 27].

Kelly, John, Vladimir Barash, Karina Alexanyan, Bruce Etling, Robert Faris, Urs Gasser, and John G. Palfrey. 2012. "Mapping Russian Twitter." SSRN Scholarly Paper. Accessed April 29, 2014. http://papers.ssrn.com/abstract=2028158. [24].

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): 326–343. [4, 13].

Ko, Kyungmin, Heejin Lee, and Seungkwon Jang. 2009. "The Internet dilemma and control policy: political and economic implications of the Internet in North Korea." *Korean Journal of Defense Analysis* 21 (3): 279–295. [4, 11].

Kobak, Dmitry, Sergey Shpilkin, and Maxim S. Pshenichnikov. 2012. "Statistical anomalies in 2011-2012 Russian elections revealed by 2D correlation analysis." *arXiv:1205.0741 [physics, stat].* arXiv: 1205.0741. [27].

Koltsova, Olessia. 2006. *News Media and Power in Russia.* Routledge. [16].

Komaromi, Ann. 2004. "The Material Existence of Soviet Samizdat." *Slavic Review* 63 (3): 597. [17].

Kononov, Nickolay. 2014. "The Kremlin's Social Media Takeover." The New York Times. March 10. Accessed May 17, 2014. http://www.nytimes.com/2014/03/11/opinion/the-kremlins-social-media-takeover.html. [9, 32].

Lange, Sarah. 2014. "The End of Social Media Revolutions." *The Fletcher Forum of World Affairs* 38 (1): 47–68. [6].

Lipman, Maria. 2009. "Media manipulation and political control in Russia." Chatham House. Accessed May 16, 2014. https://www.chathamhouse.org/sites/default/files/public/Research/Russia%20and%20Eurasia/300109lipman.pdf. [16, 18].

Lunden, Ingrid. 2014. "Durov, Out For Good From VK.com, Plans A Mobile Social Network Outside Russia." Techcrunch. April 22. Accessed May 17, 2014. http://techcrunch.com/2014/04/22/durov-out-for-good-from-vk-com-plans-a-mobile-social-network-outside-russia/. [9, 32].

Lunden, Ingrid. 2015. "Intel Shuts Down Russian Developer Forums To Comply With Russia's 'Blogger Law'." TECHCRUNCH. January 5. Accessed January 12, 2015. http://techcrunch.com/2015/01/05/intel-shuts-down-russian-developer-forums-to-comply-with-russias-blogger-law/. [31].

Macfarquhar, Neil. 2014. "Russia Quietly Tightens Reins on Web With 'Bloggers Law'." The New York Times. May 6. Accessed May 17, 2014. http://www.nytimes.com/2014/05/07/world/europe/russia-quietly-tightens-reins-on-web-with-bloggers-law.html. [8].

MacKinnon, Rebecca. 2011. "China's "Networked Authoritarianism"." *Journal of Democracy* 22 (2): 32–46. [4].

Meyer, Henry. 2011. "Putin Revives Gorbachev Glasnost Paper to Widen Election Appeal." Bloomberg. March 30. Accessed January 7, 2015. http://www.bloomberg.com/news/2011-03-30/putin-revives-gorbachev-glasnost-paper-to-widen-election-appeal.html. [26].

Mills, Elinor. 2009. "Twitter, Facebook attack targeted one user." CNET. August 6. Accessed May 27, 2014. http://www.cnet.com/news/twitter-facebook-attack-targeted-one-user/. [25].

Morozov, Alexander. 2011. ""Moskovskie Novosti" Space." OPENSPACE.ru. March 28. Accessed January 7, 2015. http://os.colta.ru/media/projects/18065/details/21397/. [26].

Morozov, Evgeny. 2011. "Whither Internet Control?" *Journal of Democracy* 22 (2): 62–74. [5, 9, 12].

Nabi, Zubair. 2013. "The Anatomy of Web Censorship in Pakistan." *CoRR* abs/1307.1144. [4].

