DISCUSSION PAPER SERIES

DP16233

Demand for Online News under Government Control: Evidence from Russia

Andrey Simonov and Justin Rao

DEVELOPMENT ECONOMICS

INDUSTRIAL ORGANIZATION

PUBLIC ECONOMICS

ISSN 0265-8003

Demand for Online News under Government Control: Evidence from Russia

Andrey Simonov and Justin Rao

Discussion Paper DP16233
Published 08 June 2021
Submitted 03 June 2021

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

Demand for Online News under Government Control: Evidence from Russia

Abstract

We examine the nature of consumer demand for government-controlled online news outlets in Russia, testing whether such demand reflects a preference for pro-government ideological coverage, or other factors unrelated to outlets' ideological positions. We detect government-sensitive topics and measure outlets' news reporting decisions from news article texts, and estimate a structural model of demand for news using detailed browsing data that traces individual-level consumption. The average consumer has a distaste for pro-government ideology but a strong persistent taste for state-owned outlets, primarily driven by third-party referrals and non-sensitive news content. We discuss implications for online media control and media power.

JEL Classification: C11, C55, D72, L15, L82, L86, M31, P26

Keywords: media, media capture, Censorship, Demand for News, product differentiation, Text as Data

Andrey Simonov - as5443@gsb.columbia.edu
Columbia University and CEPR

Justin Rao - justin.rao@booking.com
Booking.com BV

Acknowledgements
We thank Charles Angelucci, Ruben Enikolopov, Andrey Fradkin, Matt Goldman, Elisabeth Honka, Sanjog Misra, Chris Nosko, Andrea Prat, Miklos Sarvary, Mehdi Shadmehr, Bradley Shapiro, Raluca Ursu, Ali Yurukoglu, participants of the 16th Trans-Atlantic Doctoral Conference, 2016 AMA-Sheth Doctoral Consortium, 2016 Economics of Media and Communications Conference, 2017 Marketing Science Conference, 15th Media Economics Workshop and 4th Media Bias Conference, Behavioral Insights from Text at Wharton, 2020 Wallis Institute Annual Conference, seminar participants at the University of Chicago, Stigler Center at Chicago Booth, HKUST, University of Minnesota-Carlson, University of Rochester-Simon, UCSD-Rady, Duke-Fuqua, LBS, Columbia, Southern Methodist University-Cox, UC Berkeley-Haas, Stanford, CMU, UCLA-Anderson, University of Southern California-Marshall, Moscow State University, University of Washington-Foster, NYU-Stern, the FCC, and Northwestern for advice and suggestions. We thank Sofi a Avramenko, Yulia Kalinnikova, Dmitry Leshchuk, Lily Shevchenko, Semyon Tabanakov for excellent research assistance. Andrey is grateful to his advisors and dissertation committee members Pradeep Chintagunta, Jean-Pierre Dube, Matthew Gentzkow and G�unter Hitsch for their guidance and support. All opinions represent our own and not those of our current or past employers. All remaining errors are our own.

Demand for Online News under Government Control: Evidence from Russia*

Andrey Simonov
Columbia University & CEPR

Justin Rao†
Booking.com BV

May 22, 2021

Abstract

We examine the nature of consumer demand for government-controlled online news outlets in Russia, testing whether such demand reflects a preference for pro-government ideological coverage, or other factors unrelated to outlets' ideological positions. We detect government-sensitive topics and measure outlets' news reporting decisions from news article texts, and estimate a structural model of demand for news using detailed browsing data that traces individual-level consumption. The average consumer has a distaste for pro-government ideology but a strong persistent taste for state-owned outlets, primarily driven by third-party referrals and non-sensitive news content. We discuss implications for online media control and media power.

Keywords: Media, Media Capture, Censorship, Demand for News, Consumer Heterogeneity, Differentiated Products, Text as Data

JEL Classification: C11, C55, D72, L15, L82, L86, M31, P26

---

*Simonov (corresponding author): asimonov@gsb.columbia.edu; Rao: justin.rao@booking.com. We thank the editor and anonymous referees for detailed and helpful feedback throughout the review process. We also thank Charles Angelucci, Ruben Enikolopov, Andrey Fradkin, Matt Goldman, Elisabeth Honka, Sanjog Misra, Chris Nosko, Andrea Prat, Miklos Sarvary, Mehdi Shadmehr, Bradley Shapiro, Raluca Ursu, Ali Yurukoglu, participants of the 16th Trans-Atlantic Doctoral Conference, 2016 AMA-Sheth Doctoral Consortium, 2016 Economics of Media and Communications Conference, 2017 Marketing Science Conference, 15th Media Economics Workshop and 4th Media Bias Conference, Behavioral Insights from Text at Wharton, 2020 Wallis Institute Annual Conference, seminar participants at the University of Chicago, Stigler Center at Chicago Booth, HKUST, University of Minnesota-Carlson, University of Rochester-Simon, UCSD-Rady, Duke-Fuqua, LBS, Columbia, Southern Methodist University-Cox, UC Berkeley-Haas, Stanford, CMU, UCLA-Anderson, University of Southern California-Marshall, Moscow State University, University of Washington-Foster, NYU-Stern, the FCC, and Northwestern for advice and suggestions. We thank Sofia Avramenko, Yulia Kalinnikova, Dmitry Leshchuk, Lily Shevchenko, Semyon Tabanakov for excellent research assistance. Andrey is grateful to his advisors and dissertation committee members Pradeep Chintagunta, Jean-Pierre Dubé, Matthew Gentzkow and Günter Hitsch for their guidance and support. All opinions represent our own and not those of our current or past employers. All remaining errors are our own.

†Work conducted while a Senior Researcher at Microsoft Research.

1 Introduction

On May 2, 2014, an unprecedented outbreak of violence between the supporters and opposers of the new Ukrainian government in the city of Odessa led to 48 deaths. The story was widely covered by the Russian media, but the coverage differed drastically across news outlets. Independent Russian news outlets reported that both supporters and opposers of the new Ukrainian government were throwing Molotov cocktails, and that most likely the fire was started due to the carelessness of the government opposition members inside the building. In contrast, government-controlled (GC) Russian news outlets reported that the blame laid squarely with radical Ukraine government supporters. The ideological slant in the GC outlets' coverage was mixed with objectively false information, exemplified by the title of one of the articles: "116 people burned alive by fascists in Odessa" (vesti.ru, 2014).

These two drastically different takes on a particular story characterize a typical choice set of the online news consumers in many authoritarian countries like Russia. In these markets, both independent news outlets – the ones that are neither owned by nor influenced indirectly by the government – and government-owned or influenced outlets co-exist, with independent outlets being less constrained in their choice of the ideological slant. And yet, despite this competitive advantage in product differentiation, many consumers choose the GC over independent outlets. In the case of the Russian online news market in late 2014, 4 out of the top 5 outlets were either government-owned or potentially influenced.[1]

In this paper, we separate out two potential drivers of demand for the GC news outlets. First, consumers might read the GC outlets because of tastes for the pro-government bias in sensitive news coverage. Such tastes can stem from a preference for like-minded news (Gentzkow and Shapiro, 2010) – for instance, Russian pools have shown an 80% approval rating of President Putin in 2014-2016 (Economist, 2016) – or from the "conscientious" readers' interest in the pro-government news framing (Mullainathan and Shleifer, 2005). Second, consumers might be driven to the GC outlets by other features of these outlets unrelated to their ideology – for instance, modern website design, video content, referrals by news aggregators, or accumulated brand capital. Collectively, such "persistent preferences" of consumers for outlets could allow governments to exercise media capture even in competitive markets – the government can invest in outlets' features that increase consumers' outlet-level tastes, allowing to capture the readers' attention and to control their ideological news diet, without capturing all news producers in the market (Besley and Prat, 2006).

[1] This is based on publicly available statistics (liveinternet.ru, 2014) and confirmed by our browsing data. For the news outlet classification, see Table 1.

We ground our investigation in a simple model of news production and consumption. In the model, consumers choose news outlets based on their outlet-level persistent preferences and tastes for outlets' ideological coverage. The importance of outlets' ideological coverage changes from day to day, depending on the volume of realized sensitive news events. At the extreme, on days with no sensitive news events to report, news outlets' ideological positions do not matter and consumers choose outlets only based on their outlet-level persistent preferences. In contrast, on days with a lot of realized government-sensitive news events, outlets' ideological positions become prominent and might change readers' choices. These changes in consumers' outlet choices in response to changes in sensitive news volume are the key variation that identifies readers' ideological tastes when we estimate the empirical model.[2]

The focus of our empirical study is the online news market in Russia in 2013-2015. At that time, online news consumers in Russia had a broad choice of ideologically-different news outlets – the market contained a large number of GC and independent outlets, as well as some outlets in-between (formally independent but with ties to the government), and Russian-language versions of Ukrainian or other international outlets. Overall, our main sample contains top 48 online news outlets classified into one of these groups.[3]

Our empirical strategy relies on two novel datasets. First, for the top 48 online news outlets in Russian language, we collect all accessible publication records – 3.9 million online news articles – written between March 2013 and April 2015. News articles data include the article URL, date, title and text. Using a simple classification algorithm that we develop – which compares article texts published by the GC and independent outlets, looking for differences in news coverage that apply to all or most GC and independent outlets – we detect two major government-sensitive news topics.[4] The first topic covers the events that are systematically underused by the GC outlets, likely due to censorship. These events mainly correspond to political protests, opposition and corruption (hereafter "POC" news). The second topic is the Ukraine crisis of 2013-2015, where the GC outlets use the pro-Russia ideological framing – systematically different language to describe the events (Prat and Strömberg, 2013) – in their news reporting. For instance, the GC outlets report that

[2]Using such revealed-preference measure of readers' ideological tastes is particularly important in the context of countries with limited freedom of speech (Kuechler, 1998).

[3]We use the information on ownership structure (e.g., Djankov et al., 2003) and reports of alleged government influence to classify the outlets.

[4]The objective of our algorithm is different from common methods of text classification used in the literature (Gentzkow and Shapiro, 2010; Gentzkow et al., 2019) that search for language that is most *predictive* of the outlet type. Our goal is to find language differences that apply to all GC and independent outlets. In the validation exercise done with manual word coding, we show that our algorithm outperforms all feasible alternative methods of sensitive news detection.

Russia has "reunited" with Crimea, while Ukrainian and international outlets characterize this reunion as an "annexation" and "occupation."[5]

Detected government-sensitive topics and text data provide two important ingredients for our empirical strategy. First, they give us a measure of the relative importance of sensitive news over time, which we construct as a share of articles about the sensitive news topic on a given day across all outlets. We treat this measure as an exogenous variable determined by day-to-day news realizations. Second, we characterize the reporting and ideological positions of the news outlets by their share of coverage of sensitive news and share of articles with ideological framing. News outlets hold relatively stable reporting and ideological positions, showing a limited reaction to changes in the relative importance of sensitive news over time.

The second novel dataset that we use is a large panel of browsing records from Internet Explorer (IE) Toolbar data. In the data, we observe around 285 thousand IE Toolbar users who visited at least one of the top 48 news outlets between November 2013 and April 2015, providing us with the individual-level news consumption data. The aggregate outlet-level news consumption by the IE users closely tracks the population-level metrics, with an average correlation of 85.8% across outlets.[6]

The model-free evidence strongly suggests that the average consumer prefers the coverage of sensitive news by the independent outlets and visits the GC outlets due to high persistent tastes. On days with more POC and Ukraine-crisis news events, consumers are more likely to navigate to news outlets that cover more of the (censored) POC news and have less pro-government ideological slant in the Ukraine crisis coverage, which tend to be independent, Ukrainian, and international outlets. In contrast, the GC outlets get more traffic due to various sources of persistent preferences, such as referrals by third parties (e.g. Yandex News, the largest news aggregator in Russia), consumers' landings on pages with video content and special projects, and landings on the non-sensitive news articles (e.g. about celebrities) that later spillover into sensitive news consumption.

We quantify consumer preferences by estimating a structural model of demand for news. The majority of consumers in the Russian online news market, 58.85% and 67.2%, prefer

[5]The language differences that we find fit well with the reports of independent journalists monitoring the news coverage of the Ukraine crisis, and the implied ideological positions of news outlets closely track the ideological positions based on a manual classification done by two independent research assistants.

[6]Based on the top 7 online news outlets for which we observe the population-level data. The data further suggests that IE Toolbar users are older, less interested in entertainment websites and more likely to visit business-focused news websites than an average internet user in Russia. Anecdotally, this is due to a common usage of IE browser in the office setting. Section 3.3.1 discusses the data differences and their implications in detail.

more coverage of the POC and Ukraine-crisis news, respectively, and 54.98% of consumers prefer less pro-government framing in the Ukraine-crisis news. Only a minority (39.9%) of consumers behave like conscientious readers. Thus, a preference for less pro-government framing suggests that independent – not GC – outlets have a more like-minded ideology with the majority of online news consumers. However, the vast majority (87.95%) of consumers have higher persistent preferences for the GC than independent outlets. As a result, GC outlets have a one-third market share advantage over independent outlets on days with no sensitive news. Correlationally, these high persistent preferences for the GC outlets are driven by indirect traffic – Yandex News in particular – and non-sensitive news topics coverage.

In a series of counterfactual simulations, we use preference estimates to assess the effect of GC outlets' suboptimal ideological positions and superior persistent preferences on their market share and media power (Prat, 2018). GC outlets sacrifice 15.3% of market share due to their pro-government ideological positions, translating to as much as $15.6 million in foregone display advertising revenues per year – a tiny loss compared to $1.21 billion in subsidies to mass media in Russia in 2015 (rbc.ru, 2015). In contrast, without the high persistent preferences of consumers, GC outlets would lose 54.3% – a 3.5 times more compared to an effect of their ideological positions. These high persistent outlet preferences substantially increase the attention share and media power of the GC outlets – they currently command a 33.8% attention share, but would only obtain 17.92% in the absence of persistent preferences. The attention share of 33.8% could enable the government to swing the outcome of the 25-75% vote share election through media persuasion (Prat, 2018).

Our work contributes to the literature on the political economy of mass media. We show that by investing in controlled outlets' quality and non-sensitive content, governments can exercise censorship in relatively competitive online news markets – documenting a new strategy of online media capture and censorship (e.g., see Roberts, 2020, and Zhuravskaya et al., 2020 for recent surveys).[7] An exposure to sensitive news content – with potential later changes in consumers' beliefs and actions – as a by-product of other news consumption resembles the unintended effects of consumption of entertainment media documented in the contexts of crime (Dahl and DellaVigna, 2009), family choices (La Ferrara et al., 2012), nationalism (DellaVigna et al., 2014), education (Kearney and Levine, 2015), and voting

[7] More broadly, our results add to the empirical literature (e.g., Durante and Knight, 2012; King et al., 2013, 2014; Roberts, 2014; Bai et al., 2015; Garcia-Arenas, 2016; Hobbs and Roberts, 2018; Knight and Tribin, 2019; Szeidl and Szucs, 2021) and inform the theoretical literature (e.g., Besley and Prat, 2006; Petrova, 2008; Egorov et al., 2009; Prat and Strömberg, 2013; Edmond, 2013; Gehlbach and Sonin, 2014; Shadmehr and Bernhardt, 2015; Prat, 2018) on media capture.

outcomes (Durante et al., 2019), among others.[8] We identify preferences for slant from changes in news consumption in response to exogenous news event realizations, adding to other natural experiments used to identify ideological preferences of consumers (Gentzkow and Shapiro, 2010; Martin and Yurukoglu, 2017).[9] The estimated outlet-level preferences align well with the recently documented persistence in news consumption and its effect on the consumed ideological slant, beliefs and behavior of news readers (Knight and Tribin, 2019; Chen and Yang, 2019; Levy, 2021). Our analysis of GC outlets' ideological positions complements other work on media control in autocracies (e.g., Qin et al., 2017, 2018).

Our paper also contributes to the literature on the industrial organization of news markets. To our knowledge, we are the first to use individual-level data to estimate a news demand model that separates out consumers' ideological preferences from (heterogeneous) persistent tastes for outlets.[10] Our estimates contribute to the growing empirical literature studying consumption, production, and competition in online news markets (e.g., Flaxman et al., 2016; Sen and Yildirim, 2016; Athey et al., 2017; Cagé et al., 2020; Calzada and Gil, 2020; George and Hogendorn, 2020).[11] We use text data to measure product differentiation (e.g., Groseclose and Milyo, 2005), contributing to the literature on product differentiation in media markets (e.g., Berry and Waldfogel, 2001; Fan, 2013; Sweeting, 2013; Jeziorski, 2014) and in economics and marketing more broadly (e.g., Mazzeo, 2002; Hortaçsu and Syverson, 2004; Seim, 2006; Draganska et al., 2009; Eizenberg, 2014; Sullivan, 2017; Wollmann, 2018).

The next section builds a stylized model of demand for news. Section 3 describes the Russian online news market and our data. Section 4 describes the classification of the sensitive news and characterizes outlets' reporting. Section 5 presents model-free evidence on consumer preferences. We describe our empirical specification in Section 6 and present the preference estimates and counterfactual simulations in Section 7. Section 8 concludes.

