GLOBAL FOCUS

#ThinkResilience

RUSSIAN DISINFORMATION IN SLOVAKIA:

AN EXPERIMENTAL APPROACH*

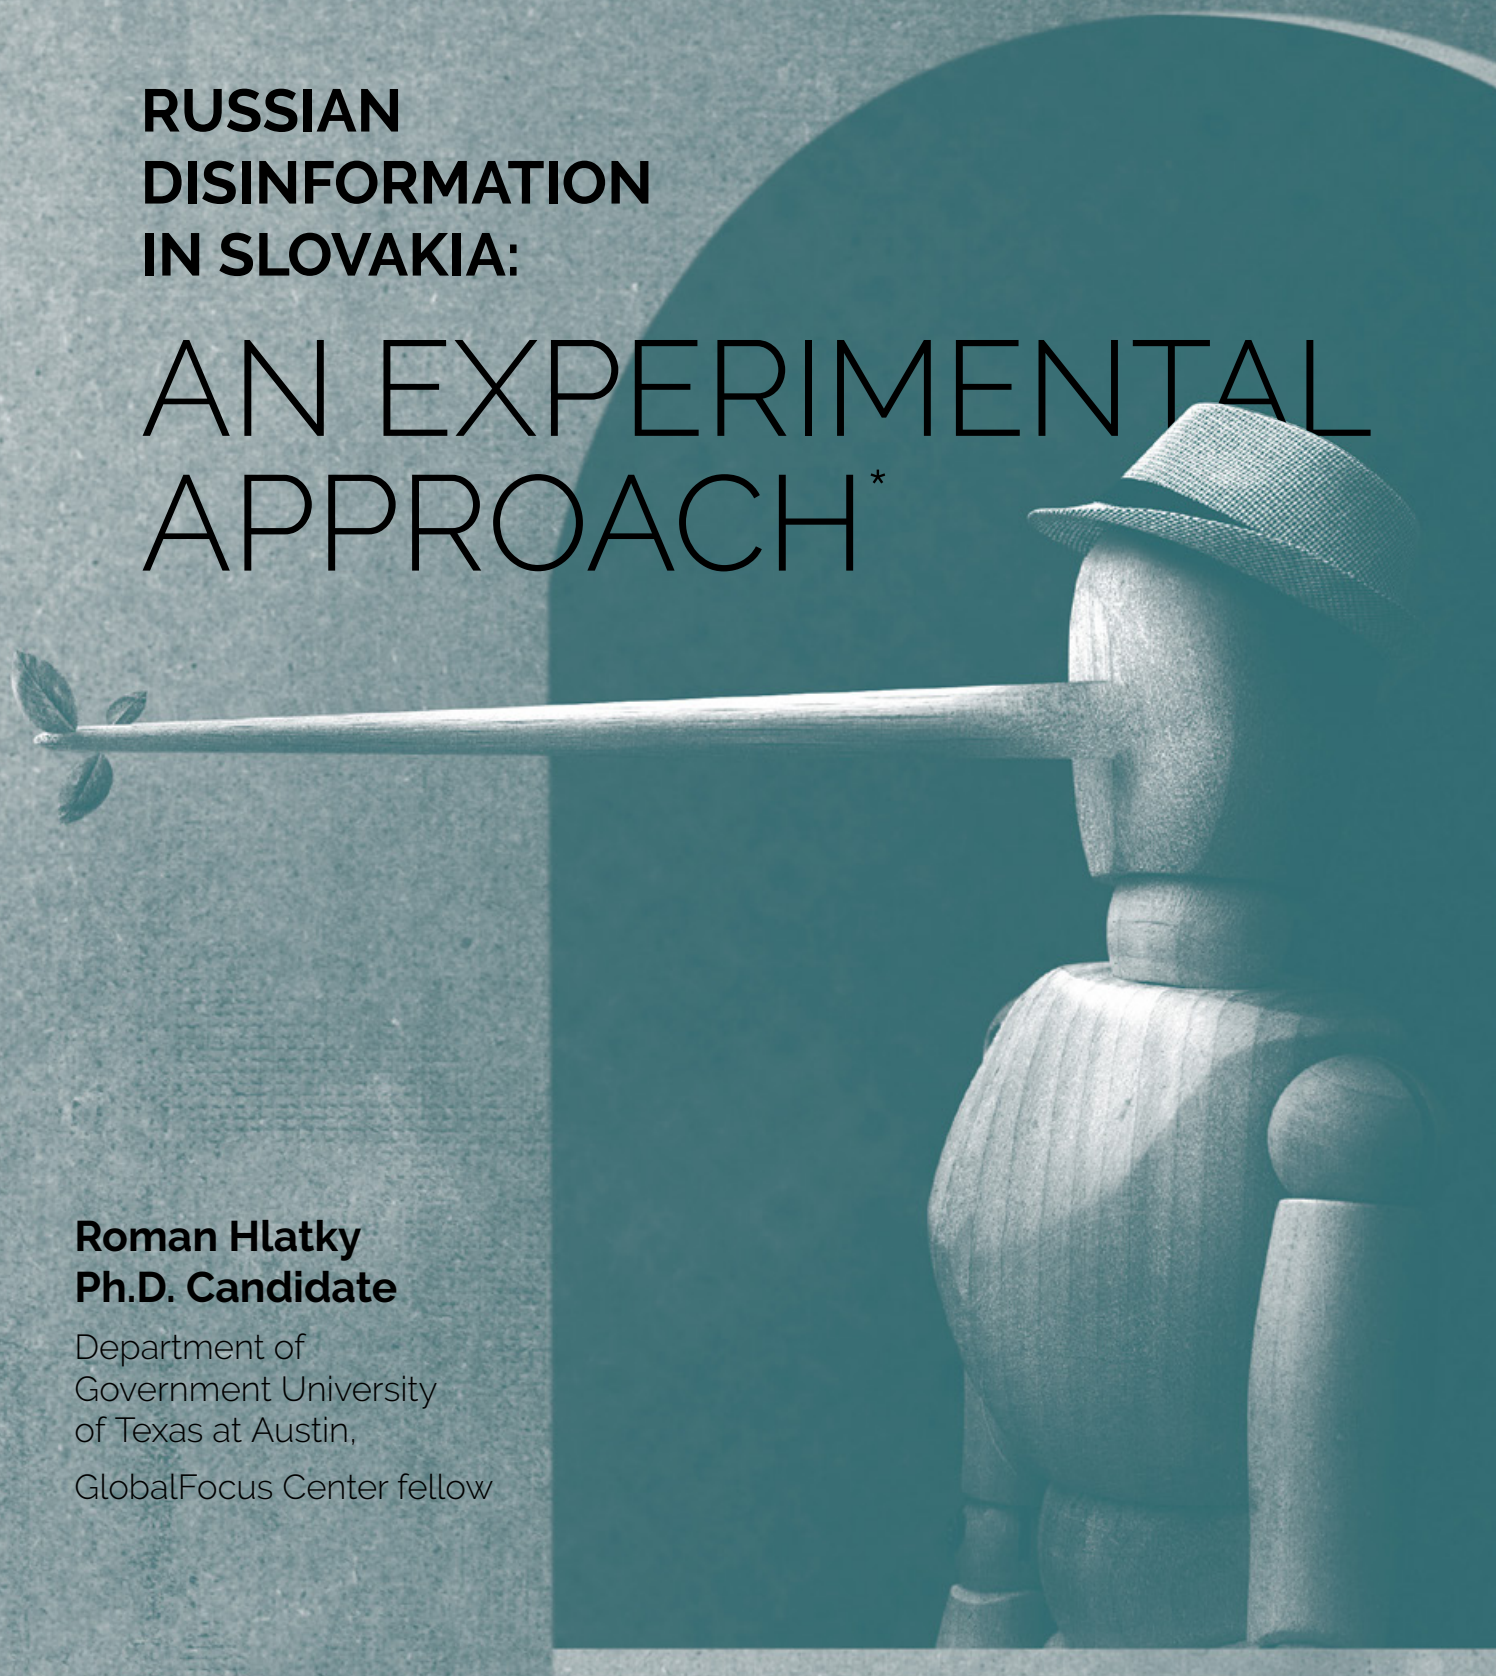

Roman Hlatky
Ph.D. Candidate

Department of Government University of Texas at Austin,

GlobalFocus Center fellow

A project by

Bucharest, Romania
June, 2021
The report can be accessed at www.global-focus.eu
or ordered at office@global-focus.eu
+40 721 259 205
Str Dimitrie Orbescu 8, et. 2, ap. 5, Bucharest, Sector 2

GlobalFocus Center is an independent international studies think-tank which produces in-depth research and high quality analysis on foreign policy, security, European affairs, good governance and development. It functions as a platform for cooperation and dialogue among individual experts, NGOs, think-tanks and public institutions from Central and Eastern Europe and Euro-Atlantic partners.

A project supported by

B | S | T The Black Sea Trust for Regional Cooperation
A PROJECT OF THE GERMAN MARSHALL FUND

OPEN INFORMATION PARTNERSHIP

*This report presents a subset of data and findings from a larger academic study by the author. Upon publication of the study, this report will be updated with the relevant citations, and data and associated code will be made publicly available. The University of Texas at Austin Institutional Review Board approved this research under study number: 2020-04-0091.

The views expressed in this report belong to the author of the study and do not necessarily reflect the position of the GlobalFocus Center.

Layout by Florin Vedeanu

Table of contents

Introduction

In the past fifteen or so years, Russia has shifted towards the sustained use non-kinetic means to achieve geopolitical goals. While a healthy scholarly debate has emerged about what to call this combination of non-kinetic and kinetic methods (e.g., Galeotti 2018), it is clear that the Russian foreign policy toolbox has expanded. In this report, I focus on the effects of one non-kinetic tool: disinformation. Specifically, I ask: (1) to what extent do individuals read and trust disinformation news outlets; (2) does Russian disinformation influence individual attitudes, and if so, to what extent; (3) does exposure to disinformation convince individuals to take steps towards combatting it; and (4) does demographic and ideological variation matter in how individuals consume and respond to disinformation?

Disinformation refers to the deliberate provision of false or misleading information for the purpose of influencing a target. At the highest level, Russia's primary foreign policy objective is to undermine (U.S.-led) unipolarity and reinstate Russian interests in both the post-Soviet space (the Russian "near abroad") and on the global stage. Disinformation helps Russia achieve this goal. It does so by (1) discouraging other states from challenging Russia's kinetic actions through the production of alternative narratives and interpretations; (2) creating doubt and scepticism about the global system of governance; (3) fomenting discord within and amongst other states, hampering their ability to act against Russia; and (4) producing a positive image of Russia, while challenging the image of Western institutions (see Čížik 2017, pp.8-28, Richey 2018, Šuplata & Nič 2016).

