Research Article

Research and Politics
April-June 2022: 1–7

DOI: 10.1177/20531680221100440
journals.sagepub.com/home/rap
SAGE

Protest and digital adaptation

Rebecca Strauch and Nils B. Weidmann

Abstract

Autocratic governments routinely interfere in digital communication technology for political purposes. However, citizens can use different technologies to bypass government interference. This article examines how political protest influences the use of anonymity-preserving digital services in autocracies. Citizens should be more likely to use these tools during high political tension because they fear governmental surveillance or censorship. The analysis combining data on the Tor anonymization network with protest event data demonstrates noticeable increases in Tor usage after days with many protest events but not days with single protest events.

Keywords

information technology, repression, protest

Over the years, autocratic governments have become more apt in the use of digital communication technologies for political purposes. Government-controlled communication can suppress collective unrest (Weidmann and Rød, 2019), for example, through government-produced content (King et al., 2017), fast and pervasive online censorship (Roberts, 2018), or the full shutdown of Internet services (Deibert et al., 2010). A common assumption in this literature is that autocratic governments enjoy a high level of control over digital communications vis-à-vis the population. However, the notion that citizens are left entirely at the mercy of the government may be premature. Different technologies now permit accessing the Internet such that the user's identity and the transmitted content is protected. Citizens living under autocratic rule may have become more skilled in using these technologies, for example, in China, where the use of VPN clients has helped them to bypass online censorship (Hobbs and Roberts, 2018). In this article, we study this "digital adaptation" in a cross-national comparison, by focusing on one particular technology, *The Onion Router* (usually referred to as "Tor"). While mostly considered to be a way to access the Darknet, there are good reasons to expect that it is also a useful device for political activists, journalists and citizens in authoritarian countries to remain anonymous on the Internet and to escape censorship.

Anonymity-granting technology such as the Tor network allows users to access the Internet safely and without tracking or surveillance. Tor is based on a network of computers run by volunteers, which redirect a user's request to visit a website over a series of intermediate servers. There are two different ways to access the Tor network, either via (1) relays or (2) bridges. When users decide to access the network through relays, the Tor browser downloads a list of accessible Tor servers that act as a random connection through which information is processed (see Figure A1 and Section 1 in the Appendix for more details). External observers can see the information entering and leaving the relay, but they cannot identify which host a user is communicating with. As the list of relays is, however, public, external observers can download the list and block respective relays in return. As an alternative, Tor offers a list of "bridges" which makes censorship from external observers less likely. Bridges are not publicly listed and Tor only distributes a small fraction of bridges for one specific day and user. This makes it much more difficult for external

University of Konstanz, Germany

Corresponding author:
Rebecca Strauch, Universitaetsstrasse 10, Konstanz 78457, Germany.
Email: rebecca.strauch@uni-konstanz.de

observers to block the chain of servers. While we primarily look at the use of relays and bridges in combination, we also examine differences between the two for the outcome of our study.

Prior research on Tor has focused on the distribution of users across different political regimes (Jardine, 2018), showing that people of illiberal countries rely on the network to secure their digital freedom. This article presents a more fine-grained analysis. We explore whether political events on the ground are reflected in day-to-day usage patterns of the Tor network, which would indicate that "digital adaption" of citizens actually occurs. We argue that anonymity-granting technology may be of particular importance during times of political unrest and collective protest. Protesters might have different intentions for using Tor around protest events. First, users may fear governmental prosecution on the Internet. It is possible to monitor Internet communication and identify particular individuals through techniques such as location tracking, online surveillance, and digital fingerprinting. During times of political unrest, autocratic governments are likely to ramp up surveillance. Anticipating this, Internet users, particularly political opponents and activists, are more likely to use anonymity-preserving tools to circumvent political prosecution on the Internet. Second, in times of political unrest, users are more likely to exchange sensitive information online. Contentious events such as protests might be of particular relevance to political activists to spread anti-governmental information on the Internet, both nationally and internationally. By this, activists are able to educate the population about governmental activities and to gain support through digital communication channels. These mechanisms all suggest that we expect to see an increase in Tor usage in the aftermath of protest events. Therefore, our main hypothesis is that *the occurrence of anti-government protests increases Tor usage in the respective country.*

Apart from protest occurrence, we take into consideration the severity of political unrest, as measured by the number of protest events and the number of participants. Many protest events are small and remain entirely local. Others, however, attract large numbers of participants, and spread to many cities in a country. The latter ones are those that should be more visible nationally, and we should see changes in Tor usage in particular as a response to these events. In addition to protest occurrence and magnitude, we argue that Internet users are more likely to turn to the Tor network when the government cracks down protesters with physical intervention. If activists experience repression, they are more afraid of the government and should thus be more likely to turn to Tor. Hence, we expect that *Tor usage in a country increases in particular after larger episodes of anti-government protest (in terms of event numbers and participants), and after physical governmental intervention.*

Finally, we should also see variation in protest-driven digital adaption across regimes. We can assume that the use of Tor should be particularly likely in those countries where autocratic governments tightly restrict access to digital communication channels and are actively influencing their content. Therefore, *the effect of protest on the use of anonymity-granting technology should be most pronounced in those countries with high restrictions of digital communication.*