[8]DellaVigna and La Ferrara (2015) provides an overview of the related literature, not covering politics. For politics, the literature on media persuasion includes DellaVigna and Kaplan (2007); Enikolopov et al. (2011); Yanagizawa-Drott (2014); Adena et al. (2015); Martin and Yurukoglu (2017), among others; see DellaVigna and Gentzkow (2010) and Strömberg (2015) for reviews.

[9]Eisensee and Strömberg (2007) use exogenous news event realizations to examine how media coverage of natural disasters affects the U.S. government response, while Durante and Zhuravskaya (2018) use it to examine the strategic timing of Israeli attacks on Palestine. Dahl and DellaVigna (2009) use variation in media content over time to identify the effects of violent movies on crime rates. In a recent paper close to our empirical context, Melnikov (2019) uses changes in the U.S. dollar/Russian ruble exchange rates to examine the effects of propaganda on government popularity in Russia.

[10]Our demand model builds on the models in Strömberg (2004) and Gentzkow and Shapiro (2015).

[11]Broader literature on news markets include George and Waldfogel (2006); Gentzkow et al. (2011, 2014); Zhu and Dukes (2015); Angelucci and Cagé (2019); Cagé (2020); Angelucci et al. (2020), among others.

2 A Stylized Model

In this section, we present a stylized model of the news supply and demand in the markets with partial government control.

2.1 Basic Model

Suppose there are two news outlets in the market, A and B. Every day, outlets produce one unit of news product, such as a newspaper or a set of online articles. The news product consists of commodities of two types: news articles about sensitive and not sensitive topics for the incumbent government. For now, assume that any publications about sensitive events are bad for the government; the government is indifferent about the non-sensitive publications.

Consumers have stable and heterogeneous preferences for sensitive and non-sensitive news articles. Assume that on day t consumers choose at most one outlet or decide not to consume the news altogether. Consumer i chooses an option with the highest utility among

$$\begin{aligned} U_{ijt} = \beta_i x_{jt}^S + \lambda_i x_{jt}^{NS} + \epsilon_{ijt} \ : \ j \in \{A, B\}, \{x_{jt}^S, x_{jt}^{NS}\} \in [0,1], \\ U_{it0} = \epsilon_{it0}, \end{aligned} \tag{1}$$

where x_{jt}^S and x_{jt}^{NS} are the amounts of sensitive and non-sensitive news in outlets j's coverage, respectively, and ϵ_{ijt} is an unobserved idiosyncratic shock to the utility. Following the standard discrete-choice literature (e.g., Train, 2009), we can derive consumer demand for news outlets' products $\{D_A, D_B\}$, which is driven by the distribution of consumer preferences, $\{\beta, \lambda\}$, commodity choices of the news outlets, $\{x_{jt}^S, x_{jt}^{NS}\}$, and the distribution of the idiosyncratic shocks, ϵ_{ijt}.

News outlets make daily production decisions on the amount of sensitive and non-sensitive news commodities in their product, x_{jt}^S and x_{jt}^{NS}. The news commodities are costly to produce as they require journalists to investigate news topics. However, it is less costly to produce news about a certain topic on days when a lot of topic-related events happen. For example, writing sensitive news is more costly on the days when no sensitive news events have happened, as production requires more investigation. More formally, news production costs $c_t^S(x_{jt}^S, V_t^S)$ and $c_t^{NS}(x_{jt}^{NS}, V_t^{NS})$ are decreasing in the the amount of the events of the same type that happen on day t, $\{V_t^S, V_t^{NS}\} \in [0, 1]$.

Finally, suppose that outlet A is government-controlled and outlet B is independent. Given that the government dislikes sensitive news publications, it exercises censorship by imposing additional costs of production of sensitive news on outlet A, $c^G(x^S_{At})$.[12] The shape of the $c^G(\cdot)$ function is determined by the government's objective function.

Two observations follow from this setting. First, controlled outlet A would choose to produce less sensitive news than independent outlet B, $x^{S*}_{At} \leq x^{S*}_{Bt}$, as it faces higher marginal costs of sensitive news production.[13] Second, unless the shape of $c^G(\cdot)$ function is highly concave – meaning that the government mainly cares about the first few sensitive stories reported by outlet A – the difference in the amount of sensitive news produced by the outlets, $x^{S*}_{Bt} - x^{S*}_{At}$, is increasing in V^S_t. Intuitively, when there is no sensitive news to report, $V^S_t = 0$, it can be very costly for both news outlets to produce sensitive news (high c^S_t), so both outlets produce very low x^{S*}_{jt}. In contrast, when there is a lot of sensitive news to report, the cost of sensitive news production is low and c^G plays a more important role. In Section 4.2, we confirm that the difference in the sensitive news reporting between the GC and independent outlets increases with V^S_t. We further show that news outlets tend to report a fixed share of sensitive news that does not depend on V^S_t, allowing us to decompose $x^S_{jt} \approx V^S_t \bar{x}^S_j$, where $\bar{x}^S_j$ is the share of sensitive news reported by the news outlet j.[14]

2.2 Extensions

Persistent preferences. Apart from the news commodities supplied, outlets can differentiate themselves in a variety of ways, such as website design, overall quality of the news coverage, other content of the website, and promotion by third parties (Strömberg, 2004). Consumers can like or dislike these attributes of the outlets,

$$U_{ijt} = \alpha_{ij} + \beta_i x^S_{jt} + \lambda_i x^{NS}_{jt} + \epsilon_{ijt} \ : \ j \in \{A, B\}, \{x^S_{jt}, x^{NS}_{jt}\} \in [0, 1], \quad (2)$$

where α_{ij} represent the matching value between consumer i's preferences and features of the news outlet j. These persistent preferences might also include the effects of habit formation and inertia in news consumption.

Space constraints. Up to this point we have assumed that news outlets make two separate choices of x^S_{jt} and x^{NS}_{jt} that only depend on the realizations of V^S_t and V^{NS}_t. In

[12] For example, a government that instructs a news outlet not to cover a story or omit some facts from a story about a corruption scheme organized by some officials is censorship. Media economics literature refers to censorship as "issue and fact bias" (Prat and Strömberg, 2013) or as "filtering or selection of news" (Gentzkow et al., 2016). Censorship works through the effects of agenda setting (McCombs and Shaw, 1972) and priming (Iyengar and Kinder, 1987).

[13] Online Appendix A presents an extended discussion of the news outlets' optimization problem.

[14] Since our focus is on estimating consumer preferences, we stop short of estimating the shape of the cost functions. Instead, we take the editorial strategies of the news outlets as given.

practice, outlets operate under capacity constraints; their coverage cannot exceed a certain number of articles, for example, because of a fixed amount of space in the newspaper or a limited amount of journalists. We simplify the model by assuming that the news outlets always have to fill a strict amount of space, $x_{jt}^S + x_{jt}^{NS} = 1$, so the only thing that varies over time is the ratio of the produced sensitive and non-sensitive news commodities.[15]

Using this simplification, we can re-write consumer utilities as

$$U_{ijt} = \alpha_{ij} + \beta_i x_{jt}^S + \lambda_i(1 - x_{jt}^S) + \epsilon_{ijt} = (\alpha_{ij} + \lambda_i) + (\beta_i - \lambda_i)x_{jt}^S + \epsilon_{ijt}, \tag{3}$$

where $\alpha_{ij} + \lambda_i$ is the persistent preference of the consumer i for an outlet j only with non-sensitive news, and $\beta_i - \lambda_i$ is the relative preference of the consumer i for sensitive over non-sensitive news.[16] With a slight abuse of notation, we redefine utility to get rid of λ_i,

$$U_{ijt} = \alpha_{ij} + \beta_i x_{jt}^S + \epsilon_{ijt}, \tag{4}$$

where α_{ij} is the persistent preference of the consumer i for an outlet j only with non-sensitive news, and β_i is the relative preference of the consumer i for sensitive over non-sensitive news.

Ideological framing. So far, we have assumed that the only method of government control over sensitive news reporting is censorship. Apart from censorship, governments can frame the sensitive news reporting (Prat and Strömberg, 2013), making it more aligned with the government's ideology. This implies that the sensitive news reporting can have an ideological stand bias, such as supporting, opposing, or being neutral about the government.[17] We extend the model and allow news outlets to choose the ideological framing in their sensitive news reporting, $\mathrm{sl}_j \in [-1, 1]$,

$$U_{ijt} = \alpha_{ij} + (\beta_i + \gamma_i \mathrm{sl}_j)x_{jt}^S + \epsilon_{ijt}, \tag{5}$$

where γ_i captures consumer's preference for the ideology of the reporting – for instance, driven by their taste for like-minded news (Gentzkow and Shapiro, 2010).

Conscientious news consumption. Consumers' preferences for the ideological framing in the news coverage might also be driven by "conscientious" news consumption (Mullainathan and Shleifer, 2005; Xiang and Sarvary, 2007). Conscientious readers sample alter-

[15]This assumption is useful in our empirical specification of the model since we observe only the relative importance of sensitive and non-sensitive news over time, V_t^S and V_t^{NS}.

[16]Note that λ_{ij} can include any persistent difference in the non-sensitive news reporting between outlets A and B, capturing their differentiation in the non-sensitive news reporting.

[17]The literature refers to this ideological bias as 'framing and ideological stand bias" (Prat and Strömberg, 2013) and "distortion of news" (Gentzkow et al., 2016).

native ideological positions to filter out the ideological framing in the news reporting. This predicts that conscientious readers will consume more ideologically-diverse news outlets on days with a lot of sensitive news coverage, whereas consumers with a preference for like-minded news will read more similar outlets when there is a lot of sensitive news. We capture this difference in behavior by allowing for the ideological variety-seeking behavior in product choice (McAlister and Pessemier, 1982; Kim et al., 2002),

$$U_{\tau ij} = \begin{cases} \alpha_{ij} + (\beta_i + \gamma_i \mathrm{sl}_j) x^S_{jt} + \epsilon_{\tau ijt} & \text{if } \ \tau = 1, \\ \alpha_{ij} + (\beta_i + \gamma_i \mathrm{sl}_j + \rho_i |\mathrm{sl}_j - s_{i\tau}|) x^S_{jt} + \eta_i |\mathrm{sl}_j - s_{i\tau}| + \epsilon_{\tau ijt} & \text{if } \ \tau > 1, \end{cases} \tag{6}$$

where τ is the choice occasion of consumer i on day t, and $s_{i\tau}$ is the ideological framing of the outlet that was consumed on $\tau - 1$. A positive coefficient ρ_i signals an increase in the ideological "variety-seeking" of consumer i on days with a lot of sensitive news, signaling conscientious news consumption. In contrast, a negative ρ_i means that consumers read more ideologically-similar news outlets on days with a lot of sensitive news coverage, consistent with the like-minded news consumption.[18]

3 Empirical Context and Data

3.1 Online News Market Structure in Russia in 2013-2015

Despite high government control over all major TV news channels starting around the year 2000, the online news market – the second most important source of news in Russia[19] – enjoyed relative freedom up until 2013. Starting around 2013, political pressure has forced a number of top online news outlets to change their editorial and management teams, including prominent cases like changes of the editor-in-chief at RIA Novosti, a major state news agency with balanced news coverage, and lenta.ru, one of the largest independent news outlets.[20] Government control further intensified in February of 2014 with the beginning of the Ukrainian crisis – the government reacted by blocking the websites of some opposition leaders in March 2014 (bbc.com, 2014) and implementing a law to limit the foreign ownership of Russian news outlets to 20% (squirepattonboggs.com, 2014). Still, by April 2015, online

[18]This stylized model ignores any forward-looking behavior consumers might have when choosing whether to read another article within a day. We also refrain from incorporating and testing potential complementarities across outlets into the demand (Gentzkow, 2007) and supply (Xiang and Sarvary, 2007) models.

[19]In 2014, 23% named internet as their main news source, compared to 60% for TV news. By 2017, the importance of internet has increased to 32% and the importance of TV dropped to 52% (VTsIOM, 2017).

[20]Online Appendix B.1 list the changes and the corresponding outlets.

news market in Russia had a large number of independent players.

Table 1: Russian-language online news media by the type of influence in December 2014.

(i) GC	(ii) Potentially Influenced		(iii) Independent	(iv) International and Ukrainian
vz (5.17%)	lenta (6.48%)	rosbalt (1.49%)	rbc (15.3%)	**International**:
tass (5.15%)	regnum (6.4%)	echo (1.46%)	newsru (1.67%)	bbc (1.63%)
vesti (4.24%)	gazeta (3.66%)	izvestia (0.94%)	tvrain (1.47%)	svoboda (0.77%)
rg (4.22%)	utro (2.83%)	bfm (0.91%)	vedomosti (0.8%)	reuters (0.01%)
ntv (3.41%)	interfax (2.38%)	sobesednik (0.81%)	slon (0.75%)	meduza (0.00%)
aif (3.06%)	kommersant (2.38%)	polit (0.40%)	novayagazeta (0.74%)	dw (0.00%)
ria (2.52%)	kp (2.32%)	znak (0.27%)	forbes (0.68%)	**Ukrainian**:
dni (1.9%)	mk (1.93%)	ng (0.26%)	snob (0.59%)	korrespondent (1.97%)
rt (1.5%)	fontanka (1.91%)	ridus (0.15%)	the-village (0.24%)	unian (1.73%)
1tv (0.66%)	lifenews (1.86%)	trud (0.12%)	newtimes (0.10%)	liga (0.78%)

We simplify the domain names; for instance, "1tv" stands for "www.1tv.ru". Most domains have the www.*.ru structure, with some exceptions. Outlet-to-type classification is done based on the media ownership information and evidence of the indirect influence listed in the Online Appendix B. We present outlet market shares computed based on news article visits in IE Toolbar data in parentheses.

Table 1 presents the top 48 Russian-language news outlets classified by the degree and type of government influence.[21] The classification of news outlets is done according to the ownership structure (Djankov et al., 2003) and evidence of the indirect influence.[22] The first group contains outlets that are owned by the government or members of the incumbent political party, which we classify as being directly controlled by the government. The second group includes the "potentially influenced" outlets, ones that are formally independent but that can be indirectly influenced by the government – for instance, by the government's pressure on the news outlets' owners. Given the ambiguous degree of control over the "potentially influenced" outlets, we exclude them from the sensitive news classification. The third group contains independent outlets, the ones with no indication that they could be under an indirect government control. Most of these news outlets are owned either by journalists, international media companies or the government opposition. The largest independent news outlet, rbc.ru, is owned by a Russian billionaire Mikhail Prokhorov, who ran for president in the 2012 elections. The final two groups in the last column present the outlets with Russian language news coverage that are international, separating out the Ukrainian outlets.

[21] We have tried to include all significant news outlets, so the set contains even international outlets with relatively low popularity in Russia at the time.

[22] Examples of indirect influence include the removal of news articles and firing journalists under political pressure. Online Appendix B presents more detailed information on the ownership structure and evidence of the indirect influence for each news outlet.

3.2 Publication Records

We collected publications records of the 48 outlets described above for April 1, 2013 – March 31, 2015. The data are collected directly from archives on news outlet websites and from the media archives *medialogia.ru* and *public.ru*. The resulting panel contains 3.9 million news articles. For each article, we collect the title, text, URL link, and timestamp.[23] We process texts using standard techniques such as stemming and removing the stop words. Online Appendix C provides details about the data collection and processing.

Table A1 in Online Appendix D presents the summary statistics of published news articles, split by outlets' types. Twenty potentially-influenced outlets publish almost half (47.37%) of all the news in the sample, GC and independent outlets publishing the other 30% and 11%, respectively. These shares are relatively stable over time – the standard deviation of the shares of articles (computed across weeks for each type) is between 0.5 and 2.4 percentage points, and the implied coefficients of variation are between 0.03 and 0.18. The news coverage of an average GC outlet is more extensive compared to the coverage of an average independent outlet; the GC outlets publish more news articles on an average day (161 versus 80) and have more words per article (205 versus 179).

3.3 News Consumption Records

To measure news consumption, we use browsing data from the Internet Explorer (IE) Toolbar. The IE Toolbar data includes complete browsing histories for a subset of IE users.[24] IE Toolbar data contain information about each webpage consumers visited (URL), websites where consumers came from (referral URL), timestamp of the visit, number of seconds spent, browsing session ID, user ID, the browser's language, and other information. We focus the analysis on Toolbar users who specified Russian as the language of their browser.[25]

Our browsing data covers the period between November 15, 2013, and March 31, 2015, for all users with the IE language set to Russian and who navigate to news websites in our sample at least once.[26] The resulting panel consists of 284,574 users who make 20.27 million

[23]For five news outlets ("meduza," "newtimes," "ridus," "snob," "the-village"), article texts were not collected for technical reasons. While we use these outlets for the sensitive news detection (exploiting titles instead of article texts), we drop them from the descriptive analysis and demand estimation due to an unreliable measure of slant estimate.