While Russia employs disinformation globally, Central and Eastern Europe (CEE) is a vector of particular and active interest. Note that though I discuss the region uniformly here, historical legacies and contemporary dynamics do lead to clear differences amongst these states in terms of both Russian strategy and domestic vulnerability to disinformation (see, for example, Milo & Klingová 2017). Yet, despite these differences, several commonalities underlie Russian interests in the region. First, Russia considers CEE a weak spot in the Euro- Atlantic political structure (Juhasz & Szicherle 2017). By exacerbating existing Eurosceptic, nationalist, conservative, and anti-American attitudes amongst certain segments of the CEE public and elite, Russia can create dissensus within both the European Union (EU) and the North Atlantic Treaty Organization (NATO). Second, historical, cultural, and ideological affinities between (some) CEE states and Russia also magnify the reach and efficacy of disinformation. Third, many CEE states have limited – and an only recently developed – institutional capacity to combat Russian disinformation (see Galeotti 2017 and Janda et al. 2017). Thus, CEE is very much an active and often receptive target of Russian disinformation efforts.

While existing work on the region – much of it cited in this report – has provided a wealth of valuable information about Russian disinformation efforts in CEE, we still lack rigorous, causally identified work investigating the effects of Russian disinformation. This report attempts to remedy this gap. It does so by analyzing the results of an experimental survey conducted with a representative sample of the Slovak population (N=1400).

Research context: The Slovak Republic

Slovakia is an important case for understanding Russian disinformation efforts in CEE. First, compared to the rest of CEE, Slovaks are fairly pro-Russia. According to a 2020 Pew Research Poll, 60 percent of Slovaks have a favorable image of Russia. This is second to only Bulgaria (73 percent). Moreover, the Globesec Trends 2020 Survey found that 11 percent of Slovaks believe Slovakia should belong to the East. This was the second highest percentage amongst the nine countries surveyed (Montenegro was first with 25 percent). According to the same survey, 78 percent of Slovaks believe that Russia is Slovakia's traditional Slavic brother/sister nation, showcasing the strength of Pan-Slavic ties between the two countries[1]. In fact, much of the current scholarship attributes strong pro-Russian sympathies to precisely these Slavic cultural and identity ties (e.g., Kreko et al. 2017), but the pro-Russian rhetoric of political elites certainly plays a role as well (see Kreko et al. 2015 and Mesežnikov & Bránik 2017 for further details and Russian connections beyond mainstream political parties).[2]

Second, the Slovak public has somewhat conflicting opinions about the EU. While the Globesec Trends 2020 Survey found that 75 percent of Slovaks would stay in the European Union, it also found that 40 percent of Slovaks think the EU is destroying traditional values and identity, and that 55 percent of Slovaks believe that Slovakia does not have a say in EU decision-making. These data illustrate a salient divide in Slovak attitudes towards the EU. On the one hand, Slovaks appreciate the economic benefits of EU membership, and value being a part of European community. On the other hand, perceived EU influence in migration and other identity-related polices spurs Euroscepticism (e.g., Henderson 2017).

