

Independent Media Under Pressure: Evidence from Russia *

Tom Paskhalis	Bryn Rosenfeld	Katerina Tertytchnaya
<i>New York University</i>	<i>Cornell University</i>	<i>University College London</i>
tom.paskhalis@nyu.edu	brr59@cornell.edu	k.tertytchnaya@ucl.ac.uk

May 12, 2021

Abstract

Existing literature recognizes growing threats to press freedom around the world and documents changes in the tools used to stifle independent press. However, few studies investigate how independent media respond to state pressure. Do independent outlets comply, orienting coverage to favor regime interests? Or does repression encourage more negative coverage of the regime instead? To shed light on these questions, we investigate how the abrupt removal of independent outlet TV Rain from Russian television providers influenced its coverage. We find that shortly after TV Rain was dropped by providers, the tone of its political coverage became more positive and its similarity with state-controlled television increased. However, these effects were short-lived. Additional evidence suggests that subscription revenue contributed to the station's resilience. These findings, from the first causal test of how attacks influence independent media coverage in a nondemocracy, add to our understanding of media manipulation and authoritarian endurance.

Keywords: independent media, electoral autocracies, text analysis, Russia

Word Count: 9034

*We are grateful to Jack Blumenau, Alexandra Cirone, Arturas Rozenas, Denis Stukal, Joshua Tucker, Christopher Wratil, Yuri Zhukov, and seminar participants at Cornell, Columbia, HSE, and UCL for their feedback on earlier drafts. This paper also benefited from comments and suggestions received at the EPSA 2019 Annual Meeting, 2019 ASEEES Annual Convention and APSA 2020 Annual Meeting. We acknowledge financial support from Harriman Institute/Carnegie Corporation, University of Southern California and UCL for supporting the research for this article.

1 Introduction

Most nondemocratic regimes today allow some independent alternatives to state media. Yet illiberal incumbents' coexistence with independent outlets – those that do not depend on the state's support or sponsorship – is at best uneasy.¹ Nuanced coverage of developments at home and abroad and revelations of official malfeasance threaten to erode a regime's legitimacy and challenge the stability of authoritarian rule (Levitsky and Way 2002; Egorov, Guriev, and Sonin 2009; Petrova 2011). To stifle critical reporting, governments in autocracies and fragile democracies have in recent years deployed a diverse set of strategies (Repucci 2019). While illiberal regimes can force independent outlets to close, the costs of doing so are high (Knight and Tribin 2020; Carter and Carter 2021). As a result, incumbents often prefer more indirect strategies of control. Legal attacks as well as regulatory and financial pressure, are often used to limit outlets' access to distribution networks and advertising revenue (Repucci 2019). What impact does such pressure have on independent outlets?

We present two competing sets of hypotheses linking state pressure to political coverage on independent outlets in nondemocratic regimes. The first argues that state pressure *increases* outlets' incentives to pander to the regime. To retain advertising revenue, maintain access to state sources, or reduce the threat of future retaliation independent media may respond to pressure by providing more positive coverage of the regime and content that more closely resembles that of state outlets. The second suggests that state pressure actually *reduces* outlets' incentives to pander to the incumbent. By bolstering editors' and journalists' willingness to mobilize against a repressive state, or by increasing outlets' reliance on their liberal audience as a source of independent, subscription-based revenue, state pressure may fail to elicit more sympathetic coverage of the regime. We provide empirical tests of these hypotheses and discuss underlying mechanisms using evidence from contemporary Russia, a prominent example of an electoral authoritarian regime.

1. Following Petrova (2011, p.791), state and independent outlets primarily differ in the extent of state influence over news and editorial decisions.

Our research design exploits the abrupt removal of Russia’s most prominent independent television channel, TV Rain (*Dozhd*) from cable and satellite providers in January 2014, following the publication of a politically controversial poll. Although TV Rain’s access to advertising revenue plummeted, it continued to provide live broadcasts through its website. The empirical analysis uses original data from a corpus of 96,648 stories and utilizes a regression discontinuity in time (RDiT) design to causally examine whether and how the attack affected TV Rain’s political coverage in the weeks and months that followed. We also investigate whether TV Rain’s content became more similar to that of a state-owned channel during this period. Being able to exploit both within and between-outlet variation in the tone of coverage and its content helps us to address many threats to inference. Comparing coverage on TV Rain and state Channel 1 allows us to investigate how independent and state-run outlets respond to the same changes in the news cycle and other temporal trends. What is more, to evaluate the mechanisms associated with changes in coverage, we draw on a variety of secondary sources including press releases, interviews by the managers and editors of TV Rain, and available financial data from the station.

To preview the results, we find that shortly after TV Rain was disconnected from providers, the tone of its political coverage became more positive. In subsequent weeks, the similarity of content broadcast on TV Rain and Channel 1 also increased. Both effects, however, were short-lived. Within weeks, TV Rain’s political coverage became more negative. By summer 2014, the the content broadcast on the two outlets was again distinct. Through careful substantive discussion of potential mechanisms driving these changes, we propose that independent sources of subscriber revenue may have allowed the channel to return its more critical and distinct editorial line and address alternative explanations.

By investigating the consequences of regulatory and financial restrictions on independent outlets in a nondemocracy, this work contributes to literature on press freedom and repression. While existing scholarship has focused on forced media closures, and explored

their effect on remaining outlets' coverage in nondemocracies (Qin, Strömberg, and Wu 2018; Kronick and Marshall. 2019), less attention has been paid to how economic and legal attacks influence besieged outlets that remain in operation, an outcome that may in fact be more common around the world than outright closure (WAN-IFRA 2019). Indeed, the types of repression we consider, namely economic and regulatory attacks on independent outlets, are increasingly used by leaders in fragile democracies and electoral autocracies not only in Eastern Europe but also in Latin America, the Middle East, and Asia (WAN-IFRA 2019). To the best of our knowledge, this is the first study of how independent media in electoral autocracies adjust their coverage in response to state sanctions that stop short of shuttering an outlet. Second, our work joins scholarship that examines how various revenue sources such as advertising (Di Tella and Franceschelli 2011; Petrova 2011), or changes in media ownership, and government partisanship (Durante and Knight 2012) influence media coverage in democracies, and extends this agenda to electoral autocracies. Third, our findings have implications for studies on the factors that sustain a free press (e.g. Petrova 2011; Gehlbach and Sonin 2014). While advertisers helped to build independent media in countries such as Russia, revenue from subscribers may be key to sustaining it. Finally, if independent media can bounce back under truly autocratic institutions, resilience seems more likely in freer media environments where illiberal leaders have similarly sought to stifle the free press.

2 Independent Media in Electoral Autocracies

While illiberal incumbents dislike critical media, many allow independent TV outlets to operate. Permitting alternatives to state-run media provides several benefits. Input from independent and commercial outlets can aid incumbents in identifying poorly performing officials and disciplining bureaucratic excess (e.g. Egorov, Guriev, and Sonin 2009). Independent outlets may also help governments identify pockets of dissent, address social grievances (Huang, Boranbay-Akan, and Huang 2019), bolster the credibility of the

government's rhetoric, and provide cover for government propaganda (Stockmann and Gallagher 2011).

At the same time, independent outlets also potentially threaten the stability of authoritarian rule. During the color revolutions of Serbia, Georgia, Ukraine, and Kyrgyzstan, independent media offered opposition forces and activists a platform to express grievances, and helped them gain support for their demands (Levitsky and Way 2002, 58; Egorov, Guriev, and Sonin 2009). Alternative narratives in independent outlets can help citizens detect overt bias in state-run media (Gehlbach and Sonin 2014) and dampen support for the regime. In Russia's 1999 parliamentary elections, for example, independent TV decreased the ruling party's vote share by 9 percentage points (Enikolopov, Petrova, and Zhuravskaya 2011).

To mitigate threats to their stability, autocrats can nationalize independent media, shutdown stations, or force changes in editorial teams. Such actions have wide-ranging consequences that are both direct and indirect (Qin, Strömberg, and Wu 2018; Kronick and Marshall. 2019). Not only do they silence the outlet attacked, they may also change the behavior of other outlets that observe repression but are not directly affected (Kronick and Marshall. 2019). By observing repression, journalists learn what is permitted, and accordingly adjust their reporting (Crabtree, Fariss, and Kern 2015).

To date, however, we know less about the effect attacks by autocrats that stop short of forcing independent outlets to close. This gap in our understanding of the consequences of subtler attacks on press freedom is problematic, since illiberal leaders increasingly use a more sophisticated toolbox to control the free media (WAN-IFRA 2019).

With the costs of overt coercion rising, 21st century illiberal regimes often opt for more indirect forms of repression, delegating to private entities and using parliament-approved legislation to stifle independent voices (Guriev and Treisman 2019; Bakke, Mitchell, and Smidt 2020). Economic and legal tools such as arbitrary tax investigations, financially draining lawsuits or the abuse of regulatory and licensing practices are increasingly used against independent outlets in contemporary autocracies and fragile democracies alike.

For example, control over access to distribution networks and legislation restricting access to advertising revenue have been used to side-line independent outlets in Belarus, Russia, Hungary, and Serbia (Freedom House [2017](#); Repucci [2019](#)).

Such forms of indirect state pressure can have several simultaneous effects on independent media. Short of shuttering independent outlets, attacks may affect a station's revenue and viewership, the morale of editors and journalists, as well as access to sources and government figures. How this bundle of effects impacts reporting by an independent outlet subject to attack is the central question that we aim to address in this paper. To the best of our knowledge, this is the first study to causally explore how independent media in electoral autocracies respond to more subtle mechanisms of repression.

3 Indirect Repression and Independent Outlets

Building on studies of media capture and repression, we present two competing sets of hypotheses that link state pressure to the tone and content of independent outlets. The first hypothesis holds that state pressure increases outlets' incentives to pander to the regime. This may be due to several factors related to editors' and journalists' motivations and an outlet's resources.

First, state attacks such as legal sanctions may result in budget constraints that influence outlets' resources and dampen their ability to provide critical coverage. In nondemocratic regimes media outlets' ability to secure advertising revenue depends not only on commercial, but also political factors. Advertisers and investors are often reticent about supporting outlets perceived as hostile to the ruling regime (Levitsky and Way [2002](#)). Thus, when outlets come under state attack, revenue from advertising dries up for fear of retaliation, even when an outlet's audience does not change. The resulting fiscal pressure on independent outlets may have several implications: forcing outlets to produce less content, drop expensive programs, implement wage cuts, or even fire staff and journalists. In short, attacks that dampen outlets' access to advertising revenue

may constrain their ability to produce the kind of nuanced and investigative coverage for which they are known in nondemocratic settings.

State sanctions, or even the threat of state pressure, may thus dampen editors' willingness to challenge the state. Concerns about the consequences of critical coverage may lead editorial teams to moderate coverage either pre-emptively or in response to an attack (e.g. Gratton 2015). Drawing on evidence from Mexico, for example, Stanig (2015) shows that the adoption of repressive defamation laws dampened coverage of corruption. To avoid being prosecuted or sued, journalists may engage in self-censorship or avoid reporting certain types of information altogether.

Editors and journalist may also try to actively signal loyalty to the regime, and they can do so in a number of ways. By providing more positive coverage of the regime, reducing their coverage of the opposition, or adopting a content mix that is more similar to that of state channels—the primary source of the regime's political message—independent outlets may try to escape sanctions, or reduce their scope. In short, attacks may elicit more positive coverage by shaping editors' and journalists' motivations to self-censor.

Explanations emphasizing the effect of government sanctions on independent outlets' resources and motivations to challenge incumbents yield the following expectation: *state-pressure increases independent outlets' incentives to pander to the regime.*

Pandering to the regime, however, comes at a cost. By providing coverage which resembles that of state-controlled media, independent outlets in electoral autocracies risk losing their audience (Gehlbach and Sonin 2014). As they become more 'similar' to mainstream media, independent outlets may alienate their 'core consumers,' the viewers who value the fact that independent media provide coverage that other outlets do not.

A second argument, however, suggests that state pressure may fail to elicit the regime's desired outcome, and may even backfire. Consistent with research showing that repression mobilizes dissent (e.g. Lichbach 1987; Pan and Siegel 2020), it is possible that state sanctions may bolster editors' and journalists' willingness to mobilize against a repressive state. Moral outrage, a sense of injustice, and anger may lead journalists and editors

to radicalize, in keeping with the finding that repression can catalyze greater dissent through emotional mechanisms. Rather than reduce the willingness to confront the regime, pointed acts of repression have often been shown to increase citizens' willingness to mobilize against it (e.g. O'Brien and Deng 2015; Aytac, Schiumerini, and Stokes 2018). While many independent outlets pride themselves on giving a voice to all sides, attacks make it more difficult for journalists and editors to resist the temptation to take sides.

Changes in outlets' revenue sources may also encourage independent outlets under siege to adopt a more critical tone. To compensate for losses in advertising revenue and continue operations, outlets under strain must seek alternative sources of income. Subscriptions increasingly provide this income and allow independent media to diversify their sources of revenue (WAN-IFRA 2019). As existing research shows, it is independent outlets' ability to provide content that is distinct to state-run channels which allows them to attract audiences that seek an alternative perspective and value the importance of a free press (see e.g. Greene and Robertson 2019, 40-41). Indeed, across fragile democracies and electoral autocracies around the world, the market value of negative and critical news for independent outlets contrasts sharply with the political value of positive news for incumbents. Altogether, as the importance of subscription-based revenue increases, incentives to cater to those who are willing to pay for news that is different also increase while incentives to pander to the regime decline.