[24]The users included in the IE Toolbar data have installed a plug-in on their IE and opted-in for the data collection. Around 75% of users who installed the plug-in opt in for the data collection.

[25]Having a browser in the Russian language indicates that the user knows Russian and is potentially in the market for Russian online news.

[26]Although IE Toolbar data were collected for several years, the unique user IDs were kept only for one

URL visits of the 48 news-outlet websites defined above.[27]

The news outlet URLs visited by the users include four types of web pages – news outlets' main pages, subdirectories, news articles, and other pages (special projects and videos) – which we infer from the URL structure and the publication records data.[28] Table A2 in Online Appendix E shows summary statistics of browsing by types of the URLs. News articles account for around a half (47.5%) of URL visits, and more than half (52.9%) of the overall time spent on these webpages. A median visit to a news article URL takes 89 seconds. The main page accounts for 20.9% of all visits, and news subdirectories and other pages account for 13% and 18.7%, respectively.

Apart from the URLs visited, we observe the referrer URLs, addresses of the webpages where a user clicked a link that sent them to a news outlet website. Table A3 in Online Appendix E summarizes main types of this referrer URLs. In a majority (53.6%) of the first website visits on a day, consumers navigate to the website directly (there is no referral recorded), with Yandex being the second most common traffic source, accounting for 21.7% of the visits.[29] Other search engines, such as Google, Bing and Rambler, account for 7.5% of the first visits. Social media websites account only for 0.34% of website landings in our sample – reflecting a low role of social media in the online news market in Russia at that time.

3.3.1 IE Toolbar Representativeness

Before we proceed with the analysis, we examine whether news consumers in the IE Toolbar data are representative of the overall population of news consumers in Russia. While the market share of the IE browser in Russia in November 2013-March 2015 was a sizable 14.4%, following Chrome (42.9%) and Firefox (18.7%) browsers (statcounter.com, 2015), there might be systematic differences in news and ideological preferences between the IE Toolbar users and general population.

---

and a half years. By the time the data collection was conducted, the earliest available browsing data with user IDs were from November 15, 2013.

[27]There is a total of 2.17 million users with the IE language set to Russian. While our main sample of users is only 13% of users with the IE browser set to Russian language, they account for 77.8% of all browsing.

[28]Online Appendix E contains details on how we classify URLs into these groups.

[29]Most of this traffic is coming from `news.yandex.ru`, a popular news aggregator run by Yandex, the main search engine in Russia. Yandex News automatically aggregate news articles across various news sources in Russian language and features top 5 current news on the main page of its search engine. Yandex lists three criteria for featuring the news – news should be widely covered in the media, of current interest, and popular – the same general principles as used by other news aggregators (e.g. Google News). Yandex provides more details on Yandex News at `https://yandex.ru/support/news/index.html` (in Russian language).

To compare the IE Toolbar users and the general population, we collected population-level data on daily visits of the most popular websites in Russia using `liveinternet.ru` (LI), a website that tracks statistics for the Russian internet. Due to the layout of the website ranking on LI, we can collect reliable records of usage over the period of time studied for the 30 most popular websites in Russia, which includes seven news websites from our sample.[30]

Online Appendix F compares browsing habits of the news consumers in the IE Toolbar data to the general population recorded by LI. Results suggest that IE Toolbar users are older, less interested in streaming and entertainment websites and more interested in news than the general population. At the same time, the overall rankings of the websites are relatively similar, with the same top 5 websites in both IE and LI datasets. The main difference in news outlets' visits in IE and LI datasets stems from a higher market share of an independent outlet rbc.ru, a business-focused news agency, in the IE data, and a lower market share of a GC outlet ria.ru, a news agency competing with rbc.ru. This difference is consistent with the anecdotes that IE users are more business-focused, and it suggests that our analysis based on the IE data might overestimate the persistent preferences for the independent outlets and underestimate the persistent preferences for the GC outlets – since rbc.ru is an independent outlet and ria.ru is a GC outlet.

Figure 1: Normalized average number of weekly visitors to the top seven news outlets, IE Toolbar and LI (population-based) data.

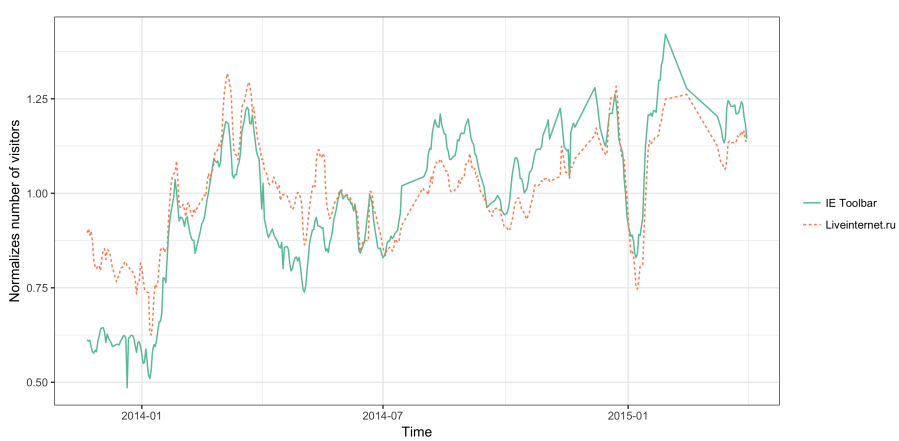

For each website and news source, the average traffic level is normalized to one, and the IE Toolbar data are corrected for the churn rate. The traffic is then averaged across the outlets.

We further compare news consumption in the IE and LI data by looking at changes in

[30] We use the digital archive "Wayback Machine" to collect historical data on website usage. The top page includes only the top 30 websites; Wayback Machine does not have frequent records for the other pages.

news website visits over time, an important variation that identifies ideological preferences of consumers. Figure 1 presents the normalized average traffic to the top seven LI news outlets based on the LI and IE Toolbar data. Changes in the news consumption in the IE Toolbar data closely track the population-level consumption in the LI data, with a correlation of 0.858. The correlations between traffic changes in the LI and IE Toolbar datasets vary from 0.52 to 0.914, with correlations of 0.914 and 0.702 for rbc.ru and ria.ru, respectively.[31] We conclude that while IE Toolbar data oversamples business-oriented news readers compared to the population of news consumers in Russia, consumption habits of the IE Toolbar users are otherwise representative of the news consumption of the population.

4 Supply of Government-Sensitive News

In this section, we use publication records data to describe the news products supplied by outlets to the market – which news topics they cover and how these topics are presented. We show that the coverage of government-sensitive news is the main dimension of product differentiation in the market, and then go on to detect main sensitive topics and describe their coverage by news outlets.

Throughout the analysis, we treat news articles as collections of words or n-grams, a "bag-of-words" approach typical in the event detection (Allan et al., 1998) and topic modeling (Blei et al., 2003) literatures. The topics of news articles are defined using named entities – words and phrases that represent information about actors (people or organizations), locations, and timing of the news, and are crucial in detecting news events (Kim et al., 2012; Hu et al., 2013).[32] As a result, we can represent the collection of topics covered by each news article or outlet with a vector of counts of named entities that appear in the texts. The usage of non-named entities by outlets pins down the ideological framing in sensitive news coverage.[33]

4.1 Detection of Government-Sensitive News

We start with examining the main dimensions of product differentiation in outlets' news coverage. For this, we extract principal components from a 21,873 by 48 matrix of normalized

[31]Figure A2 in Online Appendix F presents changes in the traffic for each of the top seven news outlets.

[32]Tracking named entities is a common approach in the information retrieval literature to extract news representations (Kumaran and Allan, 2004, 2005); named entities increase news topic coherence when getting more weight in the topic model (Krasnashchok and Jouili, 2018).

[33]We discuss the implications and extensions to this approach in Section 4.3 below, and validate our sensitive news classification in Online Appendix H.

counts of usage of popular named entities by outlets.[34] We detect named entities using a simple algorithm that searches for capitalized names in texts, finding 21,873 unigrams and 16,917 bigrams of named entities that appear more than 200 times in our data.[35]

Figure A3 in Online Appendix G summarizes scores of the first two principal components across the outlets. Results strongly suggest that the difference in sensitive news coverage is the main differentiation dimension in the online news in Russia. The first principal component almost perfectly separates out GC and independent outlets – 16 out of 24 outlets with scores above the median are either independent, international or Ukrainian, and only one is GC. In contrast, 9 out of 24 outlets below the median score are GC and 13 are potentially influenced. The second principal component differentiates outlets on the volume of coverage of news about the events in the Ukraine, which is evident from the Ukrainian outlets having the top scores. The GC outlets are clustered closely together, suggesting that the product differentiation among them is limited.

We now use the difference in coverage of the GC and independent outlets to identify government-sensitive news topics. While named entities with the highest loadings in the first principal component already suggest some sensitive news topics – for instance, the top-20 words include “Navalny”, a prominent opposition leader, and “Roskomnadzor”, a censorship agency in Russia – they also include incidental and general words (“Wikipedia”, “Putin”, “Yandex”, “Spotify”) due to the nature of PCA that pools the co-occuring words together. To recover only named entities that are indicative of sensitive news topics, we develop a simple classification algorithm that looks for unigrams and bigrams of named entities that are systematically overused or underused by *all* or *most* GC outlets. The algorithm ranks news outlets by normalized usage of named entities, computes the average rank difference between the GC and independent outlets, and compares whether this rank difference could occur by chance. While simple, our algorithm outperforms all common alternative predictive classification methods – such as the comparison of named entity usage shares, TF-IDFs, partial least squares used by Gentzkow and Shapiro (2010), article-level Lasso regression (Tibshirani, 1996), and article-level naive Bayes classification – when compared to a manual detection of sensitive topics done by three independent research assistants.[36] Online

[34]Counts are normalized by the overall usage of named entities by a news outlet to correct for differences in outlets’ size. A related strategy was used by (Qin et al., 2018) who use PCA to extract the most important variation from nine proxies for political bias in government-owned newspapers in China.

[35]Online Appendix C provides more details on named-entity detection. We keep only relatively common words to make sure that they refer to an important topic. The threshold of 200 times is chosen arbitrarily. Local changes in the threshold do not affect the results.

[36]Online Appendix H.2 presents the results of this validation. We also note that our data is too large for

Appendix H.1 provide a detailed exposition and discussion of the classification algorithm.

Table 2: Top 20 unigrams and bigrams of named entities underused by GC news outlets.

Underused named entity, English translation:	Short contextual information about the named entity	Rank Difference $\Delta \text{Rank}_v^{Ind-Gov}$
Rotenberg	A businessman, reportedly a close ally of V. Putin	-28.9
Roskomnadzor	The federal agency exercising media control	-28.2
Khodorkovsky	A former oligarch and political prisoner	-28.1
Alexey Navalny	An opposition politician	-26.9
Navalny	An opposition politician	-26.5
Lebedev	An associate of Khodorkovsky	-25.5
Sechin	CEO of Rosneft, reportedly a close ally of V. Putin	-25.5
Kudrin	The Head of the Committee of Civil Initiatives	-25.3
Kosenko	A political activist, arrested at the opposition rally	-24.9
Sergei Guriev	An economist, interrogated about the Yukos case	-24.9
Bolotnaya	A place of a large opposition rally	-24.8
Prokhorov	A businessman and presidential candidate	-24.8
Bukovsky	A political activist	-24.7
Marat Gelman	A gallerist, fired for a political exposition	-24.7
Gennady Timchenko	A businessman, reportedly a close ally of V. Putin	-24.3
Sakharova	A place of a large opposition rally	-24.3
Svetlana Davydova	A civilian investigated for treason	-24.3
Ketchum	A PR agency working for Russian government	-24
Mikhail Khodorkovsky	A former oligarch and political prisoner	-24
Gelman	A gallerist, fired for a political exposition	-23.9

Applying the classification algorithm to common named entities, we detect 208 named entities (out of 38,790 common unigrams and bigrams) that are underused by the GC outlets to a degree that could not happen by chance. After screening out named entities that relate to the profession of journalism – which may show up in the list simply due to news source citations – we get 128 named entities that define topics censored by the GC outlets.[37] To provide examples of such censored topics, Table 2 presents a list of 20 named entities that are the most underused by the GC outlets. Words in this list are related to political opposition (for instance, "Khodorkovsky" and "Navalny"), political protests ("Bolotnaya" and "Sakharova"), alleged corruption ("Rotenberg" and "Gennady Timchenko") and media

more sophisticated classification methods, such as Gentzkow et al. (2019).

[37]We use three independent research assistants to find named entities related to journalism, as well as ambiguous named entities. Online Appendix H.2 provides more details on the procedure, with Tables A6–A10 presenting the full list of underused words. We additionally validate the 128-word threshold by using the sensitivity scores from research assistants, with the results presented in Figure A6.

control ("Roskomnadzor" and "Ketchum"). We classify articles that mention one of the top 128 underused named entities as related to "POC" news, where POC stands for topics related to political protest, opposition and corruption. Examining the differences in non-named entities usage by outlets in the POC news articles, we do not find evidence of ideological framing – the algorithm picks up a handful of differences in word usage, most commonly due to misclassified named entities.[38]

The second major sensitive news topic that we find in the data is the Ukraine crisis of 2013-2014 and a subsequent conflict between Russia and Ukraine. The Ukraine crisis has been widely covered in the Russian news – the share of news articles that mention "Ukraine" jump from 2-3% to 20-30% with the beginning of the crisis – was reported to be heavily slanted by the GC news outlets (themoscowtimes.com, 2014; time.com, 2014), and is by far the largest topic (outside of the POC news) labeled by independent research assistants as government-sensitive.[39] We label any news article that mentions Ukraine as news about the Ukraine crisis,[40] and compare the usage of non-named-entity words by the GC and Ukrainian news outlets to detect ideological framing. The ideological framing is prominent; we find 101 and 27 words over- and underused by the GC outlets compared to the Ukrainian outlets, respectively.[41] The GC news outlets report that Russia has had a "reunion" (rank 1) with Crimea, and that the Ukraine military conducts a "punitive" (rank 3), "russophobic" (rank 10), and "anti-Russian" (rank 18) operation in the Eastern Ukraine. The Ukrainian outlets report an "annexation" (rank 3) and "occupation" (rank 10) of Crimea by Russia via a "pseudo-referendum" (rank 4), and label the operation in the Eastern Ukraine "anti-terroristic" (rank 5). We screen out incidental words with the help of three independent research assistants, leaving us with 26 pro-Russia-slanted words and 7 pro-Ukraine-slanted words.[42] The selected words validate well; they are remarkably consistent with the Ukraine crisis propaganda narrative described by fact-checking websites (stopfake.org, 2014), and the implied ideological positions of news outlets – which we define below – are robust to

[38]Online Appendix H.3 provides details on this analysis.

[39]Online Appendix H.4 presents the details and results of this classification.

[40]Figure A8 in Online Appendix H.5 presents the share of news articles that contain the word "Ukraine" over time. We use this classification to keep the definition broad and ensure that we do not miss any articles related to the conflict. Alternatively, we can define news articles as being about the Ukraine crisis using 23 sensitive named entities detected by the research assistants, which we list in Table A12 in Online Appendix H.4. Our results are robust to using this alternative classification. The correlation in the volume of the Ukraine-crisis news based on these two measures is 91.5%.

[41]The corpus of non-named-entity words in the Ukraine crisis news articles contains 34,395 words. Figure A9 in Online Appendix H.6 presents the distribution of rank score differences.

[42]We report the full list of under- and over-used words, as well as results of word classification, in Tables A13 and A14 in the Online Appendix H.6.

alternative definitions of the Ukraine-crisis news articles and to the manual classification of news outlets' ideological positions.[43] We label any article that mentions one of these words as having a pro-Russia or pro-Ukraine ideological framing.

4.2 Coverage of Government-Sensitive News

We use the identified POC and Ukraine-crisis news articles, and articles with the ideological framing in Ukraine crisis reporting, to characterize the coverage of sensitive news.

First, we compute the relative importance of the POC and Ukraine-crisis news on a given day. For this, for topic $l \in \{POC, Ukr\}$ we compute the share of news articles about l across all outlets on a given day, $V_t^l = \frac{\sum_j N_{tj}^l}{\sum_{l'} \sum_j N_{tj}^{l'}}$, where N_{tj}^l is the number of topic l articles from outlet j on day t. On an average day, 9.56% of news articles cover the POC news, and 19.13% – news about the Ukraine crisis.[44] There are large differences in the share of coverage across days, with the standard deviation of V_t^l of 3.75 and 11.3 percentage points, respectively.

Figure 2: Differences in the POC news reporting by the GC and independent news outlets.