Together, pro-Russian sentiment and ambivalent stances towards the EU make the Slovak Republic particularly vulnerable to Russian disinformation efforts. These attitudinal factors are further exacerbated by an only recently developed institutional capacity to combat dis- information (see Kandrík & Jevčák 2018). According to the cross-national Disinformation Resilience Index, Slovakia ranks only above Moldova in the quality of its institutional re- sponses to disinformation. Thus, we can consider Slovakia an extreme case – by analyzing Russian disinformation in the Slovak context, we can potentially identify the upper limits of disinformation's effects.

In order to identify the effects of Russian disinformation, I fielded an experimental survey with a representative sample of the Slovak population (N=1400). I sought to achieve three research objectives: (1) establish the degree of disinformation exposure amongst the Slovak public; (2) identify whether exposure to disinformation increases anti-EU and nationalist attitudes; (3) determine whether exposure to disinformation renders individuals more willing to support anti-disinformation efforts; and (4) understand how demographic and ideological characteristics condition individual responses to disinformation.

Disinformation consumption:

Media sources

To measure the degree of possible disinformation consumption amongst the Slovak public, respondents completed a short questionnaire on their media consumption patterns. Respon dents indicated whether they recognize, read, and how much they trust nine different Slovak news sources. Five of the news sources were mainstream media outlets, while four were media outlets known to publish disinformation news articles. Recognition and readership were binary choices, yes or no. Trust was measured on a five-point scale. Overall, the survey results paint a fairly optimistic picture of disinformation consumption. Not only are disinformation news outlets less recognized than mainstream ones (37.9 percent compared to 91.3 percent), they are also less read (14.5 percent to 53 percent) and less trusted (2.57 to 2.95).[3] Table 1, below, presents these results, and also shows the scores for individual news sources. The most widely read and trusted disinformation news source is

Hlavné Správy; in fact, *Hlavné Správy* is trusted to the same extent as *Plus Jeden Deň*, *Nový Čas,* and *Denník N. Hlavné Správy's* high levels of readership and trust likely result from its reporting of mainstream news alongside disinformation articles, relatively professional appearance, and its niche as a self-titled conservative daily. The least read and trusted disinformation news source is *ExtraPlus*, which publishes largely fringe and sensationalist content.

Table 1: Recognition, readership, and trust of news sources

Source	Recognition (%)	Readership (%)	Trust (1-5)
Mainstream sources	91.3	53.0	2.95
Disinfo sources	37.9	14.5	2.57
Nov´y Čas	96.9	59.1	2.72
Sme	95.6	52.6	3.07
Plus Jeden Deň	92.6	42.0	2.79
aktuality.sk	90.8	69.5	3.28
Denník N	80.9	41.6	2.88
Hlavné Správy	65.0	32.9	2.82
Zem a Vek	41.5	10.9	2.52
Slobodny Vysielač	35.0	9.7	2.56
ExtaPlus	9.9	4.6	2.38

These results present the media consumption patterns of an aggregate, representative sample. In other words, they show what media consumption looks like in the Slovak popula- tion as a whole. As such, the numbers do not tell us very much about the media consumption of different demographic and ideological subgroups within the Slovak population. Thus, to better understand the reach of disinformation news outlets, I consider the roles of age, edu- cation, geopolitical orientation, and vote choice. Scholarship suggests that these individual- level characteristics can moderate how individuals consume and respond to disinformation. There are few differences between these subgroups of respondents when it comes to recogni- tion, thus I focus on readership (measured as percent of sample) and trust (measured on a five-point scale).

Age

Studies of disinformation consumption suggest that both the young and the old can be particularly susceptible to disinformation narratives. To see whether these patterns hold in Slovakia, I compare the the media consumption of (1) individuals who are 25 or younger and (2) individuals who are 50 or older to the rest of the population. Figure 1 presents these comparisons.

Figure 1: Readership and trust: Age

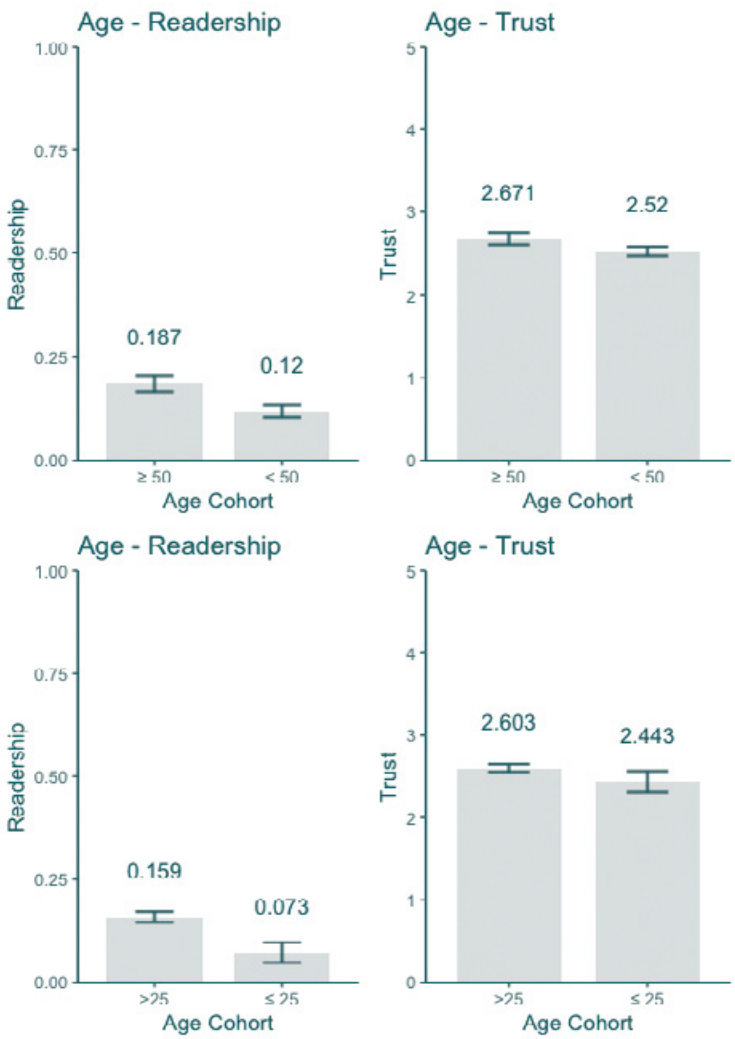

As Figure 1 shows, both readership and trust of disinformation news sources increases with age. Those who are younger than 25 read and trust disinformation less than those who are younger than 50; and those who are younger than 50 read and trust disinformation less than those who are 50 or older.[4] Thus, it seems that worries about youth exposure to disinformation are not substantiated by the data. Rather, the data suggest that interventions and media literacy programs would do well to target older individuals.