Both mechanisms, focusing on editorial teams' willingness to challenge the state, and changes in revenue sources that bolster independent outlets' incentives to pander to their core, liberal audience support the following alternative expectation: *state-pressure reduces independent outlets' incentives to pander to the regime.*

4 The Russian case

Contemporary Russia is a prominent example of electoral autocracy and shares many features in common with other illiberal regimes. While independent outlets remain active, they have faced mounting legal, regulatory, and financial pressures in recent years (Freedom House 2017).² TV Rain, one of the last remaining independent stations in Russia was established in 2010 (Greene and Robertson 2019, 40). During massive protests in 2011-2012, TV Rain gave a platform to prominent opposition leaders, but also interviewed government ministers (95-6). The protests underscored the channel's role as the country's leading independent outlet and considerably increased its viewership (150). On the eve of the events we study, TV Rain had an audience of around 10 million viewers per month and a total audience of up to 20 million, which extended well beyond Moscow.³ About one-in-four Russians had watched the channel, most on cable or satellite television (65%) (Levada-Center 2014).

For example, in late 2013, Putin dismissed the leadership of RIA Novosti, an outlet known for balanced reporting. In addition to TV Rain, two other outlets were also put under government control or taken off air in 2014. In September 2014, new legislation restricted foreigners from holding stakes of more than 20 percent in any Russian media outlet, calling the future of some of the country's leading independent publications into question (Greene and Robertson 2019, 150). Legislation enacted in July 2015 banned advertising on cable and satellite channels, depriving over 100 regional television providers of their primary source of revenue (Freedom House 2017, 560).

The attack on TV Rain was multifaceted, and involved both financial pressure and regulatory harassment. In late January 2014, cable and satellite providers—which together controlled almost 90% of the Russian market—abruptly dropped TV Rain. Their

2. For example, in late 2013, Putin dismissed the leadership of RIA Novosti, an outlet known for balanced reporting. In addition to TV Rain, two other outlets were also put under government control or taken off air in 2014. Legislation enacted in July 2015 banned advertising on cable and satellite channels, depriving over 100 regional television providers of their primary source of revenue (Freedom House 2017, 560)

3. M. Zygar, transcript of interview with Samuel Greene, June 16, 2017.

decision to drop the channel occurred on January 29, reportedly under pressure by the Kremlin (Harding 2014). Providers for their part cited commercial considerations, in particular a controversial poll about the siege of Leningrad aired by the station just three days prior. The poll, which asked whether Leningrad should have surrendered to the Nazis to avoid loss of life, was published on the eve of the 70th anniversary of the siege on January 26, and was swiftly retracted from the channel’s website and social media. In the days following the attack, the station was subject to multiple audits and investigations, including an investigation into alleged extremism.

The financial impact of the attack on TV Rain was swift. Following the scandal, advertisers withdrew their support. Advertising income, which made up about 80% of the channel’s revenue in late 2013, was cut by 80%. Within a month, viewership also declined seven-fold, from 17.4 to 2.5 million households (Levada-Center 2014). TV Rain moved online, and increasingly relied on support from subscribers. Over the coming months, TV Rain’s subscription fee, its main source of revenue, increased roughly five-fold. These trends continued in subsequent years: online subscriptions made up 60–70% of TV Rain’s revenue, compared with just 30–40% from advertising and distribution through smaller cable systems, between 2014 and 2017.⁴

5 Data

Our analysis relies on news stories from TV Rain and state-controlled Channel 1, covering the period from January 2013 to December 2014. Channel 1 provides a point of comparison for general trends in the media ecosystem at the time and state-sanctioned coverage of events.⁵

Our data, described in detail in Appendix A.1, comes from the two channels’ web

4. Source: vc.ru. 19 April 2017. “We ended up in hell and lost 80% of our viewers in a day – Interview with the general director of TV Rain, Natalia Sindeeva”.

5. We use Channel 1 rather than another independent outlet for comparison, because there was no other country-wide independent TV station in Russian in 2014. Furthermore, the attack on TV Rain could reasonably be expected to have ‘spillover’ effects that would impact coverage on other independent media (e.g. Crabtree, Fariss, and Kern 2015)

archives. In total, the corpus includes 96,648 news items on a diverse set of topics – 42,725 for TV Rain and 49,923 for Channel 1. In general, stories correspond to content presented online and broadcast on-air on the two channels. In keeping with standard practice, we treat the transcripts obtained as unbiased sources of information about the content of both outlets (Rozenas and Stukal 2019).⁶

For each individual story, we collect information about its length (in words), and its topic—whether on domestic politics, economics, other domestic news, foreign news, entertainment or sport. As discussed in detail in Appendix A.3, stories published on Channel 1 are accompanied by categorical, thematic tags, and we use those as training data to also tag TV Rain stories. Applying the same content labels used by Channel 1 to TV Rain facilitates comparison in topic dynamic across the two channels.

To evaluate the mechanisms associated with changes in coverage, we also draw on a variety of secondary sources, including press releases, published and unpublished interviews with the station’s founder and editor-in-chief, as well as revenue figures made public by the station’s leadership in 2014.

6 Empirical Analysis

6.1 Measuring news sentiment

We begin by studying changes in the tone of both channels’ political coverage. As political news, we classify stories that cite the names of prominent government figures, such as the Russian President, the Head of the Investigative Committee, the mayors of Moscow and St Petersburg and the Secretary of the Security Council. There are 16,789 such stories across both outlets (18% of the total): 8,306 (19%) on TV Rain, and 8,483 (17%) on Channel 1.⁷

6. TV Rain’s web-archive includes news bulletins and news shows, which we include in the analysis. However, as we show in Appendix C.2, findings are consistent when we restrict the sample of TV Rain to stories that appear in the ‘news’ section of the website, and to the ‘Here and Now’ news show.

7. See Appendix A.2 for more details.

We focus on this subset of news stories for a number of reasons. Stories that mention government figures are politically salient and likely to attract scrutiny, as they are carefully monitored by the Kremlin and relevant ministries. For this reason, these stories are also often carefully reviewed by editors prior to publication. We would thus expect outlets’ critical or conciliatory stance toward the government to be clearly reflected in this subset of articles.

To measure the tone of each article, we adopt a standardized dictionary approach and use a translated version of Lexicoder Sentiment Dictionary, as described in Appendix B. Dictionary analysis is commonly used for the analysis of large and complex corpora, such as news articles published over long periods of time. Lexicoder is one of the most commonly used dictionaries in political science (Young and Soroka 2012; Soroka, Stecula, and Wlezien 2015), and has been recently used for analysis in 20 different languages (Proksch et al. 2019). Furthermore, it is one of the most comprehensive sentiment analysis dictionaries with 4,567 words stems in English. Lexicoder is well-suited to our task of obtaining a continuous measure of the positivity/negativity of news stories.

We validate the sentiment indicator in a number of ways, described in detail in Appendix B. First, we sampled 25% of all articles in the corpus, translated them into English using Google Translate API, and applied the original Lexicoder dictionary. The correlation between the two sentiment scores is approximately 0.6. Next, we compared the sentiment of 1,000 randomly chosen news stories to scores assigned by human coders recruited on the crowd-sourcing platform Yandex Toloka. The correlation between the sentiment scores estimated with Lexicoder and crowd-labelers is 0.53 for articles longer than 200 words (59% of the corpus).⁸ This is roughly consistent with Proksch et al. (2019), who find correlations of about 0.6 for several other Slavic languages using similarly translated Lexicoder dictionaries. Moreover, the tone of Russian media stories, covering a wide range of topics and including direct quotes by citizens and politicians as well as interviews, reports, and studio debates, is arguably more difficult to capture than the tone of

8. The correlation falls to 0.47 for stories with fewer than 200 words, indicating that sentiment is more ambiguous in shorter texts.

the parliamentary speeches that Proksch et al. (2019) analyze. Finally, we read through the most positive and negative stories from each outlet that also mentioned government figures to ensure that these accurately reflect positive and negative coverage. Jointly, these checks suggest that the dictionary performs adequately, and that the estimated tone offers a good approximation of the stories’ sentiment.

6.2 Estimating the causal effect of Kremlin pressure

To estimate the effect of TV Rain’s abrupt removal from cable and satellite providers, we employ a regression discontinuity in time (RDiT) design (Sekhon and Titiunik 2016; Hausman and Rapson 2018). We use the day of the poll about the siege of Leningrad (26 January 2014) as the cut-off point, and days before and after the poll as the running variable. While the exact day that the channel was dropped from cable and satellite subscription packages varied across providers and TV Rain continued to broadcast on some providers for approximately a week after the poll, the risk of imminent sanctions was made abundantly clear immediately following the poll’s publication.

It is important to note that the treatment we study is a ‘bundled’ one: in the space of a few days, TV Rain faced investigations, was gradually disconnected from cable providers, and experienced a sharp decline in advertising revenue. This bundle of treatments cannot be easily disentangled. Our study estimates the overall effect of these treatments rather than isolate the effect of each individual treatment. In short, we use the phrase ‘effect of government pressure’ to describe the bundled effect of all treatments on TV Rain, and January 26th as the point of discontinuity.

Our main empirical analysis is based on OLS regressions of the following form:

$$y_{i,t} = \alpha + \tau Poll_{i,t} + \beta' X_{i,t} + \epsilon_{i,t} \tag{1}$$

where $y_{i,t}$ is the tone of individual story i published on day t , which contains mention of a government figure. Here, $Poll_{i,t}$ is the indicator function of whether the story was

published before or after the poll about the siege of Leningrad (26 January 2014), $X_{i,t}$ is the set of control variables and $\epsilon_{i,t}$ is the error term. As in a typical RDiT framework, we use time as the running variable and a treatment date (the day of the poll’s publication) as our discontinuity. However, the clustered nature of our data, with multiple stories (the unit of observation) nested in days, sets our approach apart from conventional RDiT designs (Hausman and Rapson 2018). Rather than rely on just one observation per unit of time as in the most common RDiT setup, our data also capture variation over time in the density of stories and cross-sectional variation in story characteristics. This alleviates many concerns that arise with respect to conventional RDiT designs. First, variation in the daily number⁹ (density) of stories allows us to test for sorting around the threshold (McCrary 2008). It also gives us more flexibility to analyze very narrow bandwidths around the discontinuity (starting from five days) over which the identifying assumptions are more plausibly met and to consider whether treatment effects increase or dissipate over time. Because the number of observations is usually constant over time or very small, density tests are typically not feasible in RDiT designs and researchers must rely on much wider bandwidths in order to obtain sufficient power (see Hausman and Rapson 2018, for a review).

As control variables we use the natural logarithm of the length of every news story (in words) and its topic (domestic politics, economics, other domestic news, foreign news or entertainment). As discussed in detail in Appendix A3, stories published on Channel 1 are accompanied by categorical, thematic tags, and we use those as training data to also tag TV Rain stories. Applying the same content labels used by Channel 1 to TV Rain facilitates comparison of the two channels’ similarity. To account for potentially correlated errors for stories published on the same day, we cluster standard errors by date.

The design assumes that the timing of the poll was exogenous to news reporting on TV Rain. Indeed, there is no evidence to suggest that the poll, published during TV

9. For the period of study, the daily number of stories averages around 50-70. See Appendix A1.

Rain’s program ‘*Diletanty*’ on 26 January 2014 and occasioned by the 70th anniversary of the siege’s end, was either a product of prior coverage or related to other factors that might have affected subsequent coverage. According to the channel’s general director, the poll was conceived by the program’s presenter and approved by the editor in charge as the program unfolded.¹⁰ Indeed, TV Rain’s decision to run an online poll was not unusual. From 2012 onward, approximately 2–3 polls, on average, were posted on the channel’s website each month. The timing of this particular poll was more closely tied to historical events than current developments which could explain subsequent changes in the outlet’s coverage.

Moreover, given that the station’s content had long been controversial,¹¹ it seems unlikely that its leadership could have anticipated the precise timing of the government’s attack. Though the poll served as pretext, the attack might well have occurred on another date. We thus treat its timing as-if random with respect to the station’s day-to-day political tone and content mix.

To help validate the assumption that the timing of the poll was exogenous to news reporting, we count 10 days prior to the poll’s publication and test whether results similar to those reported below are obtained. Indeed, placebo regressions, using January 16 as the cut point as well as other dates prior to the poll’s publication, show no differences in the sentiment of stories published before or after this time. Figure C.11 in Appendix C.3 reports those results. Reassuringly in terms of sorting, the density of government stories is not significantly different on either side of the threshold (see Appendix C.1). While the number of stories mentioning government figures varies over time, the results of a McCrary density test for TV Rain imply no sorting (or more precisely that we cannot reject the null of no sorting at $p = 0.23$) in the 50-day window around the block.

We also examine whether, near the cut-off, treatment and control units are similar in

10. https://tvrain.ru/teleshov/here_and_now/natalja_sindeeva_ja_hochu_izvinitstva_pered_ljudmi_kotoryh_eto_dejstvitelno_zadelo_takie_ljudi_byli_i_vnutri_telekanala-361621/

11. For example, in November 2013 it aired an investigative report chronicling corruption among high-ranking government officials.

terms of observable characteristics (Cuesta and Imai 2016). We find no evidence of other changes which might be driven by extraneous factors that could also affect government sentiment. For example, the percentage of government stories on TV Rain was remarkably similar before and after the poll’s publication, as well as before and after the channel was eventually dropped from broadcasters, on January 29 (approximately 16%). Across the period covered by our data, roughly 1-in-5 stories published on TV Rain and Channel 1 cited government figures (Appendix A.2). In addition, stories published around January 26 were of similar length (see Appendix C.4, Figure C.12) and produced by the same journalists, led by the same editorial team. Although, TV Rain was eventually forced to lay off staff, layoffs did not begin to occur until close to two months later. In sum, to the best of our knowledge, there were no other events in the days surrounding the poll’s publication, which could have driven changes in government sentiment (on TV Rain but not on Channel 1) in the absence of the poll and its aftermath.