Nechepurenko, Ivan. 2014. "How Nationalism Came to Dominate Russia's Political Mainstream." The Moscow Times. August 3. Accessed December 28, 2014. http://www.themoscowtimes.com/news/article/how-nationalism-came-to-dominate-russias-political-mainstream/504495.html. [15].

Noman, Helmi. 2011. "The Emergence of Open and Organized Pro-Government Cyber Attacks in the Middle East: The Case of the Syrian Electronic Army." The Information Warfare Monitor. May 30. Accessed May 17, 2014. http://www.infowar-monitor.net/2011/05/7349/. [12].

Nossik, Anton. 2013. "11 rubles and 80 kopecks per comment." Echo of Moscow. September 10. Accessed May 17, 2014. http://www.echo.msk.ru/blog/nossik/1154616-echo/. [14].

———. 2014. "I Helped Build Russia's Internet. Now Putin Wants to Destroy It." New Republic. May 15. Accessed May 26, 2014. http://www.newrepublic.com/article/117771/putins-internet-crackdown-russias-first-blogger-reacts. [19, 20, 21].

Oates, Sarah, and Tetyana Lokot. 2013. "Twilight of the Gods?: How the Internet Challenged Russian Television News Frames in the Winter Protests of 2011-12." SSRN Scholarly Paper. Accessed May 26, 2014. http://papers.ssrn.com/abstract=2286727. [21].

Odynova, Alexandra. 2009. "Yandex to Close List That Annoyed State — News." The Moscow Times. November 6. Accessed December 15, 2014. http://www.themoscowtimes.com/news/article/yandex-to-close-list-that-annoyed-state/388969.html. [14].

Pearce, Katy. 2014. "Two Can Play at that Game: Social Media Opportunities in Azerbaijan for Government and Opposition." *Demokratizatsiya: The Journal of Post-Soviet Democratization* 22 (1): 39–66. [4, 6].

Przeworski, Adam, Michael E. Alvarez, Jose Antonio Cheibub, and Fernando Limongi. 2000. *Democracy and Development: Political Institutions and Well-Being in the World, 1950-1990*. Cambridge University Press. [17].

Razumovskaya, Olga. 2011. "Russian Social Network: FSB Asked It To Block Kremlin Protesters." The Wall Street Journal. December 8. Accessed January 12, 2015. http://blogs.wsj.com/emergingeurope/2011/12/08/russian-social-network-fsb-asked-it-to-block-kremlin-protesters/. [32].

Remnick, David. 2011. "Putin's Television." The New Yorker Blogs. December 9. Accessed May 26, 2014. http://www.newyorker.com/online/blogs/newsdesk/2011/12/putins-television.html. [20].

Ries, Brian. 2014. "Founder of 'Russia's Facebook' Says Government Demanded Ukraine Protestors' Data." Mashable. April 16. Accessed January 12, 2015. http://mashable.com/2014/04/16/vkontakte-founder-fsb-euromaidan/. [32].

Roberts, Hal, and Bruce Etling. 2011. "Coordinated DDoS Attack During Russian Duma Elections." Internet & Democracy Blog. December 8. Accessed May 27, 2014. http://blogs.law.harvard.edu/idblog/2011/12/08/coordinated-ddos-attack-during-russian-duma-elections/. [25].

Rothrock, Kevin. 2015. "Meet Russia's 369 Kremlin-Registered Bloggers." Global Voices. January 8. Accessed January 12, 2015. http://globalvoicesonline.org/2015/01/08/meet-russias-369-kremlin-registered-bloggers/. [31].

Sakwa, Richard. 2014. *Putin Redux: Power and Contradiction in Contemporary Russia.* Routledge. [22, 28, 33].

Shleifer, Andrei, and Daniel Treisman. 2001. *Without a Map: Political Tactics and Economic Reform in Russia.* MIT Press. [16].

Sidorenko, Alexey. 2009. "Russia: Major Search Engine Closes Its Blog Rating." Global Voices. November 6. Accessed December 15, 2014. http://globalvoicesonline.org/2009/11/06/russia-major-search-engine-closes-its-blog-rating/. [14].