Education

Next, I focus on education, comparing respondents with and without university degrees. As Figure 2 shows, the university educated are slightly less likely to read disinformation than those without university degrees. This difference is not statistically significant, however. When it comes to trust, individuals with university degrees are significantly (at $p < 0.05$) less trusting of disinformation news sources compared to those without university degrees. Thus, while a university education may not play a substantive role in whether or not an individual reads disinformation news sources, it does seem to correlate with scepticism towards these outlets.

Figure 2: Readership and trust: Education

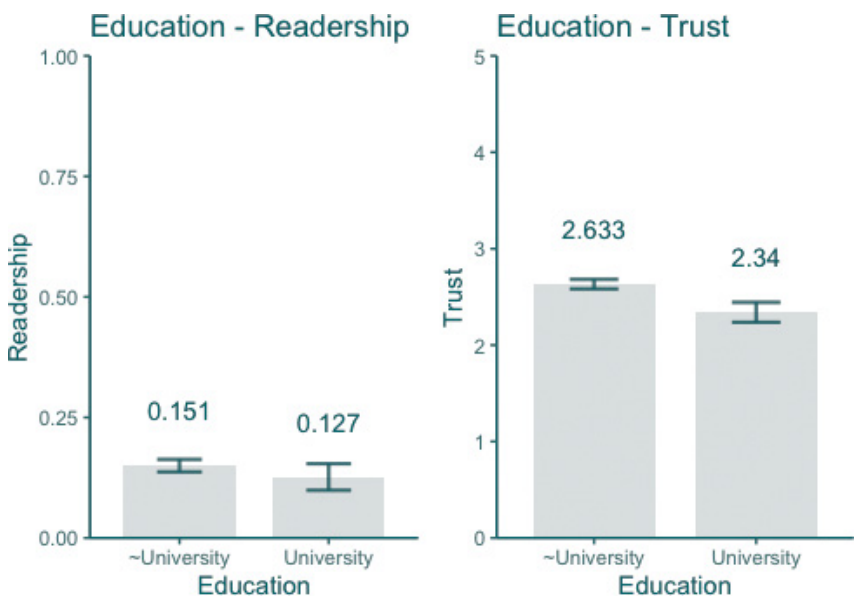

Geopolitical orientation

Strong pro-Russian sentiments may make the Slovak public particularly vulnerable to disinformation. To see whether these sympathies correlate with disinformation consumption, I compare respondents who think Slovakia should belong to the East with those who think Slovakia should belong to the West.[5] As Figure 3 shows, stark differences emerge amongst respondents with differing geopolitical orientations. Those who think Slovakia should belong to the East are both more likely to read and trust disinformation news sources than those who think Slovakia should belong to the West.

Figure 3: Readership and trust: Geopolitical orientation

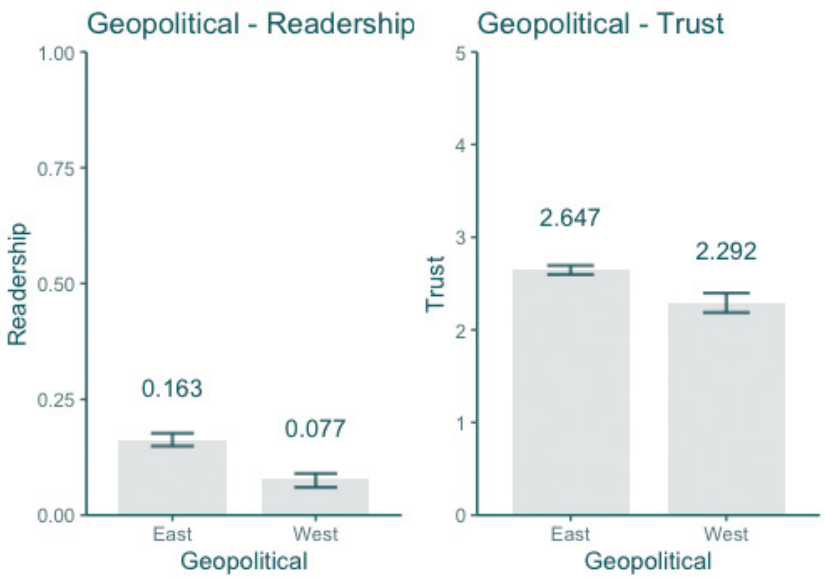

Nationalist voters

The final subgroup analysis concerns nationalist voters versus other voters.[6] Scholarship has linked disinformation to increasing polarization and support for extremist political par- ties (e.g., Conley et al. 2016). This analysis allows us to determine if the disinformation consumption differs between voters who support nationalist movements and those who do not.

Figure 4: Readership and trust: Voting

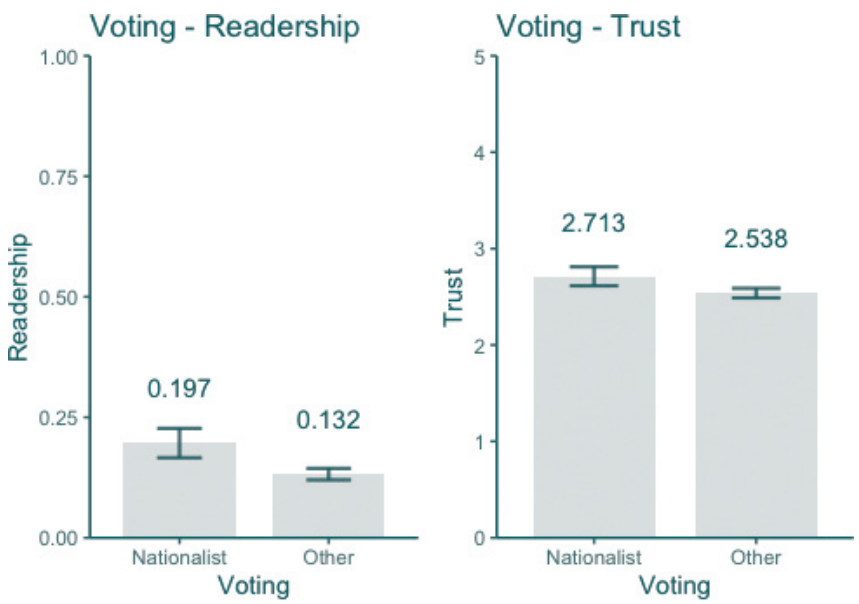

I present differences between these groups of voters in Figure 4. Clear differences are evident between nationalist and other voters in terms of both readership and trust. Unsur- prisingly, nationalist voters read disinformation sources at higher rates than other voters, and they are also more trusting of these sources.[7] These results suggest that ideology plays an important role in how individuals consume disinformation.