6.3 Government tone: Main results and robustness tests

Figure 1, based on linear models, shows differences in the tone of government stories published just before and after the poll’s publication. The left-hand plot shows the results for TV Rain. The right-hand plot shows the results for Channel 1. The unit of analysis is individual stories, and the running variable is days before and after the attack. As seen in the left-hand plot in Figure 1, the actions of Russia’s cable and satellite providers had a positive and statistically significant effect on the tone of TV Rain’s political coverage. For the five-day window, the estimated coefficient indicates approximately a 0.7 percentage point increase in the prevalence of positive over negative words in political news stories. This corresponds to roughly 30% of the standard deviation in TV Rain’s tone over 2 years. This result is also robust in the twenty and thirty-five-day windows, yet dissipates shortly after.¹²

Reassuringly, we observe no similar changes when we look at the tone of coverage

12. Autocorrelation in our series is no higher than .054 at AR(4) and even lower for other lags.

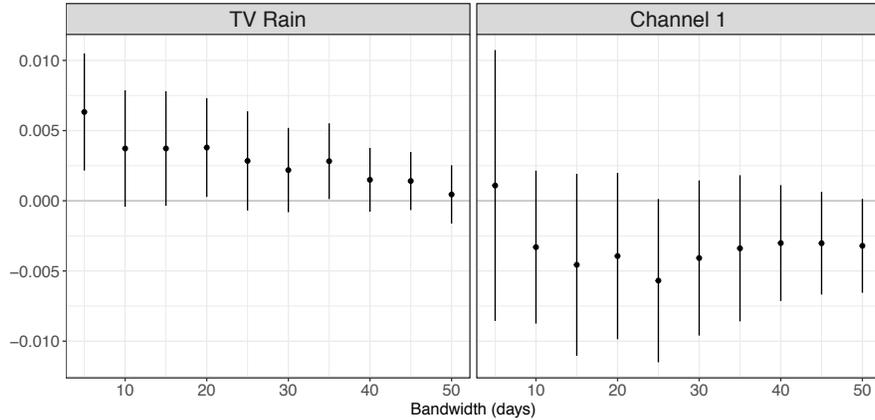


Figure 1: Effects of the poll (January 26th 2014) on government tone.

Note: The plot shows OLS regression coefficients and 95% CIs for estimates of the causal effect of the poll’s publication on government tone. The x-axis shows different bandwidths in the range between 5 and 50 days.

in state-run Channel 1, shown in the right-hand plot of Figure 1. If salient topics or other temporal trends were behind the change in tone, we would also expect Channel 1 to respond accordingly. The fact that coverage became more positive only on TV Rain increases confidence in the conclusion that regime pressure drives observed changes in tone, rather than some other factor.

As a further robustness test, we also fit local-polynomial models with automatic bin selection, using the approach described in Calonico, Cattaneo, and Rocío Titiunik (2015), and robust standard errors (Calonico, Cattaneo, and Rocio Titiunik 2014; Calonico, Cattaneo, and Rocío Titiunik 2015). The results, which are presented in Figure 2 and in tabular form in Appendix C.5, imply that TV Rain’s removal from cable and satellite packages had a significant positive effect on government tone, as in the preceding analysis. Indeed, the jump in positive coverage of the government at the moment that TV Rain came under fire can be seen clearly also in Figure 2. As elsewhere, the running variable on the x-axis is days before or after the scandal surrounding TV Rain. The y-axis displays government tone, with higher values indicating more positive coverage. The plot implies roughly a half a percentage point increase in positive political coverage.¹³

13. The somewhat smaller size of these estimates relative to those based on narrower bins in Figure

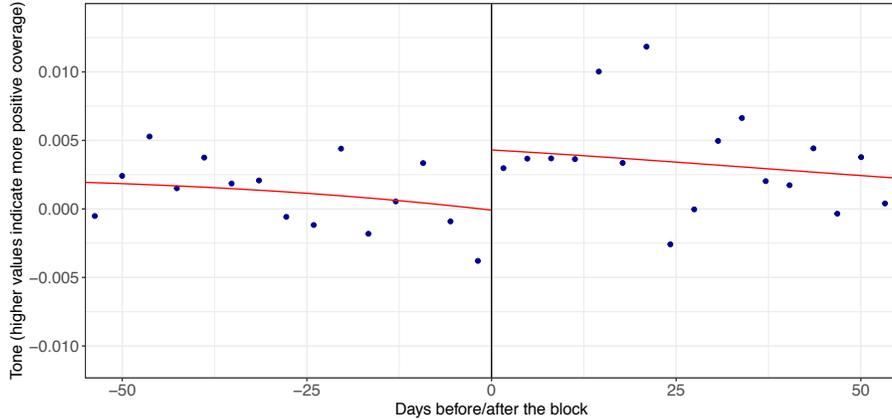


Figure 2: RD estimates using local-polynomial model.

Note: Results from a local-polynomial model with automatic bin selection and triangle kernel are shown. As elsewhere, the running variable is days before or after TV Rain published the controversial poll on January 26th.

It is worth noting that automated bin-width selection in these models results in a wider optimal bandwidth (83 days) than is shown in Figure 1. The consistency of our results across a variety of bandwidths and manual and automatic methods of bin-width selection lend further support to the main finding. However, because the automatic selection algorithms choose bin widths that are, in our judgment, too large to make credible inferences on substantive grounds, we place greater stock in models with smaller bins of less than about two months. After approximately six weeks, TV Rain began laying off staff, meaning that the journalists responsible for the stories differed pre- and post- when we focus on a larger window. In other ways as well, conditions on the ground had changed three months after TV Rain was taken off the air.¹⁴

In sum, these results show that the attack on TV Rain was initially effective and had a causal effect on the station’s coverage: journalists at first adopted a less confrontational manner of covering political news. Yet, this effect was short-lived. Approximately a month later, political coverage returned to where it was prior to the ban. What do

1 is consistent with the specification giving an estimate between the true short-run effect and long-run effect of the attack, which we argue is zero.

14. By the standards of regression discontinuity in time studies, our results rely on a very modest bandwidth. The majority of studies reviewed in Hausman and Rapson (2018) use a sample size of at least two years.

these changes mean substantively? As we discuss in Appendix D.1, the share of stories manually classified as critical of the Russian government increased from just one-in-five in the aftermath of the poll to one-in-three published in March and April, a statistically significant and substantively large difference.

The finding that political coverage on TV Rain became more positive after the ban is especially striking given that the attack effectively limited the station’s access to government sources. After TV Rain was side-lined and labeled as an opponent by the regime, it could no longer attract the government figures that had been a staple of its programming. As Mikhail Zygar, TV Rain’s editor-in-chief between 2010-2015 recalls, “for a long time, TV Rain was a station that interviewed ministers, the Prime Minister, representatives of the ruling party....that all ended in 2014 when TV Rain was marginalized”.¹⁵ In the immediate aftermath of the ban, a time when it was increasingly difficult for TV Rain to get government officials to appear on air, the tone of the station’s government coverage became more positive, and remained so for approximately a month following the attack.

6.4 Measuring Bias Based on Content

In addition to examining how government pressure shaped the tone of TV Rain’s political coverage, we also investigated how the channel’s ordeal affected its similarity with state outlet Channel 1. Now drawing on the full sample of stories published by both outlets, we analyze overtime changes in the daily similarity between TV Rain’s content and that of state-run Channel 1. To ensure that any increase in outlet similarity during early 2014 is not an artifact of winter Olympics coverage, we drop all stories on sports as well as those that mention Sochi, the Olympics etc. We assume that an outlet whose coverage is similar to that of the state channels could be seen as orienting its coverage to the regime. This approach follows existing studies that measure bias by comparing an outlet’s content and that of a source with a known political agenda (e.g. Qin, Strömberg, and Wu 2018). As a measure of content similarity, we use cosine similarity, a common measure of distance

15. M. Zygar, transcript of interview with Samuel Greene, June 16, 2017.

(see Appendix D for more details). Cosine similarity ranges between 0, meaning that stories published on any given day across the two outlets are completely different, and 1, meaning that stories on two channels are exact copies of each other. Observed differences in cosine similarity may be attributed to both changes in the topics covered across outlets and in the way those topics were covered.

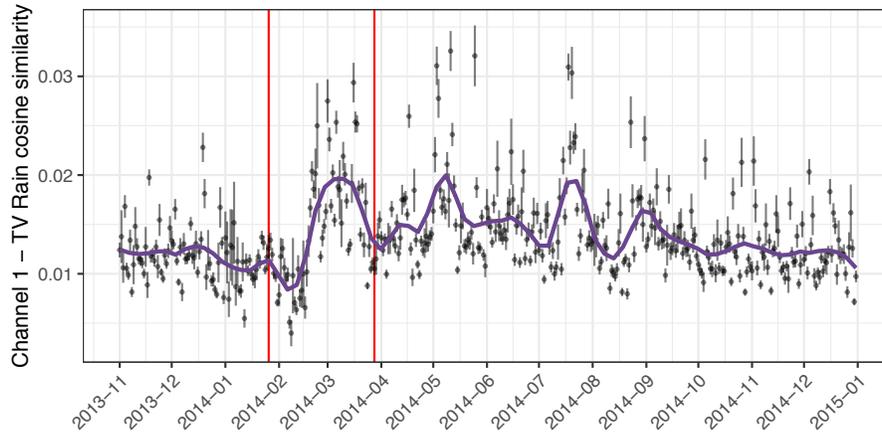


Figure 3: Trend in content similarity between TV Rain and Channel One.

Note: Estimates of daily average cosine similarity with 95% standard errors are shown. Stories on sports and mentioning Sochi are ignored. Vertical lines correspond to the date of the poll (26 January 2014) and the end of telemarathon in support of TV Rain (28 March).

Figure 3 presents the results. The similarity of coverage on TV Rain and Channel 1 begins to increase after the ban, in early February, and continues to fluctuate throughout the first half of 2014. By late July, however, content similarity between TV Rain and Channel 1 begins to decline. Altogether, this evidence of an initial rise in the similarity of content across outlets followed by a decline reinforces the preceding section’s evidence that state pressure resulted in no more than a short-lived improvement in the station’s political tone. By the second half of 2014 both types of evidence imply that TV Rain was returning to the distinct coverage it was known for before the attack.

7 Mechanisms

The evidence presented so far suggests that in the weeks following TV Rain’s ordeal, the tone of its political coverage improved. In February, content similarity with Channel 1 also began to increase. Yet, both effects were short-lived. Political coverage returned to a more negative tone by early March, while content similarity with Channel 1 declined during the summer of 2014. While we cannot entirely disentangle the bundle of factors that explain the impact of government pressure, here, we discuss several potential mechanisms that may drive changes in the tone and content of an independent station faced with state pressure.

7.1 Changes in revenue sources

As already discussed when TV Rain was dropped from TV providers, access to viewers and advertising revenue plummeted. Table 1 below shows changes in advertising revenue in the two months preceding and following the ban in January 2014.

To compensate for lost advertising revenue and continue operations, TV Rain was forced to seek alternative sources of income. One-off donations from viewers, subscriptions, as well as a live marathon in late March 2014 provided the station with much-needed revenue, generating income for more than a month of broadcasting (Ria Novosti 2014). As the second row of Table 1 shows, between December 2013 and March 2014, TV Rain’s subscription-based revenue increased approximately seven-fold¹⁶. Revenue that was independent of the state, and independent of advertisers, enabled TV Rain to continue operations. Could growing reliance on subscribers also have influenced the station’s tone of coverage and content?

Existing research shows that to maximize revenue through subscriptions and advertisement, media outlets need to produce news content that is in line with consumers’ prior beliefs (e.g. Gentzkow and Shapiro 2010). In democracies and electoral autocracies

16. See: <https://www.facebook.com/vinokurov12/posts/904727419543486>

	November 2013	December 2013	January 2014	February 2014	March 2014
Advertisement	26,454	36,387	7,834	6,560	10,030
Subscriptions	1,982	6,316	11,244	15,012	40,210

Table 1: TV Rain Sources of Revenue, November 2013-May 2014 (thousands of rubles)
Note: In January 2014, there was a sharp decline in advertising revenue that coincided with TV Rain being dropped by broadcasters. In late March, there was a large boost in subscriber revenue.

alike, stations need to cater to a target demographic’s expectations. Survey evidence suggests that unlike others, TV Rain’s core viewers—and those most able to support the channel financially through subscription revenue—believed that it was possible, and even patriotic, to criticize Russia’s leaders (Levada-Center 2014). Among TV Rain’s wealthiest viewers and most likely financial supporters, this view was held by an overwhelming majority (86%).¹⁷ Altogether then, the need to sustain the support of an increasingly narrow group of their liberally-minded audience may have well encouraged editors to pander to subscribers who had political and ideological preferences that were distinct from those of the ruling regime.

Indeed, TV Rain’s return to more negative political coverage in February coincided with rising reliance on subscription revenue, shown in Table 1. As shown in Figure 4, the tone of coverage continued to remain negative in the months that followed. Likewise, TV Rain’s similarity with state-owned Channel 1 was fleeting, returning to pre-attack levels in the summer of 2014. Again, this shift in coverage temporally coincided with TV Rain’s decision to increase its subscription fee almost five-fold.¹⁸ Subscribers who were willing to pay in order to maintain access to news coverage likely differed from those who opted not to renew their subscription in meaningful ways: they were likely to be wealthier and to place greater value on accessing distinctive content (Levada-Center 2014). Aware of these changes, TV Rain may have felt that it was more important to distinguish itself

17. Individual, story-level viewership data collected for the project also suggest that throughout 2014, stories providing more critical political coverage attracted a broader audience than others. The (Pearson) correlation between tone and web page views is -0.03 ($t = -3.4$, $df = 18318$, $p < 0.001$). As shown elsewhere, negative news sells.