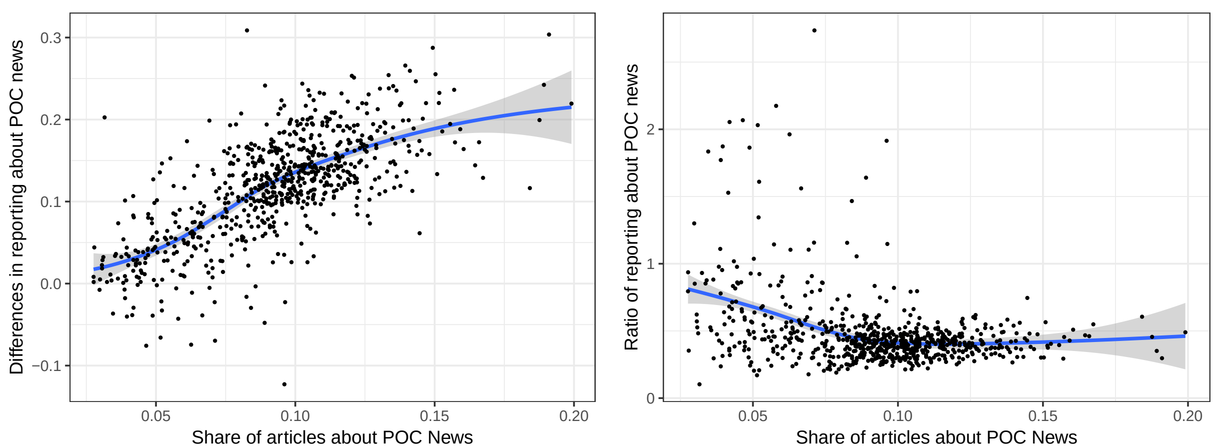

(a) Difference $x_{t,Ind}^{POC} - x_{t,GC}^{POC}$ (b) Ratio $x_{t,GC}^{POC}/x_{t,Ind}^{POC}$

Subfigure (a) plots the relationship of $x_{t,Ind}^{POC} - x_{t,GC}^{POC}$ and V_t^{POC}, and Subfigure (b) plots the relationship of $x_{t,GC}^{POC}/x_{t,Ind}^{POC}$ and V_t^{POC}. The blue line corresponds to the fitted local polynomial regression.

Second, we measure how much news about topic l each outlet reports by computing the share of reporting of outlet j on topic l on day t, $x_{tj}^l = \frac{N_{tj}^l}{\sum_{l'} N_{tj}^{l'}}$. This allows us to infer outlets' reporting positions, including the degree to which they report censored POC news. Consistent with the model predictions, the difference in coverage of the GC and independent

[43]The manual classification is described in Online Appendix H.7.
[44]A high share of news about the Ukraine crisis reflects the prominence of this topic in our data period.

news outlets, $x^{POC}_{t,Ind} - x^{POC}_{t,GC}$,[45] increases in the relative importance of POC news on day t, V^{POC}_t, as captured by Figure 2a. However, Figure 2b shows that the ratio of reporting on POC news by the GC and independent outlets is stable over time and is uncorrelated with V^{POC}_t, suggesting stable ideological positions of the outlets. On an average day, GC outlets report around 42% of POC news covered by the independent outlets. We confirm that outlets do not systematically change their reporting positions on the POC and Ukraine-crisis news – outlet fixed effects explain 30.05% and 61.69% of the variation in reporting positions of outlets, x^l_{tj}/V^l_t, while adding an interaction of outlet fixed effects and V^l_t increases the explained variation only by 2.75 and 6.21 percentage points for POC and Ukraine-crisis news, respectively. Given such limited reaction of the news outlets to V^l_t for $l \in \{POC, Ukr\}$, we approximate outlets' share of reporting of sensitive news by their average share of news reporting across t, $\bar{x}^l_j = \frac{\sum_t N^l_{tj}}{\sum_t \sum_{l'} N^{l'}_{tj}}$. Figure A11 in Online Appendix I presents the resulting ideological positions of the news outlets.

Finally, we measure the ideological framing of the Ukraine crisis news by outlets. For this, for each outlet we compute the difference in shares of articles with a pro-Russia and pro-Ukraine slant in the Ukraine crisis coverage, $\frac{N^{\text{pro-Russia}}_{tj}}{N^{\text{Ukr}}_{tj}} - \frac{N^{\text{pro-Ukraine}}_{tj}}{N^{\text{Ukr}}_{tj}}$, where $N^{\text{pro-Russia}}_{tj}$ and $N^{\text{pro-Ukraine}}_{tj}$ is the number of articles with the pro-Russia and pro-Ukraine slant in the Ukraine crisis coverage. These ideological positions are uncorrelated with V^{Ukr}_t – outlet fixed effects and their interactions with V^{Ukr}_t explain 40.61% and additional 0.47 percentage point of the variation, respectively[46] – allowing us to measure the outlet's ideological position as the overall difference in the outlet's usage of pro-Russia and pro-Ukraine slant. For this, we compute $\text{sl}^{\text{pro-Russia}}_j = \frac{\sum_t N^{\text{pro-Russia}}_{tj}}{\sum_t N^{\text{Ukr}}_{tj}}$ and $\text{sl}^{\text{pro-Ukraine}}_j = \frac{\sum_t N^{\text{pro-Ukraine}}_{tj}}{\sum_t N^{\text{Ukr}}_{tj}}$, the average shares of using pro-Russia and pro-Ukraine slant for each outlet, normalize these shares to have zero mean and a unit standard deviation, and define the ideological positions of news outlets as the difference in these normalized measures, $\text{sl}_j = \text{sl}^{\text{pro-Russia,n}}_j - \text{sl}^{\text{pro-Ukraine,n}}_j$.

Figure 3 presents the resulting Ukraine-crisis news reporting and ideological framing positions of outlets. Results line up well with our expectations; international outlets closely resemble Ukrainian outlets in their ideological slant, and independent outlets have more "neutral" ideological positions. Interestingly, some potentially influenced outlets closely resemble independent outlets, while others closely resemble GC outlets. We check the identities of these outlets and confirm that positions match anecdotal evidence on the strength of indirect influence; for instance, a potentially influenced outlet with the most pro-Russia slant

[45] Where $x^{POC}_{t,type}$ is a share of news reporting by a particular outlet type.

[46] To measure ideological framing in the Ukraine crisis news, we use data since the beginning of the crisis.

Figure 3: News outlets' ideological positions and share of reporting about the Ukraine-crisis news.

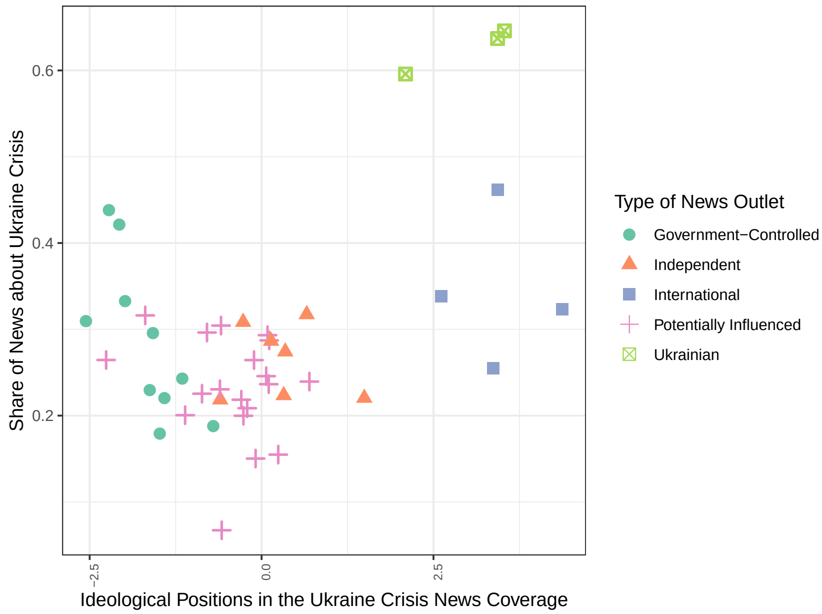

Each dot represents a position of a news outlet, with shapes and colors of the dots corresponding to the outlets' types.

is "lifenews," a website close to the Russian security services (themoscowtimes.com, 2013), while the least pro-Russia slant comes from "echo", a website known for its independent coverage despite being owned by Gazprom media.[47] Further, the resulting ideological positions of outlets validate well in a manual classification – the text-based measure closely tracks the manually-labeled measure with a correlation of 0.839.[48]

4.3 Discussion of Sensitive News Detection

We pause for an additional discussion of our method of sensitive news classification.

First, while our measures of POC and Ukraine-crisis news are based on a small subset of named entities, they proxy for a larger set of sensitive news topics. For instance, the named entity "Navalny" – a prominent political activist investigating corruption – might come up in any news related to opposition and corruption. Because of this, our measures of the volume of sensitive news and ideological positions of news outlets are robust to local changes in the

[47]Figures A12 and A13 in Online Appendix I present the ideological positions and reporting of news outlets with the corresponding outlet labels; Figure A14 presents a joint distribution of the POC news reporting (censorship) and ideological framing of the Ukraine-crisis news (propaganda) across the news outlets.

[48]Online Appendix H.7 presents further details of this validation procedure.

number of words that describe sensitive news.[49]

Second, our classification method can be extended to incorporate word inter-dependence. For instance, one can detect news topics using word co-occurrence in the news articles (e.g. Blei et al., 2003; Mikolov et al., 2013) and then run our classification algorithm on an outlet-news topic matrix. The downside of such extension is that topic detection methods will group informative and incidental words together, increasing the noise in the measure of government-sensitive news topics.

Finally, we note that our classification of sensitive news is based on a comparison of topics published in the news market and does not account for a potential self-censorship by the independent outlets. Schimpfossl and Yablokov (2014) discusses the reasons for self-censorship in the TV news market in Russia; similar logic can be applied to the online news market. Our measure of censorship is thus closest to "state censorship" in the classification of Crabtree et al. (2015).

5 Model-Free Evidence

Before estimating the empirical version of the model from Section 2, we examine model-free evidence on the direction of ideological preferences of readers and the sources of demand for the GC news outlets.

5.1 Changes in Market Shares with Sensitive News

The key variation that identifies ideological preferences of news readers in our model are changes in outlets' consumption in response to shifts in the relative importance of sensitive news in the market on day t, V_t^{POC} and V_t^{Ukr}. The intuition is that consumers will be more likely to switch to outlets with preferred ideological positions in topic l's coverage on days with high V_t^l, increasing preferred outlets' market shares. We examine the relationship between outlets' market shares and the relative importance of sensitive news by running a separate log-log regression for each outlet,

$$log(\text{share}_{jt}) = b_{0j} + b_j^{POC} log(V_t^{POC}) + b_j^{Ukr} log(V_t^{Ukr}) + b_j^{Plac} log(V_t^{Plac}) + d_j Z_t + \xi_{jt}, \qquad (7)$$

[49]For instance, if we manipulate the definition of POC news by moving around the cutoff from a more (89 censored named entities) to a less (400 censored named entities) restrictive measure, the implied measures of the POC news volume and reporting are almost unchanged – the average correlations in different measures of V_t^{POC} and $\bar{x}_j^{POC}$ are 91% and 97%, respectively.

where share_{jt} is the market share of outlet j on day t, and $d_j Z_t$ are outlet-specific week and weekday fixed effects in the main specification. The market share, share_{jt}, is defined as the share of visits that outlet j attracts on day t, where consumers can make at most one visit of an outlet per day and the outside option is defined as browsing but not visiting any news outlets.[50] Figure A15 in Online Appendix J.1 plots the evolution of outlet types' market shares over time. Apart from V_t^{POC} and V_t^{Ukr}, we include a placebo news topic variable V_t^{Plac} in regressions 7, defined as a share of news articles that mention one of the 233 not sensitive named entities.[51] Since this measure includes random words related to different topics, we do not expect V_t^{Plac} to have any systematic effect on the market shares.

Figure 4 visualizes the estimates of b_j^{POC} and b_j^{Ukr} for 42 news outlets in our sample.[52] Each point represents an estimate of b_j^{POC} or b_j^{Ukr} for one of the outlets, with blue circles for positive and red squares for negative estimates. The size of points is proportional to the absolute value of the estimate, and bold borders represent statistically significant estimates.[53]

Estimates in Figure 4a reveal that outlets with more reporting about POC news get disproportionally higher market shares on days with high V_t^{POC}; seven out of nine outlets with the highest $\bar{x}_j^{POC}$ get a statistically significant increase in their market shares on days with a high $log(V_t^{POC})$, and the other two estimates are marginally significant at 10% level (p-values of 0.106 and 0.118). The average slope coefficient for these nine outlets is 0.219, meaning that a 1% increase in V_t^{POC} leads to a 0.22% increase in these outlets' market shares.[54] In contrast, only 8 out of the other 32 outlets get significant increases in market shares when $log(V_t^{POC})$ increases, with an average b_j^{POC} estimate of 0.044.

Figure 4b visualizes the estimates of b_j^{Ukr}. News outlets that report more news about the Ukraine crisis (upper part of the plot), $\bar{x}_j^{Ukr}$, and have a more pro-Ukraine ideological position (right part), sl_j, get the highest increases in their market shares on days with a high

[50]Thus, we define news consumption of an outlet j on day t by consumer i as a visit to any page on the news outlet j. Our results are robust to alternative definitions of news consumption, such as a visit to at least one news article on outlet j, a visit to any page but the main directory, a visit to at least 5 pages on website j, and spending at least 2 and 3 minutes on website j. For observations of outlet-day pairs with zero market shares, we assign the lowest observed non-zero market share to avoid the problem of taking a logarithm of zero when estimating the regression 7.

[51]The 233 named entities were identified by research assistants as not sensitive for the government. These named entities come from 724 randomly selected named entities used in manual classification.

[52]We exclude five news outlets for which we do not have information about the text of the articles, and one news outlet (`dw.de/ru`) for which we have few (10) news consumption occasions.

[53]Tables A16–A17 and Figures A16–A20 in Online Appendix J.2 present estimates of b_j^{POC} and b_j^{Ukr}. Standard errors are heteroskedasticity and autocorrelation consistent.

[54]The slope coefficient of rbc.ru, the largest independent news outlet, is a significant 0.124 (s.e. of 0.023), well-aligned with the rest of the independent outlets. Five out of seven independent news outlets get significant increases in market shares with an increase in V_t^{POC}.

Figure 4: Estimates of correlations in the market shares of the outlets and relative importance of sensitive news news, V_t^{POC} and V_t^{Ukr}.

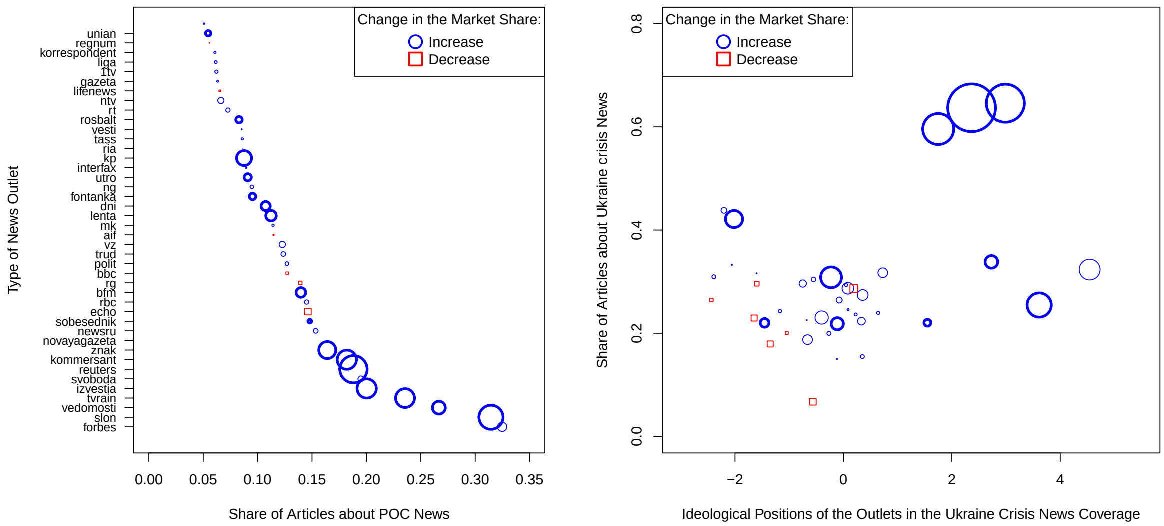

(a) POC News, $\hat{b}_j^{POC}$ (b) Ukraine-crisis News, $\hat{b}_j^{Ukr}$

Each point represents a news outlet. The size of each point represents the effect of V_t^l, $l \in \{POC, Ukr\}$, on the market shares of news outlets, measured in percentages (b_j^l coefficients of regression 7). The blue color corresponds to positive coefficients, and the red color – to negative coefficients. The bold borders of the points correspond to the 5% significance of the change in the market share.

V_t^{Ukr}. In particular, six out of seven news outlets with the most pro-Ukraine slant get a statistically significant increase in their market shares, and the remaining one is marginally significant at 5% level (p-value of 0.056). The average slope coefficient for these seven outlets is 0.267, meaning that a 1% increase in V_t^{POC} leads to a 0.27% increase in these outlets' market shares. In contrast, only 4 out of the rest of 35 outlets get significant increases in market shares on days with a high $log(V_t^{Ukr})$, with an average b_j^{Ukr} estimate of 0.041.