The effects of disinformation exposure

Next, I consider whether exposure to disinformation news headlines shifts individual atti- tudes and willingness to combat disinformation. To do so, I randomly assigned respondents into three groups: (1) the *Control group which received no stimulus; (2) the False group* which read a set of four false, Russian disinformation news headlines; and (3) the *True* group which read a set of four similar, but completely true news headlines.[8]

By comparing the survey responses of respondents in these different groups, I can identify whether exposure to disinformation news headlines leads to any changes in attitudes or behaviors.[9] Specifically, I consider two sets of attitudes: (1) evaluations of the EU and (2) nationalist orientations; and two behavioral outcomes measuring willingness to combat disinformation: willingness to sign a petition that asks either (1) the Slovak government or (2) Facebook to create a commission to combat disinformation and hoaxes in Slovakia.

Headline selection

A major focus of Russian disinformation in Slovakia (and broader CEE) concerns the EU and its supposed propagation of "liberal" identity values at odds with traditional CEE societies. Major topics include migration, LGBTQIA+ rights, and the "decadent" or immoral West more generally (see Čížik & Masariková 2018; Ižak 2018; Šuplata & Nič 2016). Thus, I focused on news headlines that concern these themes.

To identify false news headlines, I relied on the EU vs. Disinfo Database. The EU vs. Disinfo Database is a project run by the EU's External Action Service's East Stratcom Task Force. The project identifies prominent Russian disinformation stories propagated in EU member states. For this project, I chose four false news headlines that explicitly linked the EU with liberal identity politics. I avoided overly sensationalist or unrealistic headlines that respondents could easily identify as false or embellished. Next, I chose a set of similar, although completely true, headlines to ensure comparability between the two experimental groups. The survey presented all headlines to respondents as mock Facebook posts. As an example, Figure 5 shows one of the false headlines presented to respondents.

Figure 5: Example of false headline

Belief, share-ability, and interest

Prior to analyzing the effects of headline exposure on attitudes and behavior, I test whether respondents judged false versus true headlines differently. Specifically, I consider the extent to which respondents believed (four-point scale, higher values being more true), and would share (four-point scale, higher values being more likely to share) the headlines. I also recorded whether respondents wanted more information about each of the headlines presented to them (yes or no). Respondents answered these questions for each headline they saw, and I took the average for all four headlines to create an overall score for each respondent and outcome.

False and true news headlines evoked different levels of belief amongst respondents. On average, false news headlines had a belief score of 1.91 while true headlines had a score of 2.20,[10] indicating that respondents could accurately adjudicate the veracity of the headlines presented to them. Surprisingly, however, false and true headlines elicited the same willing-ness to share amongst respondents. However, the scores – 1.43 for false and 1.46 for true – were relatively low, suggesting that respondents were unwilling to share news headlines generally. Finally, false headlines piqued interest slightly more than true ones, though the difference was not statistically significant. On average 15.8 percent of respondents wanted more information about false headlines, while 12.5 percent wanted more information about true headlines.

Thus far, I have treated both true and false headlines as aggregate groups, but responses to different false headlines varied widely in terms of belief, sharing, and registration. One headline in particular stood out: a news headline suggesting that European Union funding ("Eurofunds") allocations would be linked to the acceptance of migrants by EU member states. On the four-point belief scale, this headline received a 2.20, the same score as the overall average for true headlines. This headline also evoked the highest level of sharing (1.67) and registration for further information (17 percent). The second-place headline in terms of both belief and sharing concerned the EU's enforcement of the Istanbul Convention.[11] Thus, it seems that respondents are most likely to believe disinformation covering contentious identity issues salient in public discourse. Likewise, these headlines garner the most interest. In contrast to false headlines, there were few substantive differences in belief, sharing, or registration amongst the true headlines.

Effects on EU attitudes

A major goal of Russian disinformation efforts is weakening support for the EU. Thus, I test whether exposure to the false news headlines lowers individual-level support for the EU. To measure EU support, I use an index of three questions: an EU feeling thermometer (1-4); a question about whether Slovakia should exit the EU (1-4); and a question asking whether EU membership benefits Slovakia (1-4). Higher values indicate more pro-EU stances. [12] I present the index scores and associated 95% confidence intervals for each of the three groups – control, true, and false – in Table 2 below.

Table 2: Attitudes towards the EU across experimental groups

Group	EU index score	95% CI
Control	2.597*	[2.509, 2.686]
False	2.744*†	[2.652, 2.835]
True	2.601†	[2.507, 2.695]

Note: * and † indicate a significant difference between values at $p < 0.05$.

As Table 2 shows, the *Control* group (which was not exposed to any news headlines) and the *True* group have nearly identical stances towards the EU. This suggests that true headlines, even if they cover similar topics to disinformation headlines, do little to shift indi- vidual attitudes. Importantly, we see a surprising response to the disinformation headlines. Individuals exposed to the false headlines are more pro-EU than those in both the *Control* and *True* groups. This counter-intuitive finding raises an important question: why did, on average, individuals respond to Russian disinformation by becoming more positive towards the EU?

To answer this question, I sort respondents into two groups using the survey on media sources: those that do and do not trust disinformation news outlets.[13] The purpose of this analysis is to identify whether trust in disinformation more generally impacts how individuals

respond to disinformation news headlines. Table 3 compares the EU index score for respon- dents in the *Control and False groups based on whether or not they trust disinformation* news outlets.

Table 3: Attitudes towards the EU: Trust of disinformation

Group	EU Index Score	95% CI
Control trusting	2.603	[2.488, 2.717]
Control not trusting	2.590*	[2.450, 2.730]
False trusting	2.559	[2.431, 2.686]
False not trusting	2.957*	[2.831, 3.083]

Note: * indicates a significant difference between values at $p < 0.05$.

The results presented in Table 3 help clarify the pro-EU effect we see in the entire sample. First, there is no significant difference between the *Control* and *False* groups when comparing individuals who trust disinformation. In other words, exposure to disinformation news headlines does not strengthen anti-EU attitudes amongst these individuals. Conversely, disinformation exposure leads to a significant difference amongst individuals who do not trust disinformation. For these individuals, reading false news headlines about the EU and liberal identity policy leads to a substantial pro-EU effect.