18. The fee increased from 1,000 to 4,800 roubles. Source: vc.ru. 19 April 2017. “We ended up in hell and lost 80% of our viewers in a day – Interview with the general director of TV Rain, Natalia Sindeeva”.

from other outlets. In sum, we cautiously interpret temporal changes in the tone and content of coverage on TV Rain as suggesting that the infusion of subscriber revenue added resilience to the station’s business model, and helped it regain its independent voice.

7.2 Changes in the levels of threat facing independent outlets

According to another argument, changes in the tone and content of independent outlets may reflect changes in the levels of threat, or pressure facing targeted stations, and similar outlets more broadly. In the case of TV Rain, one might argue that the initial improvement in the tone of political coverage, as well as the increase in content similarity with Channel 1 reflected editors’ willingness to appease the authorities and prevent additional retaliation. Evidence from interviews with the station’s leadership supports this interpretation. It is clear that in the early days following the poll, TV Rain still held out hope for an agreement with broadcasters. It is also clear that editors believed the station’s survival hinged on appeasing the Kremlin (Aksenov 2014).

Although the possibility of concessions may have encouraged TV Rain’s editorial team to initially show restraint in the early days following the attack, it is unlikely that the critical editorial line that re-emerged in February, or the decline in content similarity with Channel 1 in the summer of 2014, were due to the station’s belief that the threat had passed. Indeed, editors in non-democracies rarely have a clear sense of when it is ‘safe’ to resume normal operations. In Russia, 2014 was a year of growing—not decelerating—attacks on independent outlets. In April 2014, for example, when TV Rain’s online editorial line was becoming more negative, the Russian regulator was cracking down on opposition blogs. Around the same time, Russia’s most prominent independent regional broadcaster also went off the air. Yet the tone of TV Rain’s political coverage only became more negative.

Indeed, given the broader climate of intimidation for independent media throughout this period, we might wonder whether attacks on other outlets rather than the poll per se

caused TV Rain’s editorial line to shift. The dismantling of another station in December 2013 may have had a chilling effect on TV Rain, but it is unclear how those events could explain the jump that occurred nearly a month later, between late January and early February 2014. In short, given the timing of these attacks, demonstration effects alone cannot explain the patterns we observe in our data.

7.3 Editors’ response

At the same time, attacks on independent outlets may have a direct effect on journalists’ and editors’ motivation to adopt a more critical stance towards incumbents through their emotional impact. Interviews make clear that the attack provoked strong emotions among the editors and journalists impacted, which made it difficult to refrain from criticizing the government. According to the station’s editor-in-chief at the time, “[the attack on TV Rain] was a catastrophe; every journalist felt as if they were at war, as if they were taking fire, as if an atom bomb had been dropped on them. This feeling of hurt made it impossible to be...unbiased” ¹⁹. Yet rather than attack the government, TV Rain initially offered more conciliatory coverage. Crucially, TV Rain’s coverage did not return to its more critical valence and distinct coverage until some time later, once it became clear that an agreement with broadcasters was not forthcoming and with the increase in subscriber revenue.

7.4 Changes in staffing

Resource-based explanations of media capture suggest that changes in outlets’ tone and content may reflect changes in outlets’ capacity to produce critical coverage. Tightening budget constraints may force outlets to implement wage cuts, drop expensive programs, and even fire staff and journalists. Another possible alternative, therefore, is that changes in TV Rain’s news content and framing reflect changes in who was producing the stories. Indeed, by early spring 2014, loss of revenue forced TV Rain to make pay cuts, cut

19. M. Zygar, transcript of interview with Samuel Greene, June 16, 2017.

weekend shifts, and fire almost half of its staff.²⁰ Staff departures were predictably nonrandom. Journalists most likely to be hired elsewhere left the station first. If these journalists were more likely to invest effort into producing original, investigative content, their departure could explain the growing similarity in content between TV Rain and Channel 1. This explanation has at least two shortcomings. First, TV Rain did not initiate layoffs until weeks after the point at which its content became more similar to that of Channel 1 in early February. Second, this explanation cannot account for the return to more critical coverage and distinct content that occurred on the channel over time. There is no evidence that TV Rain was able to rehire the journalists it had let go in the spring. Thus, changes in who was producing the stories are not consistent with broader patterns in the data and seem unlikely to explain the changes in tone and content that we observe.²¹

8 Alternative Explanations: Case-specific factors

The changes in revenue, levels of regime repression, and independent outlets' motivation to challenge the state just discussed are common factors shaping targeted outlets' response to state pressure in any context. Yet, several factors specific to the Russian context at the time of the attack may have also shaped TV Rain's response. We consider several below.

To begin with, the approval ratings of Russia's leaders spiked as events unfolded in Ukraine. Could changes in the tone of TV Rain's coverage follow a pattern of pandering to an already popular government? The short answer is no. As the dynamic of presidential approval in Figure 4 shows, coverage on TV Rain became more positive before

20. Vc.ru. April 19, 2017. "We ended up in hell and lost 80% of our viewers in a day – Interview with the general director of TV Rain, Natalia Sindeeva"

21. Finally, could the move from cable and satellite broadcasters to an online-only format explain the changes we witness? It is also worth asking whether a change in platform could drive these changes in content, as different platforms have different norms and audiences. In brief, if the change in platform were responsible, it seems likely that TV Rain would have become *less rather than more* similar to other broadcast channels in the first half of 2014. Specifically, it is not clear why TV Rain would suddenly become more like Russia's quintessential broadcast TV channel, Channel 1, by moving online.

Putin’s approval rating began to rise and well before the annexation of Crimea. Approval remained high—and continued to rise gradually—over the summer of 2014 while TV Rain’s tone of coverage did not improve, and while the station began to regain its distinct coverage. Although TV Rain’s positive coverage of the Kremlin at the start of 2014 may have contributed to the Crimea consensus, shaping Russian opinion in ways that were favorable to the regime, changes in the tone of its coverage cannot be explained by incentives to pander to a popular leader: witness the divergence between TV Rain’s tone and popular opinion of Putin during especially the second and third quarter of 2014.

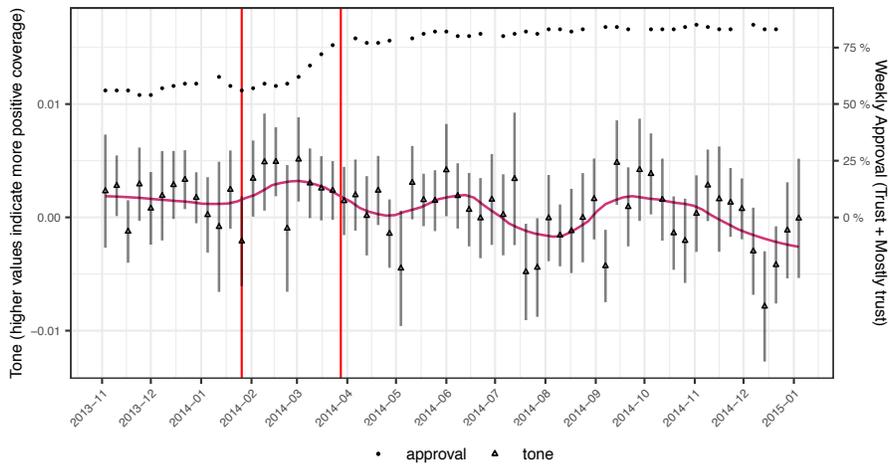


Figure 4: Tone of Government Coverage on TV Rain and Presidential Approval
Note: Estimates of weekly average ton with 95% standard errors are shown. Vertical red lines correspond to the date of TV Rain’s poll on the siege of Leningrad (26 January 2014) and the date of the end of telemarathon on TV Rain (28 March 2014). Approval ratings come from the weekly ‘FOMnibus’ surveys

Another possibility is that the changes we have documented in TV Rain’s coverage are best explained by the impact of extraordinary events and would have occurred regardless of state sanctions, or shocks to the station’s revenue. We might be concerned that TV Rain would have become more positive in its coverage of the government as a result of the Sochi Olympics or the annexation of Crimea even if it had not been dropped by providers. This too seems unlikely given the evidence. The changes on TV Rain began before coverage of the Sochi Olympics started and before the annexation of Crimea. In

fact, the uptick in positive coverage of the government began at the end of January, precisely when it was dropped by cable and satellite providers. Empirically, changes in the tone of political coverage on Channel 1 remain robust when stories related to the Olympics, or sports more broadly, are excluded from the analysis. To account for concerns regarding the impact of Sochi coverage on the similarity of coverage across TV Rain and Channel 1, we excluded stories on the winter Olympics and sports from that part of the analysis.

9 Conclusion

To the best of our knowledge, this is the first study to explore how independent media respond to government attacks that stop short of forcing an outlet's closure, and to document the efficacy of such as strategy over time. The empirical analysis relied on original data on the tone and content mix of coverage from a corpus of approximately 96,648 Russian-language news stories. We began by examining whether government pressure was effective in shaping coverage on the independent outlet TV Rain. Evidence based on an RD design suggests that when TV Rain was disconnected from television providers, the tone of its political coverage became more favorable to the government. We then analyzed over-time changes in the content of TV Rain's broadcasts relative to the content of state-controlled Channel 1. Our findings show a rise in the similarity of the content broadcast on TV Rain and state TV following the attack. Yet this 'convergence' effect was short-lived. Political coverage on TV Rain again became more negative and the similarity of the content broadcast by the two outlets declined in subsequent months. Altogether, the findings suggest that government pressure, even when exercised at arm's length, can have at least a temporary chilling effect on reporting in independent outlets.

We acknowledge that TV Rain's ordeal represents a bundle of several treatments, and that several factors drove the channel's response to state pressure. Through careful substantive discussion of the case, we made an argument for the importance of subscriber

revenue for understanding the impact of government pressure. Indeed, coverage became more critical of the government and less similar to the state channels as revenue from viewers increased. Our findings thus also highlight the dilemma autocrats' face in managing independent media. The very attacks that autocrats use to intimidate independent media can lead stations to become more reliant on viewer support, thus making their efficacy short-lived.

While evidence from a single country does not provide a general test, our results suggest, first, that government sanctions can stifle criticism from independent outlets and, second, that access to independent revenue may limit the effectiveness of such tactics. Importantly, TV Rain's response to state pressure and increased reliance on subscriber revenue is indicative of the strategies adopted by besieged outlets elsewhere. Existing research finds that in an effort to resist government pressure, outlets in Latin America are switching to a new business model, one that involves memberships and subscriptions (WAN-IFRA 2019). Independent outlets facing mounting pressures in Indonesia and Mexico have also shifted towards a subscription-based model for online services (18). Malaysiakini, a popular independent Malaysian news site, in existence since the 1990s and a frequent target of government pressure, has relied largely on subscriptions since 2002²². With government attacks undermining its ability to secure advertising revenue, Malaysiakini, like TV Rain, turned to its dedicated consumers. Similarly, Mada Masr an independent digital news site repeatedly harassed by the Egyptian government, turned in 2017 to its readers. According to commentators, "[Mada Masr's] decision to launch a membership program was more than just an attempt to develop revenue streams and sustainability; it was an attempt to maintain editorial independence and build a community."²³ SunStar in the Philippines and Hiber (7iber) in Jordan are also cases that exemplify how access to subscription revenue may help independent outlets maintain an

22. According to Malaysiakini's CEO, having struggled to secure advertising revenue due to political pressures, the outlet had to "either go subscription, or basically close down". See: <https://gijn.org/2018/07/04/after-20-years-fighting-for-democracy-whats-next-for-malaysiakini/>

23. <https://www.niemanlab.org/2017/06/mada-masr-egypts-poster-child-for-digital-news-begins-its-for>

independent voice when faced with government pressure²⁴. Recognising the importance of diversified revenue, the world association of news publishers has recently emphasized that adapting revenue models to ensure access to independent income could help safeguard press freedom in repressive contexts (WAN-IFRA 2019).

Even if the attack ultimately failed to extract from TV Rain more favorable coverage of the regime, was it successful in achieving broader aims? It is true that the attack left TV Rain with a much smaller audience—and one that was more urban and likely more liberal—than it had previously. However, TV Rain’s coverage remains relevant. When unauthorized opposition protests again broke out in Moscow in the summer of 2019 in advance of city-wide elections, TV Rain not only covered them, the station temporarily lifted its paywall to allow broad access. The move was important enough to provoke further Kremlin retaliation (Lebedeva 2019). In short, the Kremlin’s move against TV Rain ultimately succeeded neither at halting the station’s critical coverage nor at making it too obscure to be relevant.

This research opens several promising lines of inquiry. An important question concerns why regimes opt for one form of pressure over another. Under what conditions do governments choose to apply regulatory pressure, and when do they simply force an outlet to close? This question, which implicates both the international and domestic costs of stifling free speech, is beyond the scope of this study but a worthy avenue for future research. Future work could also corroborate or further contextualize findings on the importance of subscription revenue by comparing coverage across different types of independent outlets with varying degrees of dependence on viewer-based revenue in other settings. Such studies might fruitfully examine whether outlets that rely heavily on subscribers are more politically independent and critical than those that rely predominantly on advertisers, and also the related question of whether subscription-based models actually help to deter government attacks.