In contrast, we do not find any systematic correlations of outlets' market shares with the placebo topic, $log(V_t^{Plac})$; only 2 out of 42 outlets have significantly higher market shares on days with a high $log(V_t^{POC})$, and another 2 have significantly lower market shares.[55] Overall, all placebo tests confirm that the relationships we describe in Figure 4 are not incidental.[56]

We confirm that news outlets with less pro-government ideological positions gain significantly higher market shares on days with high V_t^{POC} and V_t^{Ukr} in two more ways. First,

[55] Table A18 and Figures A21 and A22 in Online Appendix J.3 present the estimates.

[56] We further check whether the estimates of b_j^{POC} and b_j^{Ukr} are correlated with the outlets' share of reporting about the placebo news topic, $\bar{x}_j^{Plac}$; both relationships are statistically insignificant.

in Online Appendix J.4 we show that the effects are not driven by small outlets by running the analysis on the outlet type level; we get similar patterns of changes in the market shares. Second, in Online Appendix J.5 we show results of a joint regression with all outlets, estimating the effect of an interaction of $log(V_t^l)$ and $\bar{x}_j^l$ for $l \in \{POC, Ukr\}$, as well as $log(V_t^{Ukr})$ and sl_j. The interaction coefficients confirm that outlets with higher $\bar{x}_j^{POC}$ gain extra market shares on days with a high V_t^{POC}, and market shares of outlets with higher $\bar{x}_j^{Ukr}$ and more pro-Ukraine sl_j gain extra market shares on days with a high V_t^{Ukr}.

The estimates of b_j^{POC} and b_j^{Ukr} can be interpreted as causal effects of sensitive news volume on outlets' market shares if the conditional independence assumption (CIA) holds, $\xi_{jt} \perp log(V_t^l)|Z_t \; \forall j, l = \{POC, Ukr\}$. While the CIA is plausible given that $log(V_t^l)$ is determined by the amount of realized sensitive news events on day t – a process hardly controlled by the market participants – changes in outlets' market shares are only suggestive about readers' ideological preferences. Apart from a tastes for a less pro-government ideology in the news, higher market shares of independent outlets might be due to consumer selection – for instance, if consumers with high persistent preferences for independent outlets also have a taste for sensitive news topics. The same pattern can be explained by conscientious news consumption. In Section 6, we estimate a structural model of demand for news that accounts for these alternative explanations, consumer heterogeneity and conscientious consumption.

5.2 Description of Online News Consumption Patterns

To understand the potential sources of readers' persistent preferences for the GC outlets, we compare how people arrive to and navigate around the GC and independent outlets.

First, consumers are more likely to arrive to the GC outlets from third-party websites. Table A22 in Online Appendix K presents shares of referral traffic – arrivals through clicks on links on third-party websites – for different types of news outlets. Direct navigation plays a lower role for the GC outlets (50.68%), especially compared to the independent ones (56.28%). In particular, the GC outlets get more than a quarter of their traffic from Yandex (25.57%), compared to only 15.5% for the independent outlets. They also tend to cross-refer each other more than other types of news outlets – 68.1% of the cross-referred traffic that lands on the GC outlets is from another GC outlet, while only 7.7% of independent outlets' cross-referred traffic is from GC outlets. These numbers are in stark contrast to consumers' outlet switching patterns, with a visit to a GC outlet preceding 33.3% and 30.4% of visits of other GC and independent outlets, respectively. Online Appendix K presents more details on these results and a discussion of the cross-referral and switching patterns.

Table 3: First visit shares by types of web pages.

				Outlet Type		
	All	GC	Potentially Influenced	Independent	International	Ukrainian
Referral From				Traffic Share (%)		
Main Page	19.41	10.94	21.53	28.29	14.16	24.01
News Articles	55.94	58.28	57.01	50.54	51.63	60.25
News Subdirectories	7.80	3.49	12.29	5.62	15.03	9.12
Other	16.85	27.28	9.18	15.55	19.18	6.62

The shares of traffic are computed conditional on the outlet type. Results are for the first visit to a news outlet on a given day. All columns sum up to 100%.

Second, consistent with the lack of direct navigation to the GC outlets, a large share of consumers arrive directly to their news article pages and "other pages", such as special projects and videos. Table 3 reports that news articles and other pages capture 58.28% and 27.28% of all arrivals on the GC outlets, in contrast to 50.54% and 15.55% for the independent outlets, respectively. Only 10.94% of arrivals on the GC outlets are on the main pages, compared to 28.29% for independent outlets. The high share of traffic navigating to "other pages" of the GC outlets reflects the availability of video content – 3 out of 10 GC news outlets in our sample are major federal TV channels, streaming their content online.[57]

Once a consumer is on the news outlet's website, our primary interest is in visits to news article URLs – the main place of news consumption. Table 4 presents the share of referrers for news article visits. Similar to the case of arrivals to news outlets, direct navigation (including visits from the same outlets' pages) and Yandex are the two primary sources of traffic, with Yandex still being disproportionally more important for the GC outlets (21.18% of article visits).Focusing more on the recurring pageviews on the same outlet, a large share of news articles are consumed after a different news article was read, with this number being 27.45% for GC and 16.85% for independent outlets. In contrast, 5 times fewer (5.44% and 3.4%) news articles are consumed after "other pages" on GC and independent outlets. This reflects that only a minority – around 10%, both for the GC and independent outlets – of consumers who land on the "other page" go on to read any news articles on the same website.

A large share of news article visits after another article was read suggests potential inertia in news consumption – that a consumer is more likely to continue reading news

[57]The top 2 "other pages" of the GC outlets are live steams of Channel One (`www.1tv.ru`) and Russia24 (`www.vesti.ru`), the two main federal TV channels in Russia. The other 3 out of 5 top "other pages" of the GC outlets are the reruns of the TV programs on the website of Channel One.

Table 4: Summary of referrals and arrivals to news article URLs, by outlet types.

	Outlet Type				
	GC	Potentially Influenced	Independent	International	Ukrainian
Referral From	Traffic Share (%)				
Direct & From This News Outlet	67.46	68.40	68.31	74.58	82.01
– Arrival From News Article	27.45	23.48	16.85	27.67	28.10
– Arrival From Other Pages	5.44	3.15	3.40	8.13	1.93
Yandex	21.18	16.53	10.77	8.71	6.44
Other Search Engines (not Yandex)	3.97	4.70	2.86	8.33	1.88
Other Aggregators (not Yandex)	1.31	3.15	0.15	0.13	0.29
Other News Outlets	1.16	0.72	0.63	1.96	0.48
Social Media	0.32	0.19	0.22	0.22	0.26
Other Websites	4.61	6.29	17.07	6.06	8.64

The shares of traffic are computed conditional on the outlet type. Results are for all news article visits. Arrival from news articles and "other pages" is recorded as consumers visiting these pages just before landing on the news article page.

on the website once she lands there. Most of the time (70%), consumers read only one article upon a website visit, and on half (51%) of the days read more than one article overall on any website. But, if a consumer continues to read news that day, with a 54.2% probability an average consumer reads the next article on the same website. In contrast, the probability that an average consumer visits the same outlet on two subsequent news consumption days is 18.3% – 35.9 percentage points lower – suggesting substantial within-day inertia in consumption. The data suggest that across-day inertia is much weaker – the probability that an average consumer makes the first visit to the same news outlet on two random (instead of consecutive) days is 17.6%, only 0.7 percentage points lower compared to back-to-back consumption days. Finally, the probability that two random consumers read the same news outlet on two random consumption occasions is only 5% – confirming the importance of the heterogeneity in consumers' outlet preferences.

High within-outlet inertia in news consumption suggests that once a consumer lands on the outlet's website, she is likely to get exposed to various published news content, even on topics different from what she has arrived for. We confirm it in Table A24 in Online Appendix L – while readers are more likely to continue reading articles on the same topic, a substantial share of readers switch from topic to topic within an outlet.[58] Interestingly,

[58]Similar results hold separately for the GC and independent outlet readers. Further, on days with a lot (above the median) of the outlet's non-sensitive news article visits, this outlet also gets more POC and Ukraine-crisis articles visits – an average of 48.2% and 45.2% more visits across outlets' types, respectively.

transitions from topic to topic are relatively similar – the difference of 6 percentage points, on average – if readers stay on the same website or switch outlets. This suggests that consumers tend to read similar news topics across outlets and do not have specialized outlets to read sensitive or non-sensitive news, instead choosing a preferred outlet as a content bundle.

We confirm that there is little outlet-topic specialization by decomposing the variation in consumers' news article visit counts across outlets, days, and topics. The interactions of consumer-outlet fixed effects explain 18.2% of the variation in article visits, while further interacting consumer-outlet fixed effects with article topics (threefold increase) increases adjusted R-squared only 5.1 percentage points. Online Appendix L further shows that consumption shares of news articles topics closely follow the publication shares of these topics by news outlets, consistent with the notion that news outlets have power over which news to expose consumers to. Overall, these results support the demand model where consumers make choices on the outlet rather than on the outlet-topic level.

The above patterns suggest several mechanisms behind readers' persistent outlet preferences and how GC outlets they can influence sensitive news consumption. First, a large share of the GC outlets' visits comes from the referral traffic, particularly Yandex. These additional referrals can be due to the GC outlets' investments into the overall quality of their websites, additional non-sensitive news coverage, or "other pages" – all of which could increase the GC outlets' positions in search engines and news aggregators search and ranking results.[59] Second, some consumers transition to news articles from "other pages" (video content and special projects), meaning that GC outlets' investments into such content could increase sensitive news consumption on their websites. Third, the GC outlets' coverage of the non-sensitive news directly increases their readership and could expose readers to the sensitive news coverage on this outlet, due to the within-outlet inertia in consumption. Fourth, the data suggest some degree of across-day inertia in the news outlet consumption, capturing brand loyalty or habit persistence – meaning that consumers might get exposed to the sensitive news coverage due to choice inertia, after reading non-sensitive news on previous days. Finally, the persistent preferences of GC outlets might be driven by the overall quality of these outlets, their websites, or news articles.[60] We examine the relative importance of these alternative mechanisms behind the persistent preferences of GC outlets in Section 7.1.1.

[59] There is little evidence that the government could directly manipulate the number of news referrals coming from Yandex to GC outlets at that time. The direct pressure on Yandex has intensified in years 2016-2017 when a new law made aggregators responsible for the news content they publish (rsf.org, 2016).

[60] For instance, the availability of video content or pictures in the top of news articles' webpages.

6 Empirical Specification

6.1 Empirical Model

Our empirical model adjusts the stylized model from Section 2 to the empirical setting of the Russian online news market.

There are three news topics covered by the outlets: non-sensitive, POC and Ukraine-crisis news. The news event realizations are driven by a stochastic process not controlled by market participants. The realized relative importance of news topics across days is captured by the overall share of news on that topic that day, V_t^{POC} and V_t^{Ukr}.

News outlet $j \in J$ makes three reporting decisions – what share of the POC and Ukraine-crisis news to report, $\bar{x}_j^{POC}$ and $\bar{x}_j^{Ukr}$, and which ideological framing to take in the Ukraine-crisis news reporting, sl_j. The importance of these ideological positions for consumer choice on day t is shifted by V_t^{POC} and V_t^{Ukr}. Finally, outlets can also choose to differentiate in terms of their persistent features, such as which non-sensitive news to report and how much money to invest in quality of the news reporting or website, among others.

There are I consumers in the market. On days when they spend time browsing online, they might choose to consume one or more news outlets, or choose not to read the news (outside option). Following Gentzkow and Shapiro (2015), we assume that consumers can read at most one news outlet at a time.[61] This setting naturally lends itself to a discrete choice model, where on a consumption occasion τ a consumer chooses an outlet j that she has not read on the previous choice occasions $1, \ldots, \tau - 1$ on this day. We define the news consumption of an outlet j as navigation to at least one news article on the outlet's j website by consumer i on day t.[62] Thus, on each day t, consumers can have at most J news consumption occasions. Unless a consumer has read all J news outlets on day t, on the last τ of the day a consumer chooses an outside option of not consuming the remaining news outlets. At each choice occasion τ on day t, a consumer chooses an outlet j such that $u_{ijt\tau} \geq u_{ij't\tau}\ \forall j' \in \{0, \ldots, J\} : j' \neq j$. We denote consumers' choices as $y_{i\tau t}$.

Adapting consumer utility defined in equation 6 to this empirical context, we get

$$\begin{aligned} u_{ijt\tau} = \alpha_{ij} + V_t^{Ukr}\bar{x}_j^{Ukr}\left(\beta_i^{Ukr} + \mathrm{sl}_j\gamma_i + |\mathrm{sl}_j - s_{i\tau}|(\tau > 1)\rho_i\right) + \\ + V_t^{POC}\bar{x}_j^{POC}\beta_i^{POC} + |\mathrm{sl}_j - s_{i\tau}|(\tau > 1)\eta_i + \mathrm{state}_{it\tau}\pi_i + \epsilon_{ijt\tau}. \end{aligned} \tag{8}$$

[61] It is impractical for people to read multiple news outlets simultaneously.

[62] This discrete-choice specification ignores the intensity of news consumption within the outlet; our results are robust to redefining a consumption occasion of an outlet to a visit to a news article on a given day, allowing for multiple articles read within a day.

The notation in this model closely follows equation 6 from Section 2 – α_{ij} denotes a persistent preference of consumer i for outlet j, β_i^{POC} and β_i^{Ukr} are relative preferences of consumer i for the POC and Ukraine-crisis news over the non-sensitive news, γ_i is a preference for the ideological framing of the Ukraine-crisis news, and ρ_i is the variety-seeking parameter that signals conscientious or like-minded Ukraine-crisis news consumption. The reduced-form parameter η_i captures the baseline variety-seeking on days with no Ukraine-crisis news. The only new term compared to equation 6, $\text{state}_{it\tau}$, is an indicator variable that captures whether a consumer i has already visited j on day t. Since, by construction, consumers never revisit the same news outlet on day t, the variable $\text{state}_{it\tau}$ serves a technical purpose of restricting the actual choice set of consumers (with a highly negative value of π_i).

6.1.1 Discussion of the Assumptions

We pause to discuss several assumptions underlying the empirical model.

First, we assume that consumers know the relative importance of news topics on day t, V_t^{POC} and V_t^{Ukr}, and the reporting and ideological positions of the news outlets, $\bar{x}_j^{POC}$, $\bar{x}_j^{Ukr}$, and sl_j. We believe that these are reasonable assumptions in our context. We define consumption as visits to news articles, meaning that consumers have some exposure to the overall set of topics that have happened on day t, either on the main page of news outlets or on news aggregators. Our estimation focuses only on frequent news consumers, who are more likely to know the average reporting positions. If these assumptions are violated, we likely overestimate the role of the persistent preferences of consumers and underestimate the role of the preferences for news reporting and ideological framing of sensitive news.

Second, we assume that ideological preferences of consumers are stable over time. If this assumption is violated, our estimates would capture only average preferences of consumers. In particular, the estimates of persistent preferences, α_{ij}, capture any long-term effects of the ideology of news outlets, as well as any unobserved differences in the sensitive news coverage other than the coverage of the POC and Ukraine-crisis news.

Third, we follow the stylized model and define the consumers' tastes for sensitive news topics as a coefficient on $V_t^{Sens}\bar{x}_j^{Sens}$. An alternative model specification is to separate out the effect of V_t^{Sens}, the relative importance of sensitive news on this day, and the effect of $\bar{x}_j^{Sens}$, the share of news on the outlets' website devoted to this topic. While separating out these effects is appealing, such alternative specification makes it hard to identify consumers' tastes for sensitive news – the model needs to estimate not only coefficients on V_t^{Sens} and $\bar{x}_j^{Sens}$, but also a correlation term between them, increasing the requirements on the number of choices observed per consumer. Separately, such alternative specification deviates from the stylized model, making it hard to interpret the nature of the estimated coefficients.

Finally, our model does not allow for the interactions between the volume of news coverage of a topic and the quality of this topic. Any horizontal or vertical differences across the news outlets are captured by the persistent preferences of consumers, α_{ij}.

6.2 Estimation and Identification

We estimate the model using only data from frequent news consumers – those who consume news at least 10 days in our data. These consumers are more likely to know the ideological positions of news outlets, and, since they make more outlet choices, the data provides more information about their preferences. There are 54,905 such news consumers in our sample responsible for 4,822,667 consumption occasions (outlet-day visits).[63] On almost half (48.6%) of the consumption days, news readers in the selected sample have only one news consumption occasion. However, conditional on having more than one consumption occasion on day t, news readers navigate to an average of 2.71 news outlets. For computational reasons, we estimate the model on a random sample of 10,000 of such frequent news consumers; all of our conclusions replicate if we re-run the model with a new random sample of consumers. As in Section 5.1, we focus on the top 42 online news outlets in the sample.