As an additional test, I also analyze whether individual reactions to the false headlines depend on whether or not individual actively read disinformation. As in the trust analysis, I divide respondents by whether they are, or are not, readers of disinformation.[14] Here, we see a similar, differential effects: the false news headlines do not lead to significant differences amongst individuals who read disinformation news outlets; however, amongst individuals who are not readers, the false news headlines again lead to a pro-EU effect. Table 4 presents these results numerically.

Table 4: Attitudes towards the EU: Readership of disinformation

Group	EU index score	95% CI
Control readers	2.465	[2.313, 2.618]
Control non-readers	2.669*	[2.561, 2.777]
False readers	2.469	[2.310, 2.628]
False non-readers	2.918*	[2.812, 3.023]

Note: * indicates a significant difference between values at $p < 0.05$.

Thus, it seems that a respondent's level of engagement with disinformation – whether in the form of trust or readership – plays an important role in how they respond to disinfor- mation about the EU. If individuals trust and read disinformation news outlets, exposure to false new headlines does not shift their attitudes. However, if individuals are not trusting consumers, exposure to false headlines leads to a pro-EU effect. These differential findings are likely a result of different underlying attitudes in the two groups. Those who consume disinformation are likely to hold anti-EU stances *ex-ante*. In contrast, individuals who do not read or trust disinformation are likely to be more pro-EU. Thus, for the former group, false headlines do little to move attitudes further in an anti-EU direction. For the latter group, false headlines present information possibly at odds with extant beliefs, evoking a pro-EU backlash. Next, I consider the effects of disinformation on nationalist attitudes.

Effects on nationalist attitudes

By attacking the EU and by focusing on topics like migration, Russian disinformation hopes to stoke identity-related grievances amongst target groups. These messages foment fears about threats to national identity, cultural homogeneity, and traditional society. In turn, these fears may empower nationalist actors, and ultimately weaken the EU from within.

To investigate the effect of Russian disinformation on nationalist attitudes, I constructed an index from three questions that asked respondents about the extent to which they agree with the following statements: (1) some races or ethnic groups are born less intelligent than others; (2) some races or ethnic groups are born less hard-working than others; and (3) some cultures are much better than others.[15] Higher scores on this index indicate more nationalist attitudes. I then measured whether the index scores differed between the experimental groups. Table 5 presents these results.

Table 5: Nationalist attitudes

Group	Nat.	Index score	95% CI
Control		2.354*ǂ	[2.268, 2.441]
False		2.092*	[1.996, 2.188]
True		2.174ǂ	[2.082, 2.266]

Note: * and ǂ indicate a significant difference between values at $p < 0.05$.

The nationalism scores for both the *False* and *True* groups statistically differ from the score of the *Control* group. Substantively, both true and false headlines led to reductions in nationalist attitudes. Again, these differences seem at odds with the intended purpose of Russian disinformation. Moreover, even true headlines about migration and related issues should evoke some sentiments of threat. To disentangle these results, I follow the same approach as with the EU attitudes, and look at the role of trust in disinformation. Since there is also a substantive and statistical difference between the *Control* and *True* groups (which was not the case for EU attitudes), I analyze the role of disinformation engagement in the *True* group as well. Table 6 presents comparisons between the groups based on whether respondents do or do not trust disinformation news sources.

Table 6: Nationalist attitudes: Trust in disinformation

Group	Nat. Index Score	95% CI
Control trusting	2.457*ǂ	[2.346, 2.568]
Control not trusting	2.222ǂ	[2.088, 2.357]
False trusting	2.241*	[2.109, 2.373]
False not trusting	1.920ǂ	[1.785, 2.055]
True trusting	2.284ǂ	[2.160, 2.408]
True not trusting	2.047	[1.910, 2.183]

Note: *, ǂ, and ǂ indicate a significant difference between values at $p < 0.05$.

In contrast to the EU results, exposure to news headlines, whether true or false, leads individuals to have less nationalist attitudes. In all but one comparison, ***Control Not Trusting*** versus ***True Not Trusting***, these differences are statistically significant. Thus, exposure to false headlines about the EU's propagation of migration and other liberal identity values leads to a reduction in nationalist attitudes amongst both those who trust and those who do not trust disinformation. Exposure to similar, but true, headlines leads to a reduction in nationalist attitudes only amongst those who trust disinformation news sources.

Together, these results suggest three conclusions. First, Russian disinformation narratives can have a potential backlash effect. Exposing individuals to false headlines that are meant to evoke identity threat can actually reduce nationalist attitudes, and it appears to do so regardless of whether respondents do or do not trust disinformation news outlets. Second, false headlines have larger effects than true headlines – i.e., they reduce nationalist attitudes to a greater degree. Third, exposing those who trust disinformation to true news headlines also reduces nationalist attitudes. This reduction likely results from high baseline levels of nationalism amongst these individuals, and indicates that corrective, true information can potentially be used to lower nationalist sentiment. An analysis of disinformation readership instead of trust leads to substantively similar conclusions.[16] Overall, these findings suggest that the ability of Russian disinformation to foment nationalist sentiment is relatively limited.

Effects on behavior: Combatting disinformation

The final outcome that I consider in this report is individual willingness to combat disinfor- mation. Specifically, the survey asked respondents if they would be willing to sign a petition that asks either (1) the Slovak Government or (2) Facebook to create a commission that in- vestigates hoaxes and disinformation in Slovakia. I expected both the *True* and *False* groups to show a greater willingness than the *Control group to sign the petitions. I also expected* the effect size to be larger in the *False group* than in the *True group*.

In the entire sample, 52.5 percent of respondents answered that they would like the Slovak government to form the commission, while 46.7 percent indicated that they would like Facebook to form one. When analyzing differences between experimental groups, no significant differences emerged. Thus, it appears that exposure to true or false headlines makes little difference in terms of respondent preferences for combatting disinformation. Nonetheless, the fact that over 50 percent of Slovaks want their government to do something about disinformation is reassuring.

Disinformation media in Slovakia: A limited audience?

Russia actively uses disinformation as a non-kinetic instrument to achieve foreign policy goals. In CEE, disinformation targets the Euro-Atlantic political structure, and exacerbates existing Eurosceptic, anti-American, and nationalist attitudes. In this report, I focused on the effects of disinformation in the Slovak Republic, a country particularly susceptible to Russian influence efforts. Specifically, I employed an experimental survey to measure disinformation consumption patterns amongst the Slovak public, and to identify the effects of disinformation exposure on anti-EU and nationalist attitudes, as well on willingness to combat disinformation. Overall, the results paint an optimistic, though nuanced, picture of disinformation's effects in Slovakia.