24. <https://www.mediasupport.org/what-does-it-mean-to-be-a-progressive-media-outlet/>

References

- Aksenov, P. 2014. *TV channel “Rain” is ready to ask for salvation from Putin*. 4 February. BBC.
- Aytaç, S. Erdem, Luis Schiumerini, and Susan Stokes. 2018. “Why Do People Join Backlash Protests? Lessons from Turkey.” *Journal of Conflict Resolution* 62 (6): 1205–1228. doi:[10.1177/0022002716686828](https://doi.org/10.1177/0022002716686828).
- Bakke, Kristin M, Neil J Mitchell, and Hannah M Smidt. 2020. “When States Crack Down on Human Rights Defenders.” *International Studies Quarterly* 64 (1): 85–96. doi:[10.1093/isq/sqz088](https://doi.org/10.1093/isq/sqz088).
- Calonico, Sebastian, Matias D. Cattaneo, and Rocio Titiunik. 2014. “Robust Nonparametric Confidence Intervals for Regression-Discontinuity Designs.” *Econometrica* 82 (6): 2295–2326. doi:<https://doi.org/10.3982/ECTA11757>.
- Calonico, Sebastian, Matias D. Cattaneo, and Rocío Titiunik. 2015. “Optimal Data-Driven Regression Discontinuity Plots.” *Journal of the American Statistical Association* 110 (512): 1753–1769. doi:[10.1080/01621459.2015.1017578](https://doi.org/10.1080/01621459.2015.1017578).
- Carter, Erin Baggott, and Brett L. Carter. 2021. *Propaganda in Autocracies*. Book manuscript.
- Crabtree, Charles, Christopher J Fariss, and Holger L Kern. 2015. *Truth Replaced by Silence: A Field Experiment on Private Censorship in Russia*. Working Paper. doi:[10.2139/ssrn.2708274](https://doi.org/10.2139/ssrn.2708274).
- Cuesta, Brandon de la, and Kosuke Imai. 2016. “Misunderstandings About the Regression Discontinuity Design in the Study of Close Elections.” *Annual Review of Political Science* 19 (1): 375–396. doi:[10.1146/annurev-polisci-032015-010115](https://doi.org/10.1146/annurev-polisci-032015-010115).
- Di Tella, Rafael, and Ignacio Franceschelli. 2011. “Government Advertising and Media Coverage of Corruption Scandals.” *American Economic Journal: Applied Economics* 3 (4): 119–51. doi:[10.1257/app.3.4.119](https://doi.org/10.1257/app.3.4.119).

- Durante, Ruben, and Brian Knight. 2012. "Partisan Control, Media Bias, and Viewer Responses: Evidence from Berlusconi's Italy." *Journal of the European Economic Association* 10 (3): 451–481. doi:[10.1111/j.1542-4774.2011.01060.x](https://doi.org/10.1111/j.1542-4774.2011.01060.x).
- Egorov, Georgy, Sergei Guriev, and Konstantin Sonin. 2009. "Why Resource-poor Dictators Allow Freer Media: A Theory and Evidence from Panel Data." *American Political Science Review* 103 (4): 645–668. doi:[10.1017/s0003055409990219](https://doi.org/10.1017/s0003055409990219).
- Enikolopov, Ruben, Maria Petrova, and Ekaterina Zhuravskaya. 2011. "Media and Political Persuasion: Evidence from Russia." *American Economic Review* 101 (7): 3253–85. doi:[10.1257/aer.101.7.3253](https://doi.org/10.1257/aer.101.7.3253).
- Freedom House. 2017. *Freedom on the Net 2017: Manipulating Social Media to Undermine Democracy*. Freedom House Annual Report. Freedom House.
- Gehlbach, Scott, and Konstantin Sonin. 2014. "Government control of the media." *Journal of Public Economics* 118:163–171. doi:<https://doi.org/10.1016/j.jpubeco.2014.06.004>.
- Gentzkow, Matthew, and Jesse M. Shapiro. 2010. "What Drives Media Slant? Evidence From U.S. Daily Newspapers." *Econometrica* 78 (1): 35–71. doi:[10.3982/ECTA7195](https://doi.org/10.3982/ECTA7195).
- Gratton, Gabriele. 2015. "The sound of silence: Political accountability and libel law." *European Journal of Political Economy* 37:266–279. doi:<https://doi.org/10.1016/j.ejpoleco.2014.09.007>.
- Greene, Samuel A., and Graeme B. Robertson. 2019. *Putin V. the People : The Perilous Politics of a Divided Russia*. New Haven: Yale University Press.
- Guriev, Sergei, and Daniel Treisman. 2019. "Informational Autocrats." *Journal of Economic Perspectives* 33 (4): 100–127. doi:[10.1257/jep.33.4.100](https://doi.org/10.1257/jep.33.4.100).
- Harding, Luke. 2014. *Russian news channel TV Rain may close after main carrier pulls plug*. 4 February. The Guardian.

- Hausman, Catherine, and David S Rapson. 2018. "Regression Discontinuity in Time: Considerations for Empirical Applications." *Annual Review of Resource Economics* 10:533–552.
- Huang, Haifeng, Serra Boranbay-Akan, and Ling Huang. 2019. "Media, Protest Diffusion, and Authoritarian Resilience." *Political Science Research and Methods* 7 (1): 23–42. doi:[10.1017/psrm.2016.25](https://doi.org/10.1017/psrm.2016.25).
- Knight, Brian, and Ana Tribin. 2020. *Opposition Media, State Censorship, and Political Accountability: Evidence from Chavez's Venezuela*. Working Paper.
- Kronick, Dorothy, and John Marshall. 2019. *Opposition media and political accountability: Evidence from Venezuela*. Working Paper.
- Lebedeva, Valeriya. 2019. *The General Director of "Dozhd" linked the channel's tax audits to the coverage of protests*. 1 August. Kommersant.
- Levada-Center. 2014. *Omnibus Survey*. February 21-25. Levada-Center.
- Levitsky, Steven, and Lucan A Way. 2002. "Elections Without Democracy: The Rise of Competitive Authoritarianism." *Journal of Democracy* 13 (2): 51–65. doi:[10.1353/jod.2002.0026](https://doi.org/10.1353/jod.2002.0026).
- Lichbach, Mark Irving. 1987. "Deterrence or Escalation?: The Puzzle of Aggregate Studies of Repression and Dissent." *Journal of Conflict Resolution* 31 (2): 266–297. doi:[10.1177/0022002787031002003](https://doi.org/10.1177/0022002787031002003).
- McCrary, Justin. 2008. "Manipulation of the running variable in the regression discontinuity design: A density test." The regression discontinuity design: Theory and applications, *Journal of Econometrics* 142 (2): 698–714. doi:<https://doi.org/10.1016/j.jeconom.2007.05.005>.

- O'Brien, Kevin J., and Yanhua Deng. 2015. "Repression Backfires: tactical radicalization and protest spectacle in rural China." *Journal of Contemporary China* 24 (93): 457–470. doi:[10.1080/10670564.2014.953849](https://doi.org/10.1080/10670564.2014.953849).
- Pan, Jennifer, and Alexandra A Siegel. 2020. "How Saudi Crackdowns Fail to Silence Online Dissent." *American Political Science Review* 114 (1): 109–125. doi:[10.1017/S0003055419000650](https://doi.org/10.1017/S0003055419000650).
- Petrova, Maria. 2011. "Newspapers and Parties: How Advertising Revenues Created an Independent Press." *American Political Science Review* 105 (4): 790–808. doi:[10.1017/S0003055411000360](https://doi.org/10.1017/S0003055411000360).
- Proksch, Sven-Oliver, Will Lowe, Jens Wäckerle, and Stuart Soroka. 2019. "Multilingual Sentiment Analysis: A New Approach to Measuring Conflict in Legislative Speeches." *Legislative Studies Quarterly* 44 (1): 97–131.
- Qin, Bei, David Strömberg, and Yanhui Wu. 2018. "Media Bias in China." *American Economic Review* 108 (9): 2442–2476. doi:[10.1257/aer.20170947](https://doi.org/10.1257/aer.20170947).
- Repucci, Sarah. 2019. "Freedom and the Media 2019: A Downward Spiral." In *Freedom and the Media 2019*. Washington, DC: Freedom House.
- Ria Novosti. 2014. "TV Rain" in five days of the marathon raised money for more than a month of broadcasting. 28 March.
- Rozenas, Arturas, and Denis Stukal. 2019. "How Autocrats Manipulate Economic News: Evidence from Russia's State-Controlled Television." *The Journal of Politics* 81 (3): 982–996. doi:[10.1086/703208](https://doi.org/10.1086/703208).
- Sekhon, Jasjeet S, and Rocio Titiunik. 2016. "Understanding regression discontinuity designs as observational studies." *Observational Studies* 2:173–181.

- Soroka, Stuart N, Dominik A Stecula, and Christopher Wlezien. 2015. "It's (Change in) the (Future) Economy, Stupid: Economic Indicators, the Media, and Public Opinion." *American Journal of Political Science* 59 (2): 457–474. doi:<https://doi.org/10.1111/ajps.12145>.
- Stanig, Piero. 2015. "Regulation of Speech and Media Coverage of Corruption: An Empirical Analysis of the Mexican Press." *American Journal of Political Science* 59 (1): 175–193. doi:<https://doi.org/10.1111/ajps.12110>.
- Stockmann, Daniela, and Mary E. Gallagher. 2011. "Remote Control: How the Media Sustain Authoritarian Rule in China." *Comparative Political Studies* 44 (4): 436–467. doi:[10.1177/0010414010394773](https://doi.org/10.1177/0010414010394773).
- WAN-IFRA. 2019. *Unlocking Journalism Resilience: Adapting a Digital Business Model To Promote Press Freedom*. Report.
- Young, Lori, and Stuart Soroka. 2012. "Affective News: The Automated Coding of Sentiment in Political Texts." *Political Communication* 29 (2): 205–231. doi:[10.1080/10584609.2012.671234](https://doi.org/10.1080/10584609.2012.671234).

Appendix

Independent Media Under Pressure: Evidence from Russia

April 23, 2021

Contents

A Descriptive Analysis	2
A.1 Corpus description	2
A.2 Stories that mention government figures	4
A.3 Content classification	5
B Sentiment Analysis	8
B.1 Translation check	9
B.2 Manual labelling: Crowd workers	9
B.3 Most positive and negative stories around the discontinuity	11
C RD Robustness Tests	14
C.1 Density of government stories around the discontinuity	14
C.2 News-only results	15
C.3 Placebo cutoff	16
C.4 Placebo outcome	17
C.5 RD results using local-polynomial models	17
D Cosine Similarity	18
E Additional Evidence	19
E.1 Manually coded critical coverage of government figures on TV Rain	19

A Descriptive Analysis

A.1 Corpus description

We scraped news items, which cover the period between January 1 2013 and December 31 2014 from the web archives of TV Rain¹ and Channel 1². In total, our corpus includes 92,648 news stories - 42,725 for TV Rain and 49,923 for Channel 1, plotted in [Figure A.1](#). News stories correspond to content presented online and broadcasted on-air on the two channels and through the cable networks, from which TV Rain was banned after January 2014. [Figure A.2](#) further shows over-time changes in the length of news stories across the two channels.

As shown below, from March 2013 to March 2014, and despite being a smaller channel, TV Rain produced as many stories as state-controlled Channel 1 did. Starting in March 2014, the number of stories TV Rain produced declined. From May to December 2014, the channel's online archive includes approximately 50 stories a day.

To make our analysis of independent media (TV Rain) more comprehensive, and due to less clearcut differentiation between news bulletins and news shows, the analysis presented in the manuscript includes all stories collected from TV Rain's archive (coming from news programmes and other news shows). The sample for Channel 1 is restricted to news stories only. However, as we show below, the main empirical results are consistent with a more narrow definition of news on TV Rain, where include only those stories that appear in *Novosti Saita* (Website News) news section³ and on *Zdes' i Seichas* (Here and Now) news show⁴.

¹<https://tvrain.ru/archive/>

²<https://www.1tv.ru/search>

³https://tvrain.ru/teleshov/novosti_sajta/

⁴https://tvrain.ru/teleshov/here_and_now/

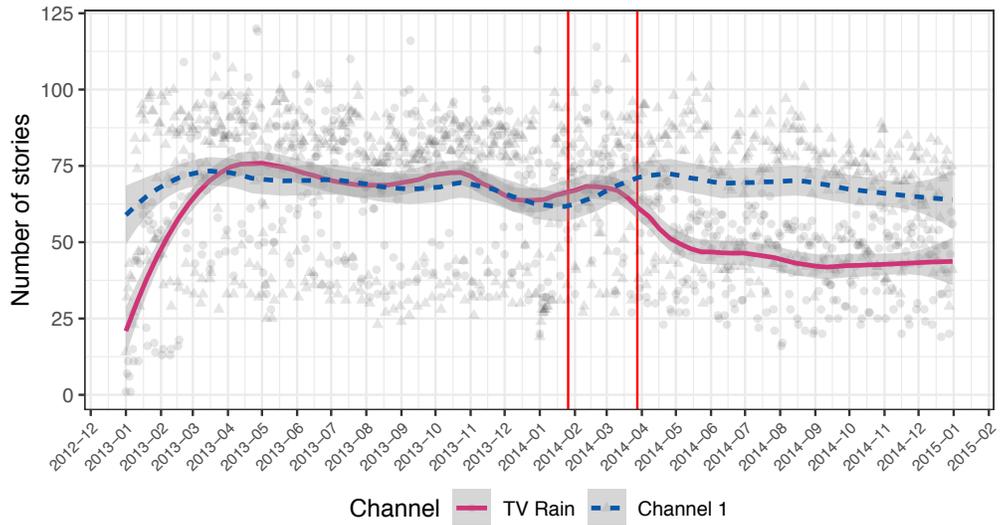


Figure A.1: **Number of stories in the corpus (2013-2014)**. Vertical red lines correspond to the date of TV Rain’s poll on the siege of Leningrad (26 January 2014) and the date of the end of telemarathon on TV Rain (28 March 2014).