We estimate the distribution of $\theta_i = \{\alpha_{ij}, \beta_i^{Ukr}, \beta_i^{POC}, \gamma_i, \rho_i, \eta_i, \pi_i\}$ using a Bayesian hierarchical model (Rossi et al., 2005). We assume that $\epsilon_{ijt\tau} \sim i.i.d.$ EV(0,1), leading to a standard logistic model, but allow for a flexible heterogeneity in consumer preferences. The probability that consumer i chooses outlet j on day t on the consumption occasion τ is

$$\pi(y_{it\tau} = j|\theta_i) = \frac{exp(u_{ijt\tau}(\theta_i))}{1 + \sum_{j'} exp(u_{ij't\tau}(\theta_i))}, \tag{9}$$

implying the likelihood of θ_i observing a sequence of choices y_i of

$$L(\theta_i|y_i) = \prod_t \prod_\tau \prod_j \pi(y_{it\tau} = j|\theta_i)^{I(y_{it\tau}=j)}. \tag{10}$$

We use a normal distribution on the first-stage prior of θ_i, a normal prior over its mean

[63] Out of 214,375 news consumers who visit a news article page at least once over the sample period. While these consumers correspond only to 24.5% of news readers in the market, they account for 92.2% of all the news articles read in the data sample period.

and an inverse Wishart prior over the covariance matrix:[64]

$$\begin{aligned} \theta_i &\sim N(\mu, \Sigma), \\ \mu &\sim N(\bar{\mu}, \Sigma \otimes a_\mu^{-1}), \\ \Sigma &\sim IW(\nu_\Sigma, \Psi_\Sigma). \end{aligned} \tag{11}$$

The flexibility of this specification comes through an unrestricted covariance matrix Σ, which allows for correlations across all outlet fixed effects and other consumer preferences. This flexibility captures alternative explanations for changes in outlet market shares discussed at the end of Section 5.1. The cost of this flexibility is that we cannot account for the potential within-day correlations of the error terms across the consumers due to computational costs; as the result, the sampling procedure might underestimate the uncertainty around the posterior point estimates. We estimate the model using a computationally costly MCMC hybrid sampling procedure; Online Appendix M provides the details.

The identification of consumer preferences, $\{\alpha_{ij}, \beta_i, \gamma_i, \rho_i\}$, relies on exogenous shifts in the relative importance of sensitive news realizations, V_t^{POC} and V_t^{Ukr}. Such shifts change the importance of the outlets' reporting positions, $\bar{x}_j^{POC}$, $\bar{x}_j^{Ukr}$, and sl_j, identifying β_i and γ_i. The distribution of the persistent preferences of consumers, α_j, is identified from the expected outlet choices when V_t^{POC} and V_t^{Ukr} are zero. The distribution of ρ, an ideological variety-seeking preference of consumers, is identified from changes in the similarity of sl_j between two subsequently consumed outlets in response to changes in V_t^{Ukr}.

7 Results

7.1 Preference Estimates

Table 5 reports the distribution of posterior point estimates of consumer preferences. The top six rows summarize the distributions of persistent preferences, α_{ij}, averaged by outlet types and demeaned to facilitate the comparison. The estimates reveal that an average consumer has the highest persistent preference for the GC news outlets ($\hat{E}(\hat{\alpha}_{GC} - \hat{\bar{\alpha}}) = 1.1033$), followed by the independent ($\hat{E}(\hat{\alpha}_{Ind} - \hat{\bar{\alpha}}) = 0.129$) and potentially influenced ($\hat{E}(\hat{\alpha}_{Inf} - \hat{\bar{\alpha}} = 0.128)$) news outlets. The heterogeneity in consumer preferences is substantial; for instance, the standard deviation of preferences for the independent outlets, $\hat{\alpha}_{Ind} - \hat{\bar{\alpha}}$, is 0.592, resulting in 41.11% of people to prefer an average outlet over the independent news outlets. In contrast,

[64]We set standard tuning parameters following Rossi et al. (2005) and Rossi (2014). Given the amount of data in our likelihood function, the results are almost unaffected by changing the tuning parameters.

the vast majority of consumers have higher persistent preferences for the GC news outlets over an average outlet (97.5%) and the independent news outlets (87.95%).

Table 5: Posterior point estimates of consumer preferences.

	Mean	S.D.	% of users > 0
$\hat{\bar{\alpha}}$	-5.872	1.101	0
	(0.016)	(0.012)	–
$\hat{\alpha}_{GC} - \hat{\bar{\alpha}}$	1.103	0.547	97.5
	(0.016)	(0.01)	(0.24)
$\hat{\alpha}_{Inf} - \hat{\bar{\alpha}}$	0.128	0.273	68.65
	(0.008)	(0.005)	(1.03)
$\hat{\alpha}_{Ind} - \hat{\bar{\alpha}}$	0.129	0.592	58.89
	(0.015)	(0.01)	(1.07)
$\hat{\alpha}_{Int} - \hat{\bar{\alpha}}$	-2.253	1.015	1.67
	(0.096)	(0.048)	(0.25)
$\hat{\alpha}_{Ukr} - \hat{\bar{\alpha}}$	-2.532	2.542	14.77
	(0.05)	(0.034)	(0.35)
$\hat{\beta}^{POC}$	0.028	0.146	58.85
	(0.002)	(0.002)	(0.7)
$\hat{\beta}^{Ukr}$	0.094	0.218	67.2
	(0.003)	(0.002)	(0.56)
$\hat{\gamma}$	0.016	0.133	54.98
	(0.002)	(0.002)	(0.75)
$\hat{\rho}$	-0.052	0.182	39.9
	(0.004)	(0.003)	(0.86)

The posterior standard deviation estimates are in parentheses.

Results are drastically different once we shift our attention to the preferences of consumers for sensitive news coverage. An average consumer prefers POC ($\hat{E}(\hat{\beta}^{POC}) = 0.028$) and Ukraine-crisis ($\hat{E}(\hat{\beta}^{Ukr}) = 0.094$) news over the non-sensitive news, and a less anti-Ukraine slant in Ukraine-crisis news coverage ($\hat{E}(\hat{\gamma}) = 0.016$). This implies that an average consumer has a distaste for the censorship (less reporting of the POC news) and the ideological framing (more pro-government slant) of the GC news outlets. Such preferences hold for the majority (58.85% and 54.98%) of consumers in the online news market in Russia. However, the utility differences that consumers get from reading sensitive news on independent media outlets relative to GC outlets are small compared to differences in persistent preferences, even on days with a high volume (2 standard deviations above the average) of sensitive news events. Table A25 in Online Appendix N computes and discusses these utility differences.

Finally, the average consumer prefers like-minded news ($\hat{E}(\hat{\rho}) = -0.052$) in the Ukraine-crisis coverage, and around 60% of consumers have a taste for like-minded Ukraine-crisis news. The rest of 40% of consumers prefer more ideologically-diverse news on days with more Ukraine-crisis events, and these tend to be consumers with a high average $\hat{\bar{\alpha}}$ and those who have a higher preference for independent and international outlets.

Table 6: Simulated market shares for different levels of POC and Ukraine-crisis news.

Outlet Types		Market Shares				
		(1)	(2)	(3)	(4)	(5)
		Volume of Sensitive News				
	V_t^{POC}:	0	Mean	Mean + 2 S.D.	0	0
	V_t^{Ukr}:	0	0	0	Mean	Mean + 2 S.D.
share_{Gov}		14.33	14.58	14.67	15.27	16.64
		(0.08)	(0.08)	(0.08)	(0.06)	(0.1)
share_{Inf}		16.11	16.37	16.53	17.23	17.79
		(0.05)	(0.06)	(0.07)	(0.04)	(0.05)
share_{Ind}		10.79	11.57	12.72	11.71	12.27
		(0.06)	(0.06)	(0.07)	(0.05)	(0.06)
share_{Int}		0.89	0.93	0.98	1.11	1.63
		(0.01)	(0.01)	(0.02)	(0.01)	(0.02)
share_{Ukr}		1.32	1.32	1.29	1.68	3.14
		(0.03)	(0.03)	(0.03)	(0.01)	(0.03)
$\text{share}_{\text{Outside}}$		56.57	55.23	53.8	53	48.53
		(0.14)	(0.14)	(0.18)	(0.1)	(0.15)
Market Share Ratios:						
$\text{share}_{Gov}/\text{share}_{Inf}$		0.89	0.89	0.89	0.89	0.94
		(0.01)	(0.01)	(0.01)	(0)	(0.01)
$\text{share}_{Gov}/\text{share}_{Ind}$		1.33	1.26	1.15	1.3	1.36
		(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
$\text{share}_{Gov}/\text{share}_{Int}$		16.16	15.66	14.9	13.72	10.2
		(0.23)	(0.25)	(0.28)	(0.14)	(0.14)
$\text{share}_{Gov}/\text{share}_{Ukr}$		10.82	11.09	11.34	9.09	5.3
		(0.22)	(0.23)	(0.24)	(0.06)	(0.07)

The market shares are percentages of the entire market. The posterior standard deviation estimates are in parentheses.

To get a better understanding of the magnitudes of consumer preferences, Table 6 computes the expected market shares of news outlets under different volumes of POC and Ukraine-crisis news. On days without sensitive news (Column 1), $V_t^{POC} = 0$ and $V_t^{Ukr} = 0$, GC outlets get 14.33% of the market, while independent outlets are getting only 10.79%. The implied market share ratio is 1.33, and the difference reflects the difference in persistent outlet preferences of consumers. As the volume of POC news (V_t^{POC}) increases (Columns 2 and 3), the market share of independent outlets grows more compared to the market share of GC outlets, reflecting consumers' preference for more POC news coverage. As a result, the ratio of GC to independent outlets' market share decreases to 1.26 and 1.15 on days with an average and 2 standard deviations above the average volumes of POC news, respectively. However, even on days with a lot of POC news, consumers are more likely to navigate to the GC outlets over independent outlets, showing the importance of GC outlets' persistent preferences. Similarly, on days with more Ukraine-crisis news (Columns 4 and 5), the market shares of the international and Ukrainian outlets grow much faster than the market shares of other outlets, reflecting the preference for less pro-government slant in the Ukraine-crisis news. Also, the market share of GC outlets grows slightly faster than the market share of independent outlets, due to their higher coverage of the Ukraine-crisis news.

7.1.1 The Nature of Persistent Preferences

The preference estimates confirm the suggestive evidence from Section 5 that persistent preferences drive consumption of GC outlets. We now provide correlational evidence on the relative importance of alternative mechanisms behind the persistent outlet preferences, and discuss how these mechanisms influence sensitive news consumption.

First, we examine the degree to which referral traffic drives persistent outlet preferences. To understand the relative importance of referral traffic of the GC outlets, we re-estimate the demand model excluding consumption sessions of GC outlets that were started due to different referral sources, and then simulate the market shares of GC outlets based on the resulting persistent preferences.[65] Horizontal bars 1-4 in Figure 5 summarize the resulting market shares of GC outlets. Under the current persistent preferences (bar 1), GC outlets get 14.33% of the market share on days with no sensitive news. If we remove all indirect traffic of GC outlets (bar 2), their market share driven by persistent preferences decreases to 7.79%, a 6.54 percentage points reduction. The traffic coming from Yandex accounts for most – 4.12 percentage points – of this reduction, with the 10.21% market share of GC outlets driven by the persistent preferences of consumers without Yandex's referrals (bar 3). Finally, referrals from other GC outlets play a much smaller role – if we remove this traffic (bar 4), the implied market share of GC outlets decreases only by 0.17 percentage points.

Second, we measure the relative importance of "other pages" and different news content in persistent preferences for GC outlets. For this, we re-estimate the model excluding con-

[65] Tables A26-A28 in Online Appendix O presents model estimates based on the data with different excluded referral traffic of the GC outlets.

Figure 5: Simulated market shares of GC outlets excluding different potential sources of persistent preferences.

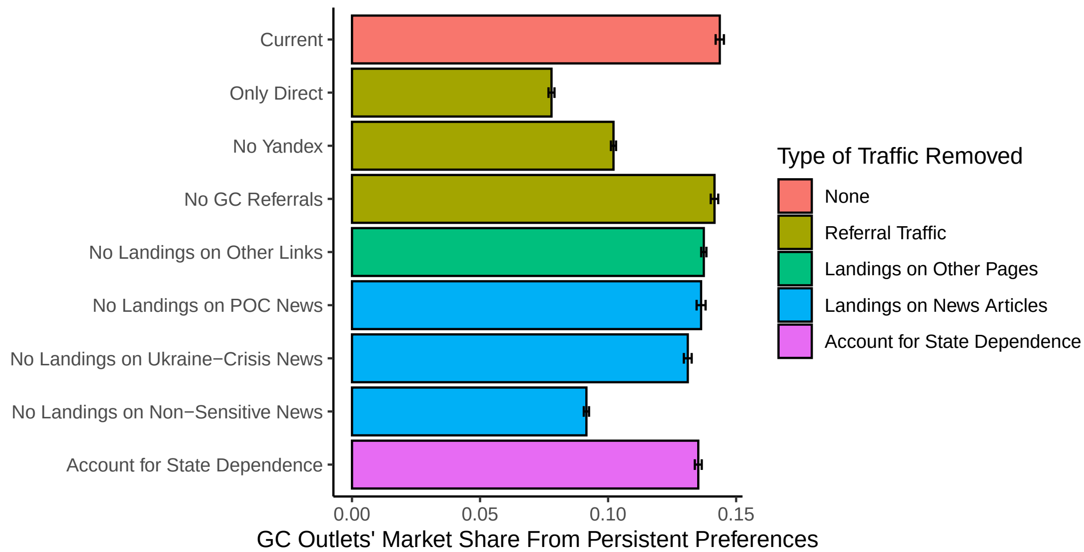

Each bar represents the estimation results with different GC outlet arrivals excluded. We simulate the market shares for days with no sensitive news, $V_t^{POC} = V_t^{Ukr} = 0$, meaning that market shares are solely driven by the persistent preferences of consumers. Error bars correspond to two standard deviations of the MCMC draws.

sumption sessions of GC outlets that started with landings on "other pages" or on news articles of different topics, and again simulate the market shares of GC outlets based on the resulting persistent preferences.[66] Horizontal bars 5-8 in Figure 5 summarize the resulting market shares. If we exclude landings on "other pages" of GC outlets (bar 5), their market share driven by the persistent preferences is 13.74%, 0.59 percentage points lower than the market share driven by the current persistent preferences. This effect is smaller compared to the effect of excluding landings on the POC (bar 6), Ukraine-crisis (bar 7), and non-sensitive (bar 8) news articles, with the 13.63%, 13.11% and 9.15% market shares of GC outlets driven by the persistent preferences in those scenarios, respectively. These results suggest that news articles about non-sensitive topics bring a large volume of consumers to GC outlets and drive a large part (5.18 percentage points) of their market shares coming from persistent preferences. In Online Appendix P, we show that readers with a high persistent preference for GC over independent outlets tend to read more non-sensitive news, including news about celebrities, sports, and international events, while consumers with a low preference for GC outlets are more likely to read articles related to Russian and Ukrainian politics and law.

[66] Tables A29-A32 in Online Appendix O presents the resulting model estimates.

Third, we examine the role of choice inertia in the persistent preferences of consumers for GC outlets. For this, we re-estimate the model adding a Markov state dependence variable – an indicator variable taking a value "1" if this GC outlet was visited on the previous day with any news consumption.[67] Online Appendix Q writes out the model specification and presents the estimation results.[68] We then simulate the market share of GC outlets driven by the persistent preferences of consumers excluding the accumulated brand loyalty mechanism. The resulting market share is 13.52%, 0.81 percentage points lower that the market share driven by the current persistent preferences of GC outlets. We conclude that while the (short-term) brand loyalty has a sizable effect on GC outlets' market share, it is smaller than the effect of referral websites (like Yandex) or non-sensitive news content.

Finally, persistent outlet preferences might be formed via long-term habit formation or reflect fixed characteristics of news outlets. While we cannot exclude these fixed outlet features and re-estimate the model, we leverage the correlation structure in α_{ij} estimates and describe the outlet features that have the highest predictive power of these correlations. If an outlet feature strongly predicts correlations in consumer tastes, it is more likely to play a role in the preference formation. To run this analysis, we compute correlations in the demeaned consumers' outlet tastes, $\alpha_{ij} - \bar{\alpha}_i$, for each pair of outlets $j, j' : j \neq j'$, resulting in 861 unique correlation estimates. We then regress these correlation estimates of news outlet pairs, $\hat{\text{cor}}(\alpha_{ij} - \bar{\alpha}_i, \alpha_{ij'} - \bar{\alpha}_i) : \; j \neq j'$, on the absolute value of the differences in outlets' characteristics z, $|z_j - z_{j'}|$, where z includes various features of outlets that reflect their overall quality (the outlet's average $\bar{\alpha}_j$, the number and length of news articles) and ideology (e.g. share of articles about sensitive news). Results show that consumers have more similar persistent tastes for outlets with more similar overall quality, $\bar{\alpha}_j$, and closer ideological positions, such as sl_j and $\bar{x}_j^{POC}$, suggesting that the overall quality and fixed outlet features might play a role in forming the persistent outlet preferences. Online Appendix R presents more details and results of this analysis.