When it comes to the consumption of disinformation, relatively few Slovaks read, recognize, and trust disinformation news outlets. Readership of disinformation news outlets averages 14.5 percent, while the average for mainstream news sources is 53 percent. Moreover, only a third of the Slovak public recognizes disinformation news outlets, while over 90 percent rec- ognize mainstream news sources. Finally, Slovaks show lower levels of trust in disinformation outlets relative to mainstream outlets, though there are some variations when considering individual news sources. For example, respondents trust *Hlavné Správy* to the same extent as several mainstream news sources. This underscores an important point: disinformation news outlets come in various forms and packages. When a media source, like *Hlavné Správy,* combines disinformation narratives with true articles and maintains a professional appearance, it garners higher levels of readership and trust. Conversely, the more fringe and sensationalist a news source, the less likely it is to be read and believed (as in the case of *ExtraPlus).*

Existing literature suggests that certain demographic and ideological characteristics make individuals more likely to seek out and believe disinformation. The findings presented in this report suggest that age, education, geopolitical orientation, and ideology play an important role in how individuals consume disinformation. Those who are older, think that Slovakia should belong to the East, and vote for nationalist parties are all more likely to read and trust disinformation news outlets. And while university education does not seem to correlate with readership, individuals without university degrees are more trusting of disinformation news outlets.

However, even when considering those who are most predisposed to reading and trusting

disinformation, the actual rates of trust and readership are still relatively low. For example, less than 20 percent of nationalist voters read disinformation news outlets (compared to 13.2 percent of all other voters). In addition, while the difference in trust between nationalist and other voters is statistically significant, substantively it represents only a quarter of a standard deviation shift. Thus, even amongst the most willing audience, the reach of disinformation appears to be relatively limited. It is also important to note that many of these demographic and ideological characteristics are not mutually exclusive. For example, a nationalist voter is probably more likely to believe that Slovakia should belong to the East, and may also fall into older age cohorts. Thus, comparisons across (instead of within) demographic or ideological categories should be treated with caution. Finally, this analysis tells us very little about cause and effect. For example, the results show that nationalist voters read disinformation at higher rates than other voters, but they do not show that disinformation causes nationalist voting. The alternate possibility is that nationalist voters seek out disinformation news sources because these sources confirm their pre-existing beliefs. Both of these possibilities may also be true: nationalist voters seek out disinformation, and disinformation further reinforces their nationalist orientations. In other words, the causal arrow runs in both directions. However, as I discuss below, even within ideologically aligned audiences, the effects of disinformation are limited.

The unintended consequences of disinformation exposure

I also investigated the extent to which false headlines, identified as Russian disinformation by the EU's East StratCom Task Force, succeeded in shifting individual attitudes. The chosen headlines focused on a common refrain in Russian disinformation narratives: the EU's propagation of liberal identity values and policies. These headlines intend to exacerbate existing ideological divisions within CEE countries, bolster the strength of Eurosceptic forces, and ultimately challenge the EU from within. Thus, to identify whether these headlines succeed in their goals, I compared anti-EU and nationalist attitudes amongst respondents who were, and were not, exposed to the disinformation.

A similar pattern emerges for both sets of attitudes. Disinformation headlines are more likely to evoke a positive backlash effect than to serve their intended purpose. When it comes to stances towards the EU, false headlines actually make individuals more pro-EU.

Even amongst consumers of disinformation, false headlines evoke a null, rather than negative, effect. Disinformation headlines reduce nationalist attitudes even more consistently, and they do so regardless of whether a respondent is or is not a consumer of disinformation.

Thus, this report shows that the effects of Russian disinformation on attitudes are not so clear cut. While the goal of disinformation is to undermine support for the EU and inflame nationalist attitudes, it may actually do the opposite. This counter-intuitive effect is largely a product of underlying factors that determine whether individuals consume and trust disinformation. As we saw in the case of EU attitudes, exposure to false news headlines did not shift the stances of disinformation consumers. The lack of an attitudinal shift could be a product of having seen similar disinformation before, or of pre-existing, entrenched anti-EU attitudes. In all likelihood, both of these factors probably play a role. When it comes to individuals who do not read or trust disinformation, we see the aforementioned backlash effect: exposure to inflammatory or embellished false headlines leads to a pro-EU attitudinal shift.

This backlash effect is even more consistent when it comes to nationalist sentiment. False headlines reduce nationalist attitudes regardless of disinformation consumption patterns. This finding is likely explained by an aversion to sensationalist disinformation content. Even those who trust disinformation may not hold attitudes or beliefs that accord with some of the more sensationalized or prejudiced messages delivered through disinformation. Thus, reading these more extreme messages may actually reduce nationalist sentiment. Alternatively, social desirability could explain the results. Individuals, when primed to consider identity-related issues, do not want to voice their true feelings. Whereas, when individuals are not primed by identity-related topics, as in the ***Control group,*** they answer more honestly. Again, it is likely that both of these factors are at play.

Finally, I considered the extent to which exposure to news headlines influenced individu- als' willingness to combat disinformation. Neither true nor false headlines shifted individual willingness to sign petitions asking the Slovak government or Facebook to create commissions to investigate disinformation and hoaxes. However, high overall rates of support for investigatory actions – over 50 percent for the government, and over 45 percent for Facebook – are reassuring.

Conclusions

Two limitations of this study warrant discussion. First, this analysis is very much about proximate, immediate reactions to Russian disinformation. The experiment showed respon- dents disinformation headlines and then asked them how they felt about a host of issues. While measuring these immediate reactions is important for understanding disinformation's effects, disinformation also shapes attitudes through less drastic exposures that occur over longer periods of time. Disinformation can be hidden amongst true news stories in Facebook feeds, shared by family members, heard while grocery shopping, and reinforced by other true or false narratives in mainstream media. Understanding all of these pathways is important for determining how disinformation shapes individual attitudes. Second, given the immediate nature of the study, results could depend on the specific disinformation headlines used. In other words, if respondents were exposed to different sets of headlines, their reactions may have also differed. However, the number of headlines used and their content mitigate against this concern. Employing four different headlines, instead of just one or two, helps alleviate any potential bias introduced by a single headline. Moreover, the headlines were chosen as representative, but not overly sensationalist, examples of Russian disinformation. In other words, they can be considered "typical" disinformation.
Ultimately, this report shows that the effects of disinformation are not so clear cut. First, we have to consider the size of disinformation consumer market – is 15 percent readership of disinformation news outlets a significant percentage of the population? Perhaps, but when presenting this number, it is also important to note that this 15 percent is primarily composed of those who already hold attitudes that accord with disinformation content – e.g., nationalists with geopolitical stances that orient towards Russia. Nonetheless, results also indicate that disinformation consumption increases with age. Thus, interventions would do well to focus on increasing media literacy amongst the elderly.
Second, the report shows that trust in disinformation news sources is, in almost all cases, lower than in the mainstream media. However, there is some variation in trust based on the specific disinformation outlet. Outlets with a professional appearance and that blend true reporting with disinformation are particularly likely to be trusted by the public. Thus, media sources like *Hlavné Správy,* which seeks to appeal to a wide audience, are a larger concern than fringe outlets like *ExtraPlus.*
Finally, this report argues for the utility of experimental studies for studying disinformation. By using experiments, scholars and policy makers can identify the causal effects of disinformation exposure. And as we saw in this report, these effects may not necessarily accord with prior expectation. Russian disinformation about EU influence in migration and other identity-related areas actually