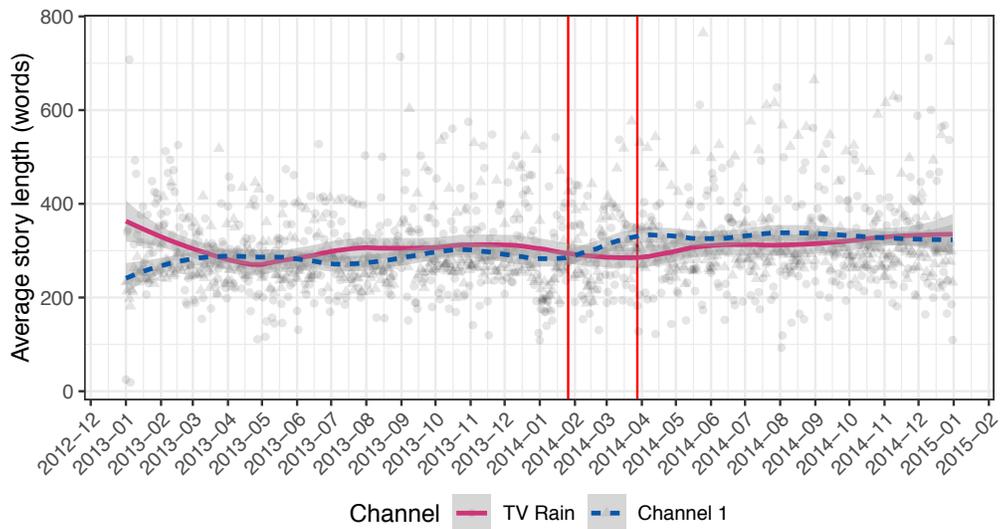


Figure A.2: **Average daily story length (2013-2014)**. Vertical red lines correspond to the date of TV Rain’s poll on the siege of Leningrad (26 January 2014) and the date of the end of telemarathon on TV Rain (28 March 2014).

A.2 Stories that mention government figures

Relevant stories are those that cite prominent members of the Russian government and the Federation Council, the Russian President, the Head of the Investigative Committee, the mayors of Moscow and St Petersburg and the Secretary of the Security Council. Table A.1 lists all the search patterns used to identify stories that contain mentions of government figures in the title or main text. There are 16,789 such stories across outlets, 8,306 (19.4% of stories) in TV Rain, and 8,483 (17%) in Channel 1. Figure A.3 shows the share of stories mentioning government figures, as a percentage of all stories in the sample.

Table A.1: Glob patterns used for identifying government figures.

бастрыкин*	матвиенко	полтавченко
голиков*	медведев*	путин*
голодец	мединск*	силуанов*
козак	набиуллин*	собянин*
колокольцев*	патрушев*	шойгу
лавров*	песков*	

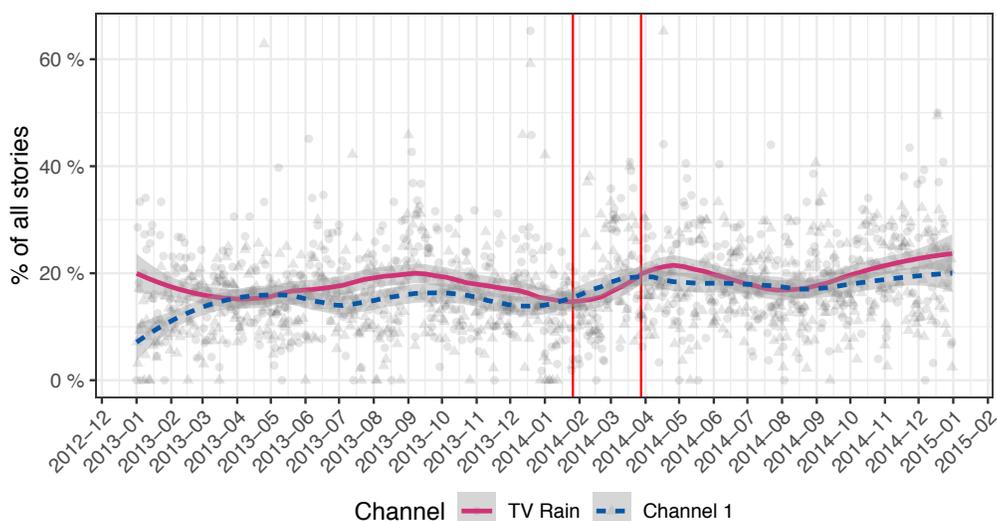


Figure A.3: **Stories mentioning government figures (2013-2014)**. Vertical red lines correspond to the date of TV Rain’s poll on the siege of Leningrad (26 January 2014) and the date of the end of telemarathon on TV Rain (28 March 2014).

A.3 Content classification

The corpus we collected covers a diverse set of topics on domestic and foreign politics. Individual articles published on the online archive of Channel 1 are accompanied by categorical, thematic tags. We collapse the 13 available categories (Politics, Economy, Society, Foreign News, Crime, Technology, Health, Culture, Sports, (program) ‘Notwithstanding’, Weather, Moscow, and the anniversary of the ‘Time’ Program) into the following topics: (i) domestic politics; (ii) the domestic economy; (iii) other domestic news; (iv) foreign news, (v) entertainment, and (vi) sport.

We use the 46,572 pre-tagged stories from Channel 1 as training data to build a supervised machine learning model for labelling 3,351 untagged stories from the same channel as well as all stories from TV Rain. To do so, we use L1-regularized (lasso) multinomial regression with stemmed unigrams and bigrams as predictors. The optimal lambda parameter was chosen through cross-validation to minimise multinomial deviance. Applying the same content labels used by Channel 1 to TV Rain facilitates comparison of the topical similarity of the two channels. The figures that follow show the daily dynamic of topical coverage for all stories in the corpus ([Figure A.4](#)), and for the subset of stories in which government figures are mentioned ([Figure A.5](#)).

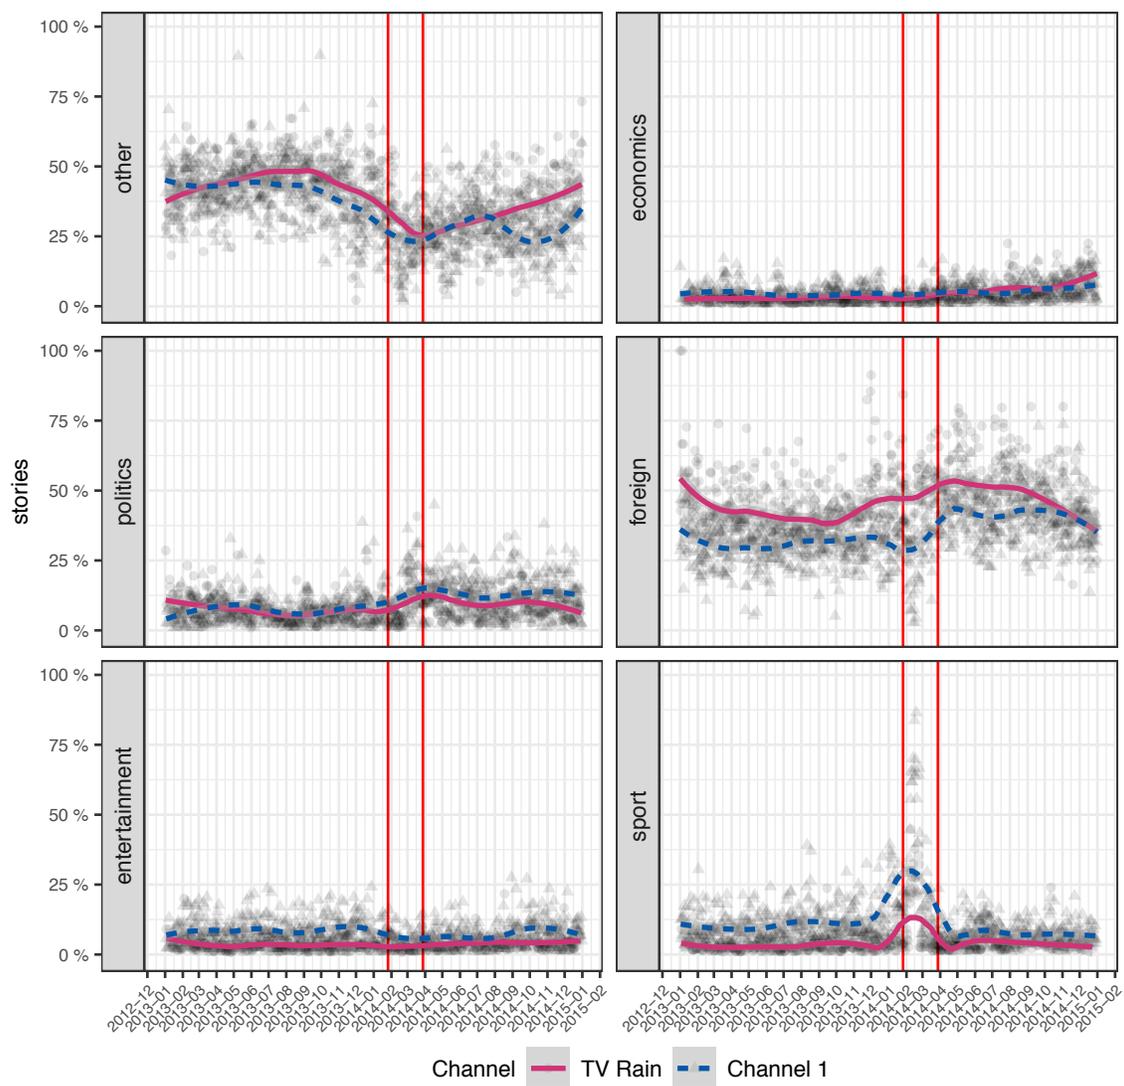


Figure A.4: **Dynamics of Topical Coverage (2013-2014)** Vertical red lines correspond to the date of TV Rain’s poll on the siege of Leningrad (26 January 2014) and the date of the end of telemarathon on TV Rain (28 March 2014).

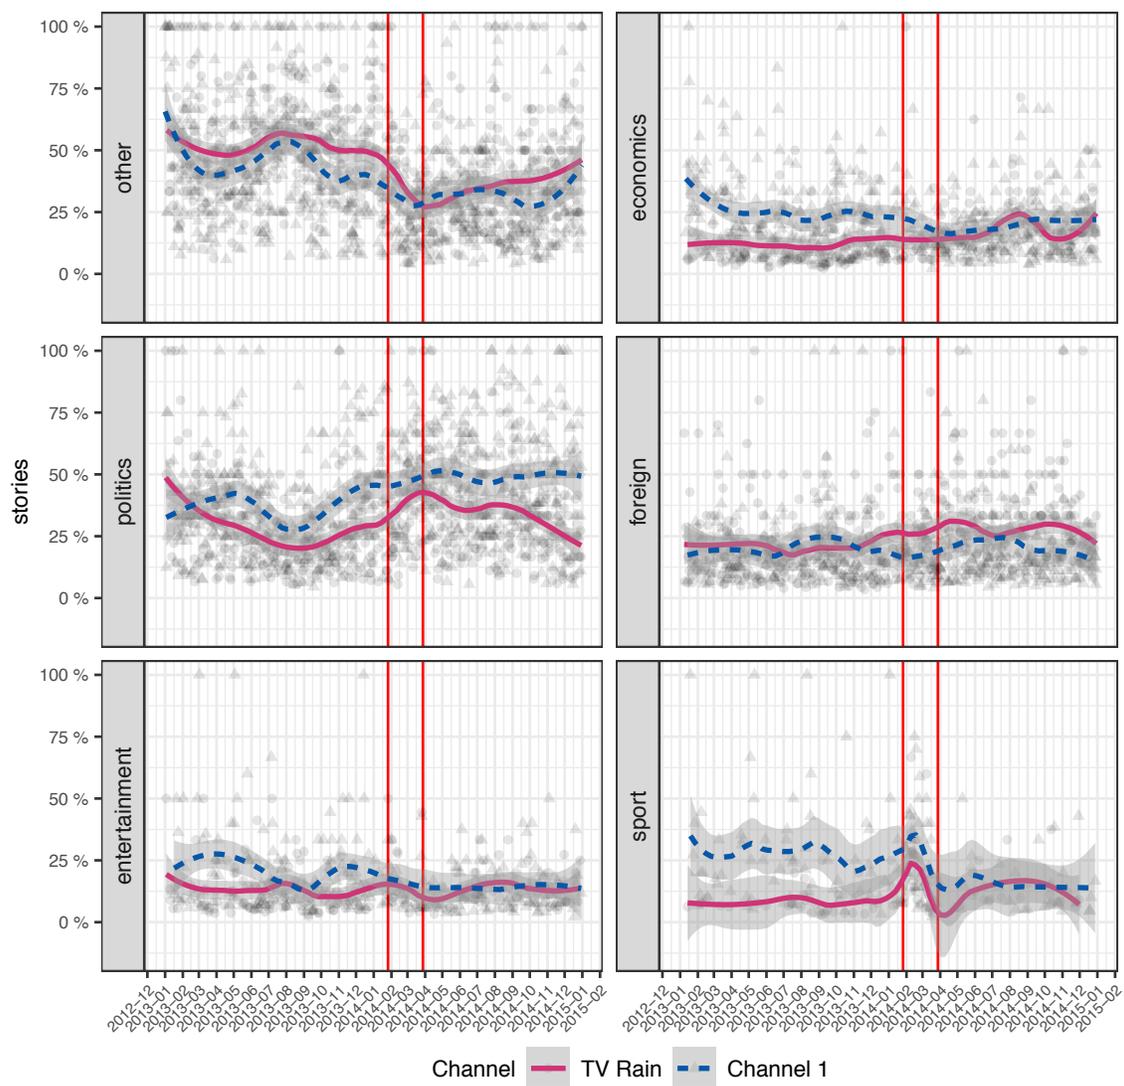


Figure A.5: **Dynamics of Topical Coverage in Government Stories (2013-2014)** Vertical red lines correspond to the date of TV Rain’s poll on the siege of Leningrad (26 January 2014) and the date of the end of telemarathon on TV Rain (28 March 2014).