Overall, results suggest that high persistent preferences of consumers for GC outlets are driven by the referrals from Yandex and consumers' landings on news articles with the non-sensitive content, and less so driven by the brand loyalty, sensitive news coverage, and "other pages" visits. Given that consumers exhibit the within-outlet inertia and tend to read news

[67]This formulation is commonly used in the literature that measures brand loyalty (Dubé et al., 2010; Bronnenberg and Dubé, 2017).

[68]We handle the initial conditions problem (Heckman, 1981) by estimating the bounds on the state dependence coefficient as proposed by Simonov et al. (2020). The difference in the upper and lower bounds on the state dependence estimate – presented in the last row of Tables A37 and A38 in Online Appendix Q.2 – is statistically insignificant, showing that our setting does not suffer from the initial conditions problem.

topics in the same proportions as covered by the outlet – as we discuss in Section 5.2 – these high persistent preferences allow GC outlets to expose consumers to their preferred ideological framing of sensitive news. This ideological framing includes the slanted Ukraine-crisis coverage and a restricted exposure to the censored POC news. We next examine the degree to which the ideological positions of GC outlets affect news consumption and measure the government's media power over consumers.

7.2 Counterfactuals

We now assess the degree to which GC outlets benefit from strong persistent tastes of consumers, as well as the cost of their sub-optimal ideological positions.

To understand the impact of government control on the outlets' market shares, we simulate market outcomes in counterfactual scenarios with different ideological positions of news outlets. The government controls its outlets with censorship – which decreases the share of news reporting about POC news (low $\bar{x}_j^{POC}$ $\forall$ $j \in$ GC) – and ideological framing in the Ukraine-crisis news (low sl_j $\forall$ $j \in$ GC). To simulate scenarios without government control, we adjust the reporting of GC news outlets on sensitive issues to match the independent outlets. More specifically, we adjust $\bar{x}_j^{POC*} = \bar{x}_j^{POC} * (\bar{x}_{Ind}^{POC}/\bar{x}_{GC}^{POC})$ and $\mathrm{sl}_j^* = \mathrm{sl}_j - \mathrm{sl}_{GC} + \mathrm{sl}_{Ind}$ for all $j \in$ GC, where $\bar{x}_{GC}^{POC}$ and $\bar{x}_{Ind}^{POC}$ represent average reporting positions of the GC and independent outlets about POC news, and sl_{GC} and sl_{Ind} represent average ideological framing positions of the GC and independent outlets in the Ukraine-crisis news. This way, we treat average ideological positions of the independent news outlets as "unbiased" and interpret simulation results as short-term reactions of the market to changes in government control.[69]

The simulated market shares under alternative levels of government control are presented in Table 7.[70] Column (1) reports the predicted market shares under the current ideological positions of the GC outlets; in this regime, they get a market share of 15.56%.

Column (2) reports the predicted market shares with adjusted $\bar{x}_j^{POC*}$ and sl_j^* for the GC outlets – a case when the government does not exercise direct control of the news market through ownership (Gehlbach and Sonin, 2014). The market share of GC outlets increases

[69] In the long run, we would expect changes both on the supply side, such as product differentiation decisions, and on the demand side, such as changes in persistent preferences. Further, when changing the reporting and ideological positions of the GC outlets, we assume that they retain their persistent preferences, which in part might driven by the high quality of their non-sensitive news coverage.

[70] We average the resulting market shares across the expected realizations of V_t^{POC} and V_t^{Ukr}. In order to speed up the counterfactual simulation, we approximate news realizations V_t^{POC} and V_t^{Ukr} by the centers of 20 clusters of these variables and simulate one choice occasion per consumer per day. Standard k-means clustering algorithm is applied to cluster the observed V_t^{POC} and V_t^{Ukr}.

Table 7: Simulated market shares for different levels of government control and persistent preferences for the GC news outlets.

Outlet Types	Market Shares					
	(1)	(2)	(3)	(4)	(5)	(6)
	Level Of Governments' Control					
	Actual	No control			More control	Low
		Direct	Indirect	Both		α_{GC}
share$_{Gov}$	15.56	17.94	15.23	17.33	15.72	7.11
	(0.04)	(0.12)	(0.04)	(0.11)	(0.04)	(0.02)
share$_{Inf}$	17.43	16.82	18.92	18	17.64	19.68
	(0.04)	(0.04)	(0.08)	(0.05)	(0.04)	(0.05)
share$_{Ind}$	12.53	12.01	12.17	11.79	11.73	13.63
	(0.03)	(0.04)	(0.04)	(0.05)	(0.06)	(0.04)
share$_{Int}$	1.2	1.13	1.15	1.09	1.22	1.32
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
share$_{Ukr}$	1.8	1.76	1.78	1.75	1.81	1.92
	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)
share$_{Outside}$	51.47	50.34	50.76	50.03	51.88	56.34
	(0.06)	(0.1)	(0.07)	(0.12)	(0.07)	(0.07)

The market shares are percentages of the entire market. The posterior standard deviation estimates are in parentheses.

from the current 15.56% to 17.94%, a 2.38 percentage points (15.3%) increase. More than half (1.13 p.p.) of this increase comes from the outside option; the rest is mainly covered by the potentially-influenced and independent outlets. This change in GC outlets' market share is smaller but comparable to the magnitudes reported by Qin et al. (2018), who find that a one-standard-deviation increase in political bias of a newspaper in China is associated with a 33% decrease in this newspaper's advertising revenues, largely determined by readership.

Similarly, in column (3) we compute the cost of government control for the potentially-influenced outlets – we adjust their average $\bar{x}_j^{POC*}$ and sl_j^* to match independent outlets. The potentially-influenced outlets are not owned but still partially controlled by the government, representing indirect control (Gehlbach and Sonin, 2014). If they were to report like the independent outlets, their expected market share would increase by 1.49 percentage points to 18.92%, an 8.5% increase. Column (4) simulates the market shares under no direct and indirect control and confirms the above results, although with smaller benefits for the GC and potentially-influenced outlets – reflecting higher competition between outlets when all of them have similar (unbiased) ideological positions.

In column (5) we examine the reverse scenario of more indirect control, a case when the independent news outlets change their average ideological positions to the ones of the potentially influenced outlets.[71] In this case, the market share of independent news decreases from 12.53% to 11.73%, a 0.8 percentage points reduction.

We assess the amount of money that news outlets lose due to inferior ideological positions driven by the government control in a simple back-of-the-envelope calculation. Almost none of the Russian online news outlets have a paywall, with display advertising being the primary source of their revenue. In 2014, the total expenditure on display advertising in Russia was 19.1 billion rubles (akarussia.ru, 2014), or around $318 million using a 60 rubles per dollar exchange rate (exchange rates.org, 2014). Even if we assume that the online news market gets all the display advertising revenues – a generous best-case scenario for the news outlets – a 1 percentage point increase in the market share of outlets converts to $318 * 0.01 / (1-0.515) = $6.56 million of display advertising revenue. This implies that GC outlets lose at most 2.38 * $6.56 = $15.6 million of display advertising revenue per year due to government control, and independent outlets would lose 0.8 * $6.56 = $5.25 million if they became controlled. For comparison, government subsidies to mass media in Russia in 2015 were $1.21 billion (rbc.ru, 2015) – several orders of magnitude more than the loss of online outlets.[72]

In the last column of Table 7, we examine the effect of GC outlets' high persistent preferences on their current market share. For this, we adjust the average persistent preference of consumers for GC outlets to matches the average persistent preference of the independent outlets – $\alpha^*_{ij} = \alpha_{ij} - \hat{\bar{\alpha}}_{GC} + \hat{\bar{\alpha}}_{Ind} \; \forall \; j \in \text{GC}$ – and compute the expected market shares under these new preferences. Under this lower persistent preference regime, the market share of the GC news outlets decreases by 8.45 percentage points, or 54.3%. Comparing this to the results in Column (2), the high persistent preferences for the GC outlets is around $8.45/2.38 \approx 3.5$ times more important in generating their market share than removing the government control of the news – highlighting the possibility for the government to expose consumers to their version of the sensitive news coverage. While we cannot causally separate out the source of high persistent preference of the GC outlets, the correlational evidence in Section 7.1.1 suggests that such high increase in market share is driven by a high share of

[71] This is perhaps a more feasible scenario given the events of 2016-2017 – by the middle of 2016, several independent news outlets had to change their ownership due to a new law (TrustLaw, 2016), and rbc, one of the top online news outlets in Russia, had to change the editorial team due to the government pressure (bbc.com, 2016) as well as its ownership later in 2017 (forbes.ru, 2017).

[72] This difference in magnitudes is partially due to the fact that Russian advertising market is relatively small, meaning that while the pro-government bias in the news has a comparable effect on readership as in China (Qin et al., 2018), the implied values of advertising revenue lost are relatively small.

referral traffic of GC outlets and their coverage of non-sensitive topics.

7.2.1 Online Media Power of the Government

While market shares and the corresponding display advertising revenues are important for the GC news outlets, the main reason for government's investments into the GC outlets is to capture the attention of news readers and potentially persuade them to support the government. To understand the ability of the government to influence readers in the online news market in Russia, we compute the degree of media power (Prat, 2018) that GC outlets have, and the role of high GC outlets' persistent preferences in this media power. Since we do not have access to cross-platform news consumption data like Kennedy and Prat (2019), we focus solely on the online news market and compute the degree of online media power.

We extend the definition of the attention share in Prat (2018) to our discrete choice model set-up. The attention share of consumer i on day t to an outlet j is

$$\text{attention share}_{ijt} = \frac{\Pr(y_{it} = j)}{1 - \Pr(y_{it} = 0)}, \tag{12}$$

where 0 is an outside option of not reading the news. Aggregating this across days, consumers, and outlet types, we get the overall attention share of the GC news outlets

$$\text{attention share}_{GC} = \sum_{j \in \text{GC}} \sum_{t=1}^{T} \sum_{i=1}^{I} \frac{\text{attention share}_{ijt}}{I * T}. \tag{13}$$

Using this last definition, we now compute the online attention shares and media power of GC news outlets under alternative persistent preferences. Column (1) of Table 8 presents the attention share estimates of GC outlets under the current persistent preferences – the resulting attention share is 33.8%. This attention share corresponds to the upper bound of 0.51 on governments' media power, meaning that the government is able to swing the 24.5-75.5% vote share election into a draw.[73]

Column (2) of Table 8 presents GC outlets' attention shares and media power under a lower level of persistent preferences of consumers – as if the average preference of consumers for GC outlets was the same as for independent outlets, a similar scenario to the results in column (6) of Table 7. In this case, the attention share of GC outlets would be 17.92%, a 15.88 percentage points reduction compared to the current scenario. The media power of

[73] The upper bound is computed based on the "worst-case scenario" assumptions – that readers are naive and do not understand that the GC news outlets are trying to persuade them (Prat, 2018).

Table 8: Attention shares and market power of GC outlets under alternative persistent preferences

	Alternative Persistent Preferences for GC Outlets							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Current	Low α	No Referrals		No	No Article Arrivals		
		(Like indep.)	Only Direct	No Yandex	Other Links	POC	Ukraine Crisis	Not sens.
att. share$_{GC}$	33.8	17.92	18.51	24.29	32.72	32.13	28.61	24.71
	(0.08)	(0.05)	(0.08)	(0.06)	(0.07)	(0.08)	(0.07)	(0.06)
Media power$_{GC}$	0.511	0.218	0.227	0.321	0.486	0.473	0.401	0.328
	(0.002)	(0.001)	(0.001)	(0.001)	(0.001)	(0.002)	(0.001)	(0.001)

the government is also substantially reduced and is only 0.218, meaning that the government can swing the 39-61% vote share election into a draw. This result once again confirms the importance of high persistent preferences of GC outlets in capturing readers' attention.

In Columns (3)-(8), we examine the relative importance of alternative mechanisms behind the GC outlets' persistent preferences in generating their attention share and media power. For this, we recompute the attention shares using the preference estimates from Section 5.2 – where we have re-estimate the model excluding alternative sources of GC outlets' traffic.[74] Consistent with our previous findings, indirect traffic (column 3) – and Yandex in particular (column 4) – play a very important role in increasing the media power of the government; the attention share of GC outlets would be 18.51% and 24.29% if the persistent preferences did not benefit from the indirect traffic overall and Yandex in particular, respectively. The availability of other pages, POC news and Ukraine-crisis news increase the GC outlets' attention share less – by 1.108, 1.67, and 5.19 percentage points, respectively (columns 5-7). The availability of the non-sensitive news substantially increases the attention share and media power of the GC outlets – the media power grows from 0.328 to 0.511.

Finally, we confirm that high persistent preferences for GC outlets allow them to capture the attention of consumers who prefer the sensitive news coverage of independent outlets, even on days with a lot of sensitive news events. The control over this group of consumers is particularly interesting since they are more likely to oppose the incumbent government in voting and engage in protests. Even on days with a lot – 2 standard deviations above average – of POC news, GC outlets have an attention share of 31.5% among consumers who prefer more POC news coverage, $\hat{\beta}_i^{POC} > 0$. The high persistent preferences of these consumers for GC outlets drives 15.15 percentage points of this attention share. Similarly, on days with a

[74]The estimates that we use here are presented in Tables A26–A32 in Online Appendix O.

lot of Ukraine-crisis news and among consumers who prefer the anti-government ideological framing in Ukraine-crisis news ($\hat{\gamma}_i > 0$), the attention share of GC outlets is 29.2%, with 14.3 percentage points driven by the high persistent preferences of consumers for GC outlets. These results show that GC outlets have a substantial power even over the ideological diet of consumers who prefer the sensitive news coverage of independent outlets.

8 Conclusions

In the new era of broad and unrestricted access to information, it is critical to understand whether governments can influence public opinion online. In this paper, we show that many consumers in the Russian online news market read the GC news outlets even though they have a distaste for the pro-government ideological coverage. Instead, the main source of demand for the GC news outlets comes from the outlet-level tastes of consumers, which are largely driven by third-party referrals and the availability of non-sensitive news content on GC outlets' websites. Once on the website, consumers are more likely to keep reading other news articles from this outlet, including the coverage of politically sensitive news. Thus, the outlet-level drivers of consumption enable the government to impose its sensitive news coverage on the readers and potentially persuade them to change their ideological preferences.

Our results on the ideological preferences of consumers should be extrapolated with care. We study online news consumption. While it is a rapidly growing segment of news readership, both in Russia and abroad (PewResearchCenter, 2017), TV is still the main news source for most news consumers in Russia (VTsIOM, 2017). Given this, it is unclear whether our estimates differ from the public opinion surveys in Russia at the time (Economist, 2016) due to a bias in the stated preferences in the surveys (Kuechler, 1998) or due to online consumers being more critical of the government. Among online news consumers, our sample of IE users should be more pro-government (based on older demographics, republic.ru (2012)) and have a lower persistent preferences for GC outlet (based on over-sampling of rbc.ru and under-sampling of ria.ru) than an average online news reader in Russia, reinforcing our conclusions.

While our results provide strong suggestive evidence that third-party referrals (such as Yandex) and non-sensitive news content are the primary mechanisms behind GC outlets' persistent preferences – that GC outlets use to expose consumers to the sensitive news coverage – our empirical strategy does not allow to causally measure the effect of these mechanisms on sensitive news consumption. Causally pinning down the mechanism behind the formation of persistent outlet preferences remains an important area for future research.

References

Adena, M., R. Enikolopov, M. Petrova, V. Santarosa, and E. Zhuravskaya (2015). Radio and the rise of the nazis in prewar germany. *The Quarterly Journal of Economics 130*(4), 1885–1939.

akarussia.ru (2014). Advertising spending by channel in 2014 in russia [in russian]. `http://www.akarussia.ru/knowledge/market_size/id5354` [Acccessed: December 16, 2019].

Allan, J., J. Carbonell, G. Doddington, J. Yamron, Y. Yang, et al. (1998). Topic detection and tracking pilot study: Final report. In *Proceedings of the DARPA broadcast news transcription and understanding workshop*, Volume 1998, pp. 194–218. Citeseer.

Angelucci, C. and J. Cagé (2019). Newspapers in times of low advertising revenues. *American Economic Journal: Microeconomics 11*(3), 319–64.

Angelucci, C., J. Cagé, and M. Sinkinson (2020). Media competition and news diets.

Athey, S., M. M. Mobius, and J. Pál (2017). The impact of aggregators on internet news consumption. *Working Paper*.

Bai, J., M. Golosov, N. Qian, and Y. Kai (2015). Understanding the influence of government controlled media: Evidence from air pollution in china. *Working Paper*.

bbc.com (2014). Websites of putin critics blocked in russia. `http://www.bbc.com/news/technology-26578264` [Acccessed: December 16, 2019].

bbc.com (2016). Management of rbc is fired [in russian]. `https://www.bbc.com/russian/news/2016/05/160513_rbc_badanin` [Acccessed: December 16, 2019].

Berry, S. T. and J. Waldfogel (2001). Do mergers increase product variety? evidence from radio broadcasting. *The Quarterly Journal of Economics 116*(3), 1009–1025.