spurred a pro-EU, anti-nationalist reaction amongst the general Slovak population. Even amongst those who regularly trust and read disinformation, false headlines had a limited effect. The immediate attitudinal effects of disinformation seem to be at odds with Russia's intended goals. Moreover, reactions to disinformation clearly depend on existing attitudes. Thus, if we wish to limit the effects of disinformation, perhaps we should focus on the underlying causes of anti-EU and nationalist attitudes.

References

Čížik, T., 2017. Russian information warfare in central Europe in *Information Warfare: New Security Challenge for Europe*. Bratislava: Centre for European and North Atlantic Affairs, pp.8-34.

Čížik, T. and Masariková, M., 2018. Cultural Identity as Tool of Russian Information Warfare: Examples From Slovakia. *Science & Military Journal,* 1(2018), pp.11-16.

Conley, H.A., Mina, J., Stefanov, R. and Vladimirov, M., 2016. *The Kremlin Playbook: Understanding Russian influence in Central and Eastern Europe*. Washington: Center for Strategic & International Studies.

Galeotti, M., 2019. The mythical 'Gerasimov Doctrine' and the language of threat. *Critical Studies on Security*, 7(2), pp.157-161.

Galeotti, M., 2017. CONTROLLING CHAOS: HOW RUSSIA MANAGES ITS POLITICAL WAR IN EUROPE. *European Council on Foreign Relations*, pp.1-18. Accessible here: https://www.jstor.org/stable/resrep21563

Henderson, K., 2017. Euroscepticism and the missing left: The Slovak case study. *Polito- logick´y časopis-Czech Journal of Political Science*, 24(3), pp.229-248.

Ižak, Š., 2019. USING THE TOPIC OF MIGRATION BY PRO-KREMLIN PROPAGANDA: CASE STUDY OF SLOVAKIA. *Journal of Comparative Politics*, 12(1), pp.53-70.

Janda, J., Víchová, V., Richter, M., Sharibzhanov, I. and Fier, J., 2017. *Overview of Countermeasures by the EU28 to the Kremlins Subversion Operations.* European Values.

Juhász, A. and Szicherle, P., 2017. *The political effects of migration-related fake news, disinformation and conspiracy theories in Europe.* Friedrich Ebert Stiftung, Political Capital Policy Research & Consulting Institute, Budapest.

Kandrík, M. and Jevčák, M., 2018. U*nprepared and vulnerable: The resilience of the Slovak republic to foreign, foremost Kremlin-led disinformation campaigns*. Strategic Policy Institute (STRATPOL). Accessible here: https://stratpol. sk/wp- content/uploads/2018/06/PP-DRI-Final.

Krekó, P., Györi, L., Milo, D., Marušiak, J., Széky, J. and Lencsés, A., M*arching Towards Eurasia: The Kremlin connections of the Slovak far-right*. Political Capital Policy Research & Consulting Institute.

Krekó, P., Györi, L. and Zgut, E., 2017. *From Russia with Hate: The Activity of Pro- Russian Extremist Groups in Central-Eastern Europe*. Political Capital Policy Research & Consulting Institute.

Mesežnikov, G. and Bránik, R., 2017. *Hatred, violence and comprehensive military training: The violent radicalisation and Kremlin connections of Slovak paramilitary, extremist and neo-Nazi groups*. Political Capital Policy Research & Consulting Institute.

Milo, D. and Klingová, K., 2017. *The vulnerability index: Subversive Russian influence in Central Europe*. Globsec Policy Institute.

Richey, M., 2018. Contemporary Russian revisionism: understanding the Kremlins hybrid warfare and the strategic and tactical deployment of disinformation. *Asia Europe Journal,* 16(1), pp.101-113.

Šuplata, M. and Nič, M., 2016. Russia's Information War in Central Europe. GLOBSEC Policy Institute.

Notes

1. The results of Globesec Trends 2021 reaffirm these results, with 11 percent of Slovaks believing Slovakia should belong to the East (second only to Bulgaria), and 47 percent considering Russia to be Slovakia's most important strategic partner (the highest percentage amongst the ten countries surveyed)

2. Perhaps the most prominent example of this is Andrej Danko, leader of the Slovak National Party (SNS). Other notable examples include Robert Fico and Luboš Blaha from Smer-SD, as well as representatives of the newly divided far right, Marian Kotleba (L'SNS) and Milan Uhrík (Republika).

3. All differences significant at $p < 0.05$.

4. All differences significant at $p < 0.05$.

5. Respondents also had a third response option – 'Somewhere In-Between" – not analyzed here.

6. I operationalize nationalist voters as those who would vote for Krestanska U´nia, L'SNS, Republika, Sme Rodina, SNS, or Vlast if parliamentary elections were to be held.

7. Both differences significant at $p < 0.05$.

8. The larger studied also involved a fourth experimental group (exposed to all eight head-lines) that will not be analyzed in this report.

9. All respondents, including those in the *Control* group were given the option of removing their data during a debrief at the end of the survey. All reported results, including the ones above, exclude respondents who opted out of the survey. In addition, all analyses presented below exclude respondents who failed a manipulation check.

10. Differences significant at $p < 0.05$.

11. The Istanbul Convention is a human rights treaty intended to combat violence against women. Ratification of the treaty is a contentious issue in many CEE states.

12. This index has a high internal consistency, with a Cronbach's score of 0.87.

13. To divide respondents, I take the mean level of trust in the four disinformation news sources (for each experimental group), and sort respondents based on whether their belief score falls below or above the mean.

14. To divide respondents, I take the mean level of readership for the four disinformation news sources, and sort respondents based on whether their readership score falls below or above this mean.

15. The index was constructed as a simple average and has a Cronbach's of 0.81.

16. Results available upon request.