B Sentiment Analysis

As noted in the article, we rely on sentiment measured machine translated Lexicoder as our main strategy for capturing the tone of political coverage in news stories.

To translate Lexicoder into Russian we proceed in two steps. First, we expand the original, partially stemmed, dictionary entries, by applying Lexicoder to a corpus of 90,266 CNN articles that cover the period from April 2007 to April 2015 (Hermann et al., 2015). This query finds 6,266 unique positive and 11,222 negative words that match string patterns in Lexicoder. Second, we translate each of these words into Russian, using Google Translate API.⁵ The final Russian version of Lexicoder, which excludes untranslated words, includes 14,402 entries.⁶ By the standard of other dictionaries (e.g. Muddiman et al., 2019), this is a very comprehensive list. We apply this new dictionary to each of the articles in the corpus of Russian media. The relative tone of each article is calculated by subtracting the total number of negative terms from the dictionary found in the text from the total number of positive terms, divided by the total number of tokens in an article.

As Table B.2 shows, Lexicoder provides better coverage (stories with detected positive/negative words) and higher correlation with manually validated sample, especially in stories shorter than 200 words than RuSentiLex (Loukachevitch and Levchik, 2016), the most comprehensive of the off-the-shelf alternatives.

Table B.2: Comparison between Russian-language sentiment dictionaries.

	Lexicoder (translated)	RuSentiLex
Positive Words	5175	3425
Negative Words	9227	9467
% Stories Covered (TV Rain)	90.9	68.7
% Stories Covered (Channel 1)	97.2	80.2
Correlation with crowd (Overall)	0.47	0.38
Correlation with crowd (>200 words)	0.53	0.52

We validate the sentiment indicator in a number of ways, described below. First, we sampled 25% of all articles in the corpus, translated them into English using Google Translate API, and applied the original Lexicoder dictionary. Next, we read through the hundred most positive and negative stories from each outlet that also mentioned government figures and found that the dictionary performed as expected. Finally, we compared the sentiment of 1,000 randomly chosen news stories to scores assigned by human coders recruited on the crowd-sourcing platform Yandex Toloka. Jointly, these validation checks suggest that the dictionary performs satisfactorily, and that the estimated tone offers a good approximation of the stories' overall sentiment.

⁵As recently demonstrated, the performance of machine translation approaches that of human translation for text analysis tasks that rely on bag-of-words assumption and disregard word order in the analysis (de Vries et al., 2018).

⁶See also Appendix B for comparison with an alternative Russian-language dictionary.

B.1 Translation check

As a validation check, we sampled 25% of all articles published between January 1 and June 30 2014 (2,094 news stories for TV Rain and 3,305 for Channel 1), translated them into English using Google Translate API, and applied the original Lexicoder dictionary. The correlation between the two sentiment scores is approximately 0.6.

B.2 Manual labelling: Crowd workers

As an alternative approach to measuring the sentiment of news stories we used Russian native speakers recruited through crowd-sourcing platform Yandex Toloka. We created two non-overlapping random samples of 500 news stories containing mentions of government figures for the period between 1 November 2013 and 30 June 2014. To facilitate the setting of an online labelling task and to use a linguistically meaningful unit of coding, we split these 1,000 stories into individual sentences, creating 27,395 sentences in total. The selection of crowd workers was restricted to Russian speakers from Belarus, Kazakhstan, Russia and Ukraine. Each sentence was labelled by 3 independent coders on the scale between -2 (very negative) to 2 (very positive). The crowd task instructions and the labellers' setup are shown in Figure B.6 and Figure B.7.

We aggregated the individual workers' labels into a single sentiment score per news article as follows. First, for each sentence, we calculated a mean score assigned by 3 labellers. Then, for each news article we averaged over all mean sentence scores comprising this article.

The correlation between the sentiment scores estimated with Lexicoder and crowd labellers is 0.47. It rises to 0.53 for longer articles (>200 words, 57% of the corpus), which indicates that tone is less ambiguous in those texts.

Инструкция

Обзор

Цель этого задания состоит в том, чтобы определить тональность новостных статей нескольких российских телеканалов. Вам будут представлены отдельные предложения из этих статей. Ваше задание состоит в оценке общей тональности (от -2, очень негативной до 2, очень положительной). Постарайтесь абстрагироваться от Вашей собственной оценки тех или иных событий и сфокусироваться на том, как эти статьи передают содержание и атмосферу событий и/или оценку действия лиц упомянутых в статье.

Шаги

1. Прочитайте целиком указанное предложение.
2. Установите является ли указанное предложение очень негативным (-2), негативным (-1), нейтральным/смешанным (0), положительным (1) или очень положительным (2) по тону.

Правила и Подсказки

Предложения могут быть классифицированы как очень негативные (-2), негативные (-1), нейтральные/смешанные (0), положительные (1) или очень положительные (2):

- Очень положительная означает, что предложение содержит явное восхваление описываемых событий и/или действий предпринятых упомянутыми действующими лицами.
- Положительная означает, что некоторые аспекты раскрывают положительный настрой, например, в виде рекомендации, благоприятного сравнения или подчеркнутое улучшение.
- Нейтральная означает, что предложение исключительно информативная по природе и не содержит каких-либо указаний на отношение автора к описываемым событиям.
- Негативная означает, что некоторые аспекты раскрывают негативный настрой. Такими, например, могут быть криминальная хроника, природный катаклизм или отрицательное сравнение.
- Очень негативная означает, что предложение подчеркивает трагическую сторону событий, содержит их прямую и резкую критику или оскорбления в адрес упомянутых в ней лиц.

Закреть

Figure B.6: Labelling instructions given to crowd workers on Yandex Toloka.

Подчеркнуто, что меры, предпринимаемые Россией, полностью адекватны сложившейся экстраординарной ситуации», - говорится в заявлении Ситуация в Украине 09.07.2018 Украину могут разменять.
Какова общая тональность этого предложения?

Очень Негативная Негативная Нейтральная Положительная Очень Положительная

Александр Баунов о переговорах Трампа и Путина 17.01.2018
Какова общая тональность этого предложения?

Очень Негативная Негативная Нейтральная Положительная Очень Положительная

Призеров летней Универсиады, которая прошла в прошлом году в Казани, чествовали сегодня в Москве.
Какова общая тональность этого предложения?

Очень Негативная Негативная Нейтральная Положительная Очень Положительная

Figure B.7: Fragment of labelling task used for crowd-coding on Yandex Toloka.

B.3 Most positive and negative stories around the discontinuity

As a validation exercise, we read through the most positive and negative stories from each outlet that also mentioned government figures and appeared within days 10 prior and following January 26, 2014. No single topic dominated positive and negative coverage this period.

TV Rain: The majority of the most positive stories that appeared on TV Rain between January 16 and 25 focused on domestic developments. For example, they discussed new investments in Dagestan, the introduction of educational reforms, and developments in the regions. One of the most positive news items from TV Rain prior to the poll engaged with Russia’s diplomatic developments and achievements in the previous year. Among TV Rain’s most negative-scored items, is a story in which an artist is quoted criticizing Russia’s new national idea. Other leading negative stories quote Vladimir Putin describing calls to boycott the Sochi Olympics as reflecting thinking of the past, a manifestation of competition, and criminal cases brought against family members of prominent politicians. Other items engaged with developments in the regions - a particular story noted that in Kazan it was impossible to put an imam through trial, as the judge could not speak Tatar. Amongst the most negative stories, two discussed developments in Ukraine.

In the 10 days following the ban, positive stories discussed police reforms and compensations given to the families of students who lost their lives during the shooting in Moscow’s 263 school. An entertainment story on a successful film project also makes the list of top ten positive news items around this time, as well as a foreign news story discussing the dismissal of Ukraine’s PM Nikolay Azarov.

A story titled ‘the Leningrad blockade (the topic of TV Rain’s poll) is not a topic for discussion: How the State Duma respondent to TV Rain’s poll’ tops the most negative story that appear on TV Rain between January 27 and February 5. Other stories discuss reforms and statements by Putin – the toughening of new legislation on extremism; a statement on Russia’s non-interference in Ukraine’s domestic affairs; and the award of a post-humous ‘Order of Courage’ award to a police officer in Volgograd who protected civilians during the terror attack in December 2013. The fifth most negative item discusses the resignation of the Deputy chief of General Staff resulting from disagreements with the head of general staff.

Channel 1: The two most positive stories that appeared on Channel 1 in the period preceding the poll discussed how the creators of children’s and youth channel “Karusel” were awarded by the Russian government, and the opening of the “Primordial/ Primitive Russia” festival in Moscow. Some of the most positive stories are also related to foreign affairs – emphasising, for example, the continuation of a conference on the settlement of the Syrian crisis in Switzerland. Another positive item notes that Russia will do everything possible to ensure that the upcoming Olympics in Sochi will be a true celebration of sports. Approximately five out of the ten most negative news items published on Channel 1 ahead of the January 26 poll discussed the crisis in Ukraine and Syria. Most of them portray Russia as ready to look for ways out of both conflicts despite international hostility. Other stories discussed challenges of school education and the situation in the financial market.

In the ten days following the poll, the majority of Channel 1 most negative stories focus on domestic issues – such as the regulation of energy consumption standards and utilities at the regional level, as well as the need to resolve issues related to traffic congestion. Other negative stories present a meeting Vladimir Putin held with members of Russia’s Security Council, during which Council members discussed the situation in Syria and Ukraine, as well as the upcoming winter Olympics. Coverage of the winter Olympics dominates positive coverage on Channel 1 in the 10 days following the publication of the poll on January 26. Four out of the ten most positive stories discuss the Olympics, while other positive news stories discuss the five-year anniversary of Patriarch Kyrill’s enthronement, on 1 February 2009.

We present several of these stories’ titles and sentiment scores below.

TV Rain

1. January 16, 2014: *«Myedialogiya»: Boris Dubrovskiy stal pyersonoy dnya poslye naznachyeniya na dolzhnost’ vrio gubernatora Chyelyabinskoy oblasti* - “Medialogia”: Boris Dubrovsky became the person of the day after being appointed to the post of Acting Governor of the Chelyabinsk Region. Sentiment Score: 0.03
2. January 21, 2014: *Syergyey Lavrov podvyol itogi diplomatichyeskogo goda*. Sergey Lavrov summed up the results, developments, of the diplomatic year. Sentiment Score: 0.03
3. February 2, 2014: *Vladimir Putin podpisal zakon o voyennoy politsii*. Vladimir Putin signed a law on military police. Sentiment Score: 0.05
4. February 4, 2014: *Syem’i pogibshih pri stryel’be v moskovskoy shkole №263 poluchat 5 mln roobley kompyensatsii*. Families of those killed in the shooting in Moscow school No263 will receive 5 million rubles of compensation. Sentiment Score: 0.02
5. January 23, 2014: *«Rossiya i vojna – eto sinonimi». Diorama «Stoyaniye na reke Ugre»*. “Russia and war are synonyms.” Diorama “Standing on the Ugra River”. Sentiment score: -0.035
6. January 19, 2014: *Putin nazval prizivi boykotirovat’ Olimpiadu rudimyentom starogo mishlyeniya*. Putin called calls to boycott the Olympics a rudiment of old thinking. Sentiment score: -0.025
7. January 27, 2014: *Blokada nye tema dlya diskussiy. Kak Gosduma otvyetila na opros Dozhdyia* "The blockade is not a topic for discussion. How the State Duma responded to the TV Rain poll. Sentiment score: -0.032
8. January 30, 2014: *Sobyanin uvolil glavu upravu Tvyerskogo rayona*. Sobyenin fired the head of the Tversky District Council. Sentiment score: -0.016

Channel 1

1. January 16, 2014: *«Sozdateli kanala «Karusel’» nagrazhdyeni pryemiyey pravityel’sstva Rossii* - The creators of the Karusel channel were awarded the Russian Government Prize. Sentiment Score: 0.085

2. January 16, 2014: *Rossiya sdyelayet vsyo vozmozhnoye, chtobi predstoyashshaya Olimpiada v Sochi stala nastoyashshim prazdnikom sporta.* Russia will do everything possible to make the upcoming Olympic Games in Sochi a real celebration of sports. Sentiment Score: 0.03
3. February 1, 2014: *Pyat' lyet nazad na pervosvyatitel'skiy prestol vstupil patriarh Moskovskiy i vseya Roosi Kirill.* Five years ago, Patriarch Kirill of Moscow and All Russia was enthroned. Sentiment Score: 0.03
4. February 4, 2014: *V Sochi otkrilas' 126-ya syessiya Myezhdunarodnogo olimpiyskogo komityeta.* The 126th session of the International Olympic Committee has opened in Sochi. Sentiment Score: 0.03
5. January 21, 2014: *Glava MID RF nazval provokatsionnimi razgovori naschyet vyeroyatnogo raskola Ookrainsi* The head of the Russian Foreign Ministry called provocative conversations about the likely split of Ukraine. Sentiment score: -0.019
6. January 19, 2014: *Myezhdoonarodnaya konfyeryentsiya po Sirii vizivayet noviye voprosi* The international conference on Syria raises new questions. Sentiment score: -0.016
7. February 03, 2014: *V Rossii nyeobhodimo sozdat' logisticheskiye tsyentri, chtobi bistro ryeshat' problyemi s zatorami na trassah.* In Russia, it is necessary to create logistics centers in order to quickly solve problems with congestion on highways. Sentiment score: -0.017
8. January 31, 2014: *Vladimir Pootin provyol sovyeshshaniye s postoyannimi chlyenami Sovbyeza RF* Vladimir Putin held a meeting with permanent members of the Russian Security Council. Sentiment score: -0.016

C RD Robustness Tests

C.1 Density of government stories around the discontinuity

We conduct a series of robustness tests to check the internal validity of our RD estimates. First, in line with literature on RD design (de la Cuesta and Imai, 2016; Cattaneo et al., 2020), we check for any sorting around the cutoff point in the running variable. As Figures C.8 and C.9 there is no evidence that TV Rain strategically chose or was forced to publish higher or lower number of stories on the days just before or just after the day of the poll. In addition we conduct a more direct McCrary test (McCrary, 2008) and find that we cannot reject the null hypothesis of no sorting in the running variable ($p = 0.23$ for all stories on TV Rain and $p = 0.30$ for news-only stories on TV Rain).