Sidorenko, Alexey. 2010. "Russia: Analysis of Hacker Attacks On Bloggers." Global Voices. June 20. Accessed December 15, 2014. http://globalvoicesonline.org/2010/06/20/russia-analysis-of-hacker-attacks-on-bloggers/. [13].

Simons, Dr Greg. 2013. *Mass Media and Modern Warfare: Reporting on the Russian War on Terrorism.* Ashgate Publishing, Ltd. [26].

Sivkova, Alena. 2014. ""We don't see much risk in blocking Twitter in Russia"." Izvestia. May 16. Accessed January 9, 2015. http://izvestia.ru/news/570863. [29].

Smyth, Regina, and Irina Soboleva. 2014. "Looking beyond the economy: Pussy Riot and the Kremlin's voting coalition." *Post-Soviet Affairs* 30 (4): 257–275. [33].

Snyder, Timothy. 2014. "Fascism, Russia, and Ukraine." The New York Review of Books. March 20. Accessed May 27, 2014. http://www.nybooks.com/articles/archives/2014/mar/20/fascism-russia-and-ukraine/. [33].

Soldatov, Andrei, and Irina Borogan. 2013. "Russia's Surveillance State." *World Policy Journal* 30 (3): 23–30. [18].

Sumlenny, Sergej. 2013. "Bad news: what does the closure of RIA Novosti mean for media in Russia?" Calvert Journal. December 12. Accessed May 17, 2014. http://calvertjournal.com/comment/show/1837/RIA-novosti-putin-russian-media-kiselyov. [8].

Tetrault-Farber, Gabrielle. 2014. "RIA Novosti Begins Cutting 1/3 of Staff." The Moscow Times. March 12. Accessed May 17, 2014. http://www.themoscowtimes.com/news/article/ria-novosti-begins-cutting-13-of-staff/495980.html. [8].

The Moscow Times. 2014. "15 Global Firms Hit by Russia's Law Limiting Foreign Ownership of Media." September 28. Accessed December 26, 2014. http://www.themoscowtimes.com/article/507968.html. [8].

Travis, Hannibal, ed. 2013. *Cyberspace Law: Censorship and Regulation of the Internet.* Routledge. [4].

Tregubova, Yelena. 2003. *The Tales of a Kremlin Digger.* Moscow: Ad Marginem. [16, 18].

West, Darrell. 2010. "President Dmitry Medvedev: Russia's Blogger-in-Chief." The Brookings Institution. April 14. Accessed January 3, 2015. http://www.brookings.edu/research/opinions/2010/04/14-medvedev-west. [23].

Womack, Helen. 2014. "Making waves: Russian radio station is last bastion of free media." *Index on Censorship* 43 (3): 39–41. [18].

Yaffa, Joshua. 2013. "Is Pavel Durov, Russia's Zuckerberg, a Kremlin Target?" Bloomberg Businessweek. August 1. Accessed May 17, 2014. http://www.businessweek.com/articles/2013-08-01/is-pavel-durov-russias-zuckerberg-a-kremlin-target. [9].

Yagodin, Dmitry. 2012. "Blog Medvedev: Aiming for Public Consent." *Europe-Asia Studies* 64 (8): 1415–1434. [23].

Yudina, Natalia. 2012. "RuNet, hate crime and soft targets: how Russia enforces its anti-extremism law." Open Democracy. October 30. Accessed May 27, 2014. http://www.opendemocracy.net/od-russia/natalia-yudina/runet-hate-crime-and-soft-targets-how-russia-enforces-its-anti-extremism-la. [27].

———. 2014. "Beware the Rise of the Russian Ultra-Right." The Moscow Times. September 11. Accessed December 28, 2014. http://www.themoscowtimes.com/opinion/article/beware-the-rise-of-the-russian-ultra-right/506876.html. [15].

Zasursky, Ivan. 2004. *Media and Power in Post-Soviet Russia.* M.E. Sharpe. [16].