Besley, T. and A. Prat (2006). Handcuffs for the grabbing hand? media capture and government accountability. *The American Economic Review 96*(3), 720–736.

Blei, D. M., A. Y. Ng, and M. I. Jordan (2003). Latent dirichlet allocation. *Journal of machine Learning research 3*(Jan), 993–1022.

Bronnenberg, B. J. and J.-P. Dubé (2017). The formation of consumer brand preferences. *Annual Review of Economics 9*, 353–382.

Cagé, J. (2020). Media competition, information provision and political participation: Evidence from french local newspapers and elections, 1944–2014. *Journal of Public Economics 185*, 104077.

Cagé, J., N. Hervé, and M.-L. Viaud (2020). The production of information in an online world. *The Review of Economic Studies 87*(5), 2126–2164.

Calzada, J. and R. Gil (2020). What do news aggregators do? evidence from google news in spain and germany. *Marketing Science 39*(1), 134–167.

Chen, Y. and D. Y. Yang (2019). The impact of media censorship: 1984 or brave new world? *American Economic Review 109*(6), 2294–2332.

Crabtree, C., C. J. Fariss, and H. L. Kern (2015). Truth replaced by silence: A field experiment on private censorship in russia. *Working Paper*.

Dahl, G. and S. DellaVigna (2009). Does movie violence increase violent crime? *The Quarterly Journal of Economics 124*(2), 677–734.

DellaVigna, S., R. Enikolopov, V. Mironova, M. Petrova, and E. Zhuravskaya (2014). Cross-border media and nationalism: Evidence from serbian radio in croatia. *American Economic Journal: Applied Economics 6*(3), 103–32.

DellaVigna, S. and M. Gentzkow (2010). Persuasion: empirical evidence. *Annu. Rev. Econ. 2*(1), 643–669.

DellaVigna, S. and E. Kaplan (2007). The fox news effect: Media bias and voting. *The Quarterly Journal of Economics 122*(3), 1187–1234.

DellaVigna, S. and E. La Ferrara (2015). Economic and social impacts of the media. In *Handbook of media economics*, Volume 1, pp. 723–768. Elsevier.

Djankov, S., C. McLiesh, T. Nenova, and A. Shleifer (2003). Who owns the media? *The Journal of Law and Economics 46*(2), 341–382.

Draganska, M., M. Mazzeo, and K. Seim (2009). Beyond plain vanilla: Modeling joint product assortment and pricing decisions. *QME 7*(2), 105–146.

Dubé, J.-P., G. J. Hitsch, and P. E. Rossi (2010). State dependence and alternative explanations for consumer inertia. *The RAND Journal of Economics 41*(3), 417–445.

Durante, R. and B. Knight (2012). Partisan control, media bias, and viewer responses: Evidence from berlusconi's italy. *Journal of the European Economic Association 10*(3), 451–481.

Durante, R., P. Pinotti, and A. Tesei (2019). The political legacy of entertainment tv. *American Economic Review 109*(7), 2497–2530.

Durante, R. and E. Zhuravskaya (2018). Attack when the world is not watching? us news and the israeli-palestinian conflict. *Journal of Political Economy 126*(3), 1085–1133.

Economist, T. (2016). Vladimir putin's unshakeable popularity. `https://www.economist.com/graphic-detail/2016/02/04/vladimir-putins-unshakeable-popularity` [Acccessed: December 16, 2019].

Edmond, C. (2013). Information manipulation, coordination, and regime change. *Review of Economic Studies 80*(4), 1422–1458.

Egorov, G., S. Guriev, and K. Sonin (2009). Why resource-poor dictators allow freer media: A theory and evidence from panel data. *American political science Review*, 645–668.

Eisensee, T. and D. Strömberg (2007). News droughts, news floods, and us disaster relief. *The Quarterly Journal of Economics 122*(2), 693–728.

Eizenberg, A. (2014). Upstream innovation and product variety in the us home pc market. *Review of Economic Studies 81*(3), 1003–1045.

Enikolopov, R., M. Petrova, and E. Zhuravskaya (2011). Media and political persuasion:

Evidence from russia. *The American Economic Review 101*(7), 3253–3285.

exchange rates.org (2014). Usd to rur exchange rate on december 31, 2014. `https://www.exchange-rates.org/Rate/USD/RUB/12-31-2014` [Acccessed: December 16, 2019].

Fan, Y. (2013). Ownership consolidation and product characteristics: A study of the us daily newspaper market. *American Economic Review 103*(5), 1598–1628.

Flaxman, S., S. Goel, and J. M. Rao (2016). Filter bubbles, echo chambers, and online news consumption. *Public opinion quarterly 80*(S1), 298–320.

forbes.ru (2017). Rbc is sold: Berezkin bought the largest media publisher from prokhorov [in russian]. `https://www.forbes.ru/milliardery/346333-berezkin-kupil-u-prohorova-rbk` [Acccessed: December 16, 2019].

Garcia-Arenas, J. (2016). The impact of free media on regime change: Evidence from russia. *Working Paper*.

Gehlbach, S. and K. Sonin (2014). Government control of the media. *Journal of Public Economics 118*, 163–171.

Gentzkow, M. (2007). Valuing new goods in a model with complementarity: Online newspapers. *The American Economic Review 97*(3), 713–744.

Gentzkow, M., B. Kelly, and M. Taddy (2019). Text as data. *Journal of Economic Literature 57*(3), 535–74.

Gentzkow, M. and J. M. Shapiro (2010). What drives media slant? evidence from us daily newspapers. *Econometrica 78*(1), 35–71.

Gentzkow, M. and J. M. Shapiro (2015). Ideology and online news. In *Economic Analysis of the Digital Economy*, pp. 169–190. University of Chicago Press.

Gentzkow, M., J. M. Shapiro, and M. Sinkinson (2011). The effect of newspaper entry and exit on electoral politics. *The American Economic Review 101*(7), 2980–3018.

Gentzkow, M., J. M. Shapiro, and M. Sinkinson (2014). Competition and ideological diversity: Historical evidence from us newspapers. *American Economic Review 104*(10), 3073–3114.

Gentzkow, M., J. M. Shapiro, and D. F. Stone (2016). Media bias in the marketplace: Theory. Vol. 2., Handbook of Media Economics.

Gentzkow, M., J. M. Shapiro, and M. Taddy (2019). Measuring group differences in high-dimensional choices: Method and application to congressional speech. *Econometrica 87*(4), 1307–1340.

George, L. M. and C. Hogendorn (2020). Local news online: Aggregators, geo-targeting and the market for local news. *The Journal of Industrial Economics 68*(4), 780–818.

George, L. M. and J. Waldfogel (2006). The new york times and the market for local newspapers. *American Economic Review 96*(1), 435–447.

Groseclose, T. and J. Milyo (2005). A measure of media bias. *The Quarterly Journal of Economics 120*(4), 1191–1237.

Heckman, J. J. (1981). The incidental parameters problem and the problem of initial conditions in estimating a discrete time-discrete data stochastic process.

Hobbs, W. R. and M. E. Roberts (2018). How sudden censorship can increase access to information. *American Political Science Review 112*(3), 621–636.

Hortaçsu, A. and C. Syverson (2004). Product differentiation, search costs, and competition in the mutual fund industry: A case study of s&p 500 index funds. *The Quarterly journal of economics 119*(2), 403–456.

Hu, L., J. Li, Z. Li, C. Shao, and Z. Li (2013). Incorporating entities in news topic modeling. In *Natural Language Processing and Chinese Computing*, pp. 139–150. Springer.

Iyengar, S. and D. R. Kinder (1987). News that matters: Agenda-setting and priming in a television age. *University of Chicago Press*.

Jeziorski, P. (2014). Effects of mergers in two-sided markets: The us radio industry. *American Economic Journal: Microeconomics 6*(4), 35–73.

Kearney, M. S. and P. B. Levine (2015). Early childhood education by mooc: Lessons from sesame street. Technical report, National Bureau of Economic Research.

Kennedy, P. J. and A. Prat (2019). Where do people get their news? *Economic Policy 34*(97), 5–47.

Kim, H., Y. Sun, J. Hockenmaier, and J. Han (2012). Etm: Entity topic models for mining documents associated with entities. In *2012 IEEE 12th International Conference on Data Mining*, pp. 349–358. IEEE.

Kim, J., G. M. Allenby, and P. E. Rossi (2002). Modeling consumer demand for variety. *Marketing Science 21*(3), 229–250.

King, G., J. Pan, and M. E. Roberts (2013). How censorship in china allows government criticism but silences collective expression. *American Political Science Review 107*(2), 326–343.

King, G., J. Pan, and M. E. Roberts (2014). Reverse-engineering censorship in china: Randomized experimentation and participant observation. *Science 345*(6199).

Knight, B. and A. Tribin (2019). The limits of propaganda: Evidence from chavez's venezuela. *Journal of the European Economic Association 17*(2), 567–605.

Krasnashchok, K. and S. Jouili (2018). Improving topic quality by promoting named entities in topic modeling. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pp. 247–253.

Kuechler, M. (1998). The survey method: An indispensable tool for social science research everywhere? *American Behavioral Scientist 42*(2), 178–200.

Kumaran, G. and J. Allan (2004). Text classification and named entities for new event detection. In *Proceedings of the 27th annual international ACM SIGIR conference on Research and development in information retrieval*, pp. 297–304. ACM.

Kumaran, G. and J. Allan (2005). Using names and topics for new event detection. In

Proceedings of the conference on Human Language Technology and Empirical Methods in Natural Language Processing, pp. 121–128. Association for Computational Linguistics.

La Ferrara, E., A. Chong, and S. Duryea (2012). Soap operas and fertility: Evidence from brazil. *American Economic Journal: Applied Economics 4*(4), 1–31.

Levy, R. (2021). Social media, news consumption, and polarization: evidence from a field experiment. *American Economic Review 111*(3), 831–70.

liveinternet.ru (2014). Web page with weekly public ranking of news websites in russia [in russian, accessed through wayback machine]. `https://web.archive.org/web/20141214080407/http://www.liveinternet.ru/rating/ru/media/week.html` [Acccessed: January 26, 2020].

Martin, G. J. and A. Yurukoglu (2017). Bias in cable news: Persuasion and polarization. *American Economic Review 107*(9), 2565–99.

Mazzeo, M. J. (2002). Product choice and oligopoly market structure. *RAND Journal of Economics*, 221–242.

McAlister, L. and E. Pessemier (1982). Variety seeking behavior: An interdisciplinary review. *Journal of Consumer research 9*(3), 311–322.

McCombs, M. E. and D. L. Shaw (1972). The agenda-setting function of mass media. *Public opinion quarterly 36*(2), 176–187.

Melnikov, N. (2019). Censorship, propaganda, and political popularity: Evidence from russia. *Available at SSRN 3276926*.

Mikolov, T., I. Sutskever, K. Chen, G. S. Corrado, and J. Dean (2013). Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, pp. 3111–3119.

Mullainathan, S. and A. Shleifer (2005). The market for news. *The American Economic Review 95*(4), 1031–1053.

Petrova, M. (2008). Inequality and media capture. *Journal of public Economics 92*(1), 183–212.

PewResearchCenter (2017). Key trends in social and digital news media. `https://www.pewresearch.org/fact-tank/2017/10/04/key-trends-in-social-and-digital-news-media/` [Acccessed: December 16, 2019].

Prat, A. (2018). Media power. *Journal of Political Economy 126*(4), 1747–1783.

Prat, A. and D. Strömberg (2013). The political economy of mass media. *Advances in Economics and Econometrics Volume 2, Applied Economics*.

Qin, B., D. Strömberg, and Y. Wu (2017). Why does china allow freer social media? protests versus surveillance and propaganda. *Journal of Economic Perspectives 31*(1), 117–40.

Qin, B., D. Strömberg, and Y. Wu (2018). Media bias in china. *American Economic Review 108*(9), 2442–76.

rbc.ru (2015). Channel one and vgtrk will get compensated for devaluation [in russian]. https://www.rbc.ru/politics/29/06/2015/55912ffa9a7947453982cda9 [Accessed: December 16, 2019].

republic.ru (2012). Eight figures: Who are these people supporting putin? [in russian]. https://republic.ru/posts/l/745894 [Acccessed: December 16, 2019].

Roberts, M. E. (2014). Fear or friction? how censorship slows the spread of information in the digital age. *Working Paper*.

Roberts, M. E. (2020). Resilience to online censorship. *Annual Review of Political Science 23*, 401–419.

Rossi, P. (2014). *Bayesian non-and semi-parametric methods and applications*. Princeton University Press.

Rossi, P. E., G. M. Allenby, and R. McCulloch (2005). *Bayesian statistics and marketing*. John Wiley & Sons.

rsf.org (2016). New russian law targets news aggregators. https://rsf.org/en/news/new-russian-law-targets-news-aggregators [Acccessed: October 31, 2020].

Schimpfossl, E. and I. Yablokov (2014). Coercion or conformism? censorship and self-censorship among russian media personalities and reporters in the 2010s. *Demokratizatsiya: The Journal of Post-Soviet Democratization 22*(2), 295–311.

Seim, K. (2006). An empirical model of firm entry with endogenous product-type choices. *The RAND Journal of Economics 37*(3), 619–640.

Sen, A. and P. Yildirim (2016). Clicks bias in editorial decisions: How does popularity shape online news coverage? *Working paper*.

Shadmehr, M. and D. Bernhardt (2015). State censorship. *American Economic Journal: Microeconomics 7*(2), 280–307.

Simonov, A., J.-P. Dubé, G. Hitsch, and P. Rossi (2020). State-dependent demand estimation with initial conditions correction. *Journal of Marketing Research 57*(5), 789–809.

squirepattonboggs.com (2014). Russia moves to limit foreign ownership in the media section. http://www.squirepattonboggs.com/~/media/files/insights/publications/2014/10/russia-moves-to-limit-foreign-ownership-in-the-2__/files/russiamovestolimitforeignownership/fileattachment/russiamovestolimitforeignownership.pdf [Acccessed: December 16, 2019].

statcounter.com (2015). Desktop browser market share russian federation. https://gs.statcounter.com/browser-market-share/desktop/russian-federation/ [Accessed: December 16, 2019].

stopfake.org (2014). A fight with the dishonest information about the ukraine events [in russian]. http://www.stopfake.org/propaganda-2-0-novye-izmereniya-v-dejstvii/ [Acccessed: December 16, 2019].

Strömberg, D. (2004). Mass media competition, political competition, and public policy. *The Review of Economic Studies 71*(1), 265–284.

Strömberg, D. (2015). Media and politics. *economics 7*(1), 173–205.

Sullivan, C. (2017). The ice cream split: Empirically distinguishing price and product space collusion. *Available at SSRN 3321948*.

Sweeting, A. (2013). Dynamic product positioning in differentiated product markets: The effect of fees for musical performance rights on the commercial radio industry. *Econometrica 81*(5), 1763–1803.

Szeidl, A. and F. Szucs (2021). Media capture through favor exchange. *Econometrica 89*(1), 281–310.

themoscowtimes.com (2013). Kremlin helps media moguls expand. https://www.themoscowtimes.com/2013/10/20/kremlin-helps-media-moguls-expand-a28748 [Accessed: December 16, 2019].

themoscowtimes.com (2014). The media war behind the ukraine crisis. https://www.themoscowtimes.com/2014/03/10/the-media-war-behind-the-ukraine-crisis-a32837 [Acccessed: December 16, 2019].

Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Methodological) 58*(1), 267–288.

time.com (2014). Russians rewrite history to slur ukraine over war. https://time.com/3545855/russia-ukraine-war-history/ [Acccessed: December 16, 2019].

Train, K. E. (2009). *Discrete choice methods with simulation.* Cambridge university press.

TrustLaw (2016). Media regulation in russia: A landscape analysis of laws and trends. https://www.trust.org/contentAsset/raw-data/4798c68a-eed1-4660-b7c9-fc16a0032cc9/file [Acccessed: December 16, 2019].

vesti.ru (2014). 116 people burned alive by fascists in odessa [in russian]. https://www.vesti.ru/doc.html?id=1550135 [Acccessed: December 16, 2019].

VTsIOM (2017). Internet against tv: The battle continues [in russian]. https://wciom.ru/index.php?id=236&uid=116190 [Acccessed: December 16, 2019].

Wollmann, T. G. (2018). Trucks without bailouts: Equilibrium product characteristics for commercial vehicles. *American Economic Review 108*(6), 1364–1406.

Xiang, Y. and M. Sarvary (2007). News consumption and media bias. *Marketing Science 26*(5), 611–628.

Yanagizawa-Drott, D. (2014). Propaganda and conflict: Evidence from the rwandan genocide. *The Quarterly Journal of Economics 129*(4), 1947–1994.

Zhu, Y. and A. Dukes (2015). Selective reporting of factual content by commercial media. *Journal of Marketing Research 52*(1), 56–76.

Zhuravskaya, E., M. Petrova, and R. Enikolopov (2020). Political effects of the internet and social media. *Annual Review of Economics 12*, 415–438.