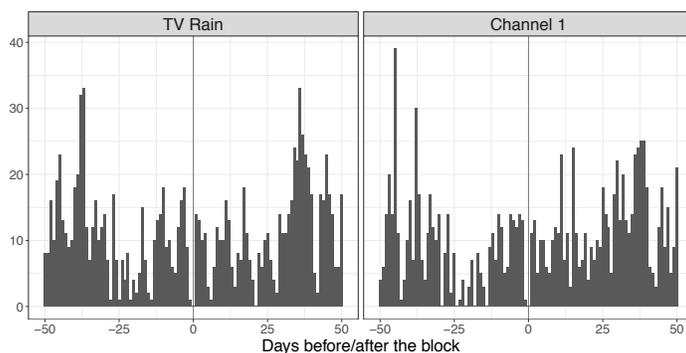


Figure C.8: Density of the number of government stories around the discontinuity (26 January 2014)

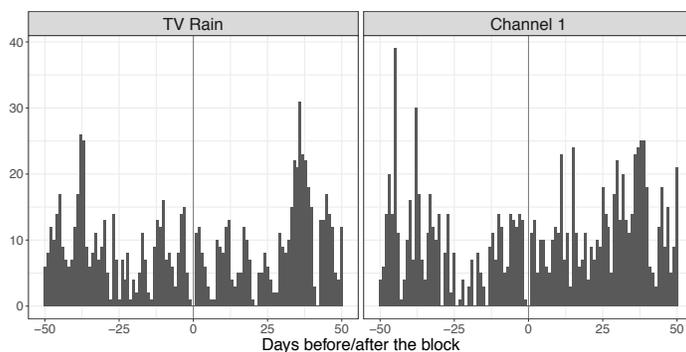


Figure C.9: Density of the number of government stories (TV Rain - news only) around the discontinuity (26 January 2014)

C.2 News-only results

As one of the robustness tests, we check whether the main empirical results hold if we restrict our analysis to a more narrow definition of news on TV Rain (see section A.1 of the present Appendix for more details). As [Figure C.10](#) shows, we observe a similar increase in sentiment of news stories mentioning government figures on TV Rain as reported in the main results.

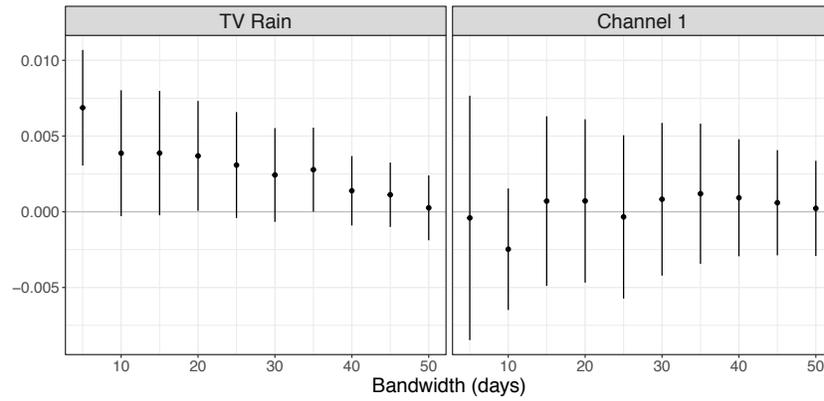


Figure C.10: RD effects of poll (26 January 2014) on government tone in news stories

C.3 Placebo cutoff

We also check whether the increase in the tone of stories mentioning government figures on TV Rain can only be observed around the time of the controversial poll and subsequent exclusion of the channel from cable and satellite packages. In order to do so, we run the same RD model around the time of a placebo poll (16 January 2014, 10 days prior to the actual poll). As [Figure C.11](#) indicates, for narrow bandwidths we do not find any significant changes in the tone of stories on TV Rain.

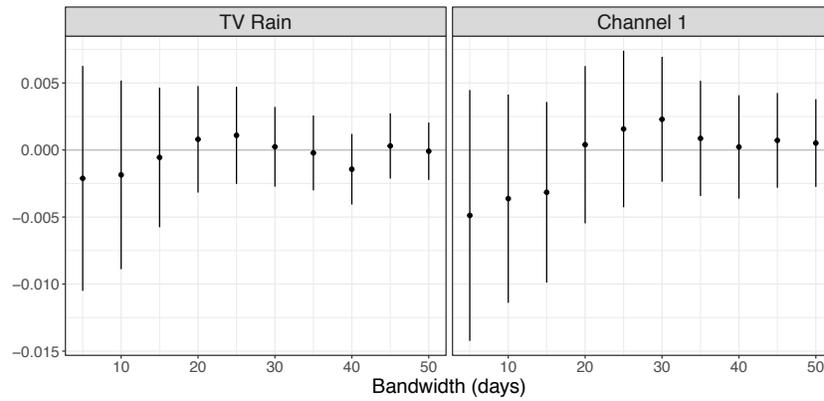


Figure C.11: **RD effects of placebo poll (16 January 2014) on government tone**

C.4 Placebo outcome

We use other attributes of stories from TV Rain and Channel 1 as our placebo outcome. In particular, we test whether the block had any effect on the length of published stories. As Figure C.12 shows, we find no significant changes in stories length for TV Rain.

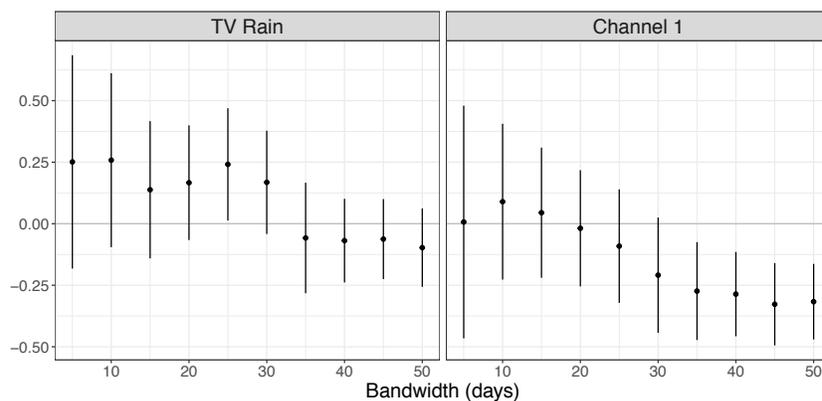


Figure C.12: RD effects of poll (26 January 2014) on government stories length (words)

C.5 RD results using local-polynomial models

Table C.3: Local-polynomial RD Estimates with Robust Confidence Intervals.

	Conventional	Bias-Corrected (conventional SEs)	Bias-Corrected (robust SEs)
Point estimate	0.0046	0.0051	0.0051
$P > z $	0.0072	0.0030	0.0095
CI Lower	0.0013	0.0017	0.0013
CI Upper	0.0080	0.0085	0.0090
Optimal Bandwidth (days)	83	83	83
Number of observations	8298	8298	8298

D Cosine Similarity

To measure similarity between TV Rain and Channel 1 we employ cosine similarity, a common measure of distance between text documents (Hager and Hilbig, 2020). This metric, originally developed in the field of information retrieval (Manning et al., 2008), allows us to capture both linguistic (analogous coverage that uses identical words) and content (coverage of the same events) similarity. More technically, cosine similarity between two documents \mathbf{d}_i and \mathbf{d}_j is defined as:

$$\text{sim}(\mathbf{d}_i, \mathbf{d}_j) = \frac{\sum_{z=1}^n d_{i,z} d_{j,z}}{\sqrt{\sum_{z=1}^n d_{i,z}^2} \sqrt{\sum_{z=1}^n d_{j,z}^2}}$$

where each of the documents \mathbf{d} is a vector representation of n unique terms (words) such that $\mathbf{d}_i = (w_{i,1}, w_{i,2}, \dots, w_{i,z}, \dots, w_{i,n})$ and $\mathbf{d}_j = (w_{j,1}, w_{j,2}, \dots, w_{j,z}, \dots, w_{j,n})$ represent two documents that are being compared. Each element of the vector $w_{i,z}$ is the number of times a term z occurs in document i . In our case we use TF-IDF (term-frequency inverse document frequency) weighted counts as opposed to raw term counts. We calculate the weighted counts separately for each of the channels, using terms' frequencies in each of the corpora as weights.

Our primary quantity of interest is the average similarity across all pairs of news articles published on TV Rain and Channel 1 in a given day. In other words, we are interested in $\bar{\text{sim}}(d_{r,t}, d_{o,t})$, where $d_{r,t}$ and $d_{o,t}$ are news stories broadcasted on TV Rain and Channel 1, respectively, on day t . Individual estimates of cosine similarity range between 0 (two documents are completely different) and 1 (two documents are exactly the same). Figure D.13 shows the distribution of raw between-story estimates of similarity.

In addition to calculating between-channel daily average cosine similarity we also estimate standard errors, presented in Figure 3 of the manuscript. To account for the increased coverage of the Sochi Olympics and its potential effect on between-channel similarity, we exclude all news stories that contain the mentions of 'Sochi' or were classified as covering sports by our model.

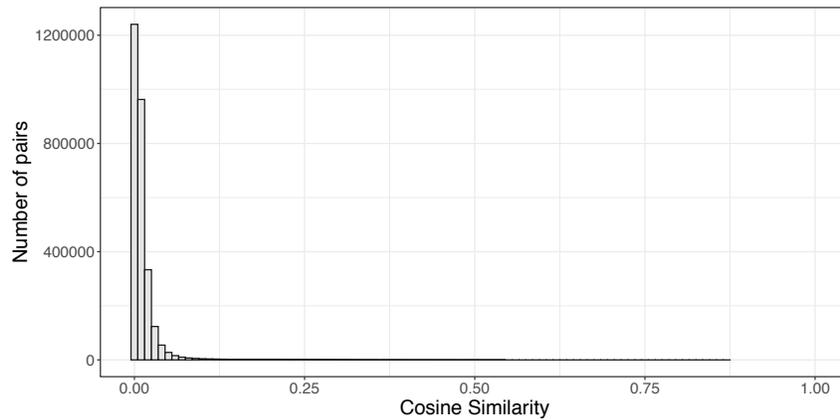


Figure D.13: **Distribution of raw cosine similarity estimates between TV Rain and Channel 1 new stories.**

E Additional Evidence

E.1 Manually coded critical coverage of government figures on TV Rain

We read through the titles of stories that mentioned the government figures listed in Table 1, and published on TV Rain between January 1 and April 31, 2014 ($N=1079$), manually coding whether these could be perceived to be critical of the Russian government or not. Stories critical to the government are those which discuss corruption scandals and those in which government figures receive criticism from others at home or abroad. Stories that discuss events with negative connotations for the regime, such as the sanctions, anti-government protest, clashes among government figures, and poor economic performance, are also coded as being critical to the government.

The share of TV-Rain stories classified as critical to the ruling regime declined two-fold between January and February 2014. Between February and March 2014, the share of critical stories tripled. While approximately one in five (23%) stories were classified as critical to the regime in January and February 2014, around one in every three stories published in March (35%) and April 2014 (29%) were of critical content. This is consistent with the sentiment estimates presented in the manuscript, implying that coverage on TV Rain became more critical in the early spring months.

What is more, in a random sample of 685 manually classified news items, the 150 stories classified as expressing critical content had an average sentiment score of $-.001$ (SD: $.014$). The remaining 535 stories had an average sentiment score of $.004$ (SD: $.017$). This evidence lends further support for the sentiment estimates obtained using the dictionary approach.

References

- Cattaneo, M. D., N. Idrobo, and R. Titiunik (2020). *A Practical Introduction to Regression Discontinuity Designs: Foundations*. Elements in Quantitative and Computational Methods for the Social Sciences. Cambridge University Press.
- de la Cuesta, B. and K. Imai (2016). Misunderstandings About the Regression Discontinuity Design in the Study of Close Elections. *Annual Review of Political Science* 19(1), 375–396.
- de Vries, E., M. Schoonvelde, and G. Schumacher (2018). No Longer Lost in Translation: Evidence that Google Translate Works for Comparative Bag-of-Words Text Applications. *Political Analysis* 26(4), 417–430.
- Hager, A. and H. Hilbig (2020). Does Public Opinion Affect Political Speech? *American Journal of Political Science* 64(4), 921–937.
- Hermann, K. M., T. Kocisky, E. Grefenstette, L. Espeholt, W. Kay, M. Suleyman, and P. Blunsom (2015). Teaching Machines to Read and Comprehend. pp. 1–9.
- Loukachevitch, N. and A. Levchik (2016, May). Creating a General Russian Sentiment Lexicon. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC'16)*, Portorož, Slovenia, pp. 1171–1176. European Language Resources Association (ELRA).
- Manning, C. D., P. Raghavan, and H. Schütze (2008). *Introduction to Information Retrieval*. Cambridge: Cambridge University Press.
- McCrary, J. (2008). Manipulation of the running variable in the regression discontinuity design: A density test. *Journal of Econometrics* 142(2), 698–714. The regression discontinuity design: Theory and applications.
- Muddiman, A., S. C. McGregor, and N. J. Stroud (2019). (Re)Claiming Our Expertise: Parsing Large Text Corpora With Manually Validated and Organic Dictionaries. *Political Communication* 36(2), 214–226.