Research design

In order to test our hypotheses empirically, we use a sample of authoritarian countries following the definition by Geddes et al. (2014) and Coppedge et al. (2019).[1] These countries are observed daily, from September 2011 to December 2018. The period of analysis is limited as data on protest events was not available after 2018 by the time we completed the study, while coverage of the Tor usage data does not start until the fall of 2011.[2]

Data

To obtain data on the use of anonymity-granting services, we rely on the Tor Metrics project (2020b). The project gathers data on the use of the Tor network and makes these estimations publicly available at the level of countries and with daily resolution. Due to the anonymity-preserving structure of the Tor network, it is not possible to count individual users; rather, Tor Metrics counts requests by clients who access the Tor network. An average number of users is then calculated from this estimate under the assumption that each client makes 10 directory requests per day. In order to identify the country from which the Tor network is accessed from, Tor Metrics resolves the Internet addresses of relay and bridge users to country codes.

For our main analyses, we use *total Tor usage*, which is the sum of *relay* and *bridge usage*. This variable varies significantly across the countries. For example, in Russia, some days see as much as nearly 400,000 Tor users per day, while in Cuba, the number never exceeds 650 users. We take this distribution into account by log-transforming this variable, but also by including a series of fixed effects in our models that (among others) net out country-specific levels of Tor usage over time.[3]

Figure 1 shows the evolution of Tor usage over time across the countries in our sample.[4] As the Tor browser is easy to install on a computer, also less technically-inclined users can benefit from this technology. With an overall growth of Internet penetration and an increase in Internet users, the Tor network is expanding worldwide. While the number of Tor usage increases over time, we observe an outstanding peak in the network in September 2013. Around this time, the Tor network suffered from a tremendous influx

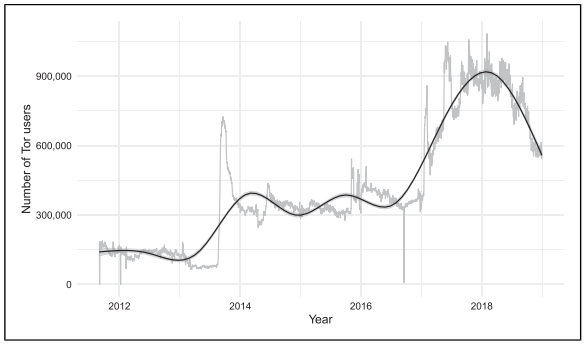

Figure 1. Tor usage across authoritarian countries in our sample over time. The gray line displays absolute daily Tor usage over time and the black line depicts the moving average.

of new clients which was associated with a coordinated attack trying to swamp the network with requests to slow it down (Tor Project, 2020a). We address potential problems due to this sudden increase in our analysis below.

Data on political protest is provided by the "Mass Mobilization in Autocracies Database" (MMAD, Keremoğlu et al., 2020; Weidmann and Rød, 2019, Ch. 4). The dataset includes event-level data on protest events in autocratic countries on a daily basis, and covers the years 2003–2018 (Version 3). Data is coded from newspaper articles published by the Associated Press, the Agence France Presse, and BBC Monitoring. In line with our theoretical framework, we limit our main analysis to anti-government protests. Our analysis operationalizes political protest in different ways: First, we use a binary variable to indicate protest occurrence on a particular day in the given country. Overall, protest is a rare occurrence, and only about 4% percent of our observations have a protest occurrence coding of 1. Second, in line with our hypotheses, we use a set of more nuanced variables to capture the size, in terms of number of events and number of participants, and the level of physical intervention. We distinguish between country-days with single protests and those with more than one protest, those with a low (<250 participants) and a high participation (250 or more) overall, as well as country-days where protest was repressed with physical intervention (where protesters were injured or killed, as coded in the MMAD) or with non-physical intervention (where security forces were only present) (Table A1 in the Appendix). In all these analyses, the base category is *no protest*. It is important to mention that information about protest participation and physical intervention is missing from many events in the MMAD, which leaves us with even fewer numbers of events in our study as compared to those analyses where only event occurrence or counts are included.

In addition, to examine variation in our main analysis across different types of autocratic countries, we include data on the extent to which digital freedoms are constrained in a country. The Digital Society Project (Mechkova et al., 2019) collects data on digital media freedom and for the purpose of our analysis, we rely on the indicators "government Internet filtering in practice" and "government social media censorship in practice." Higher values of these indicators correspond to more digital media freedom and less censorship. To facilitate interpretation, we invert these variables: In our models, higher values indicate higher levels of filtering/censorship.

Results

We test our hypotheses by applying OLS regression models with log-transformed dependent variables to estimate the effect of protest on Tor usage.[5] The unit of analysis is country-day and we lag our predictor variables by one day. We also include a lagged dependent variable to take into account serial correlation in Tor usage.[6] Particular weekdays may be preferred for protest, which is why we include weekday fixed effects in all our models. To account for country-specific and temporal trends in Tor usage, we first include separate country and year fixed effects. Since Tor usage may follow country-specific temporal trends, we also estimate models with country × year fixed effects.

Protest occurrence and number of protests

Table 1 shows the results for the first set of models with a binary predictor (protest 0/1). In Model 1 with separate country, year, and weekday fixed effects, protest occurrence has a positive and statistically significant effect on Tor usage. This result holds with a more strict estimation (country × year fixed effects, Model 2). The overall effect is small; the mere occurrence of protest increases total Tor usage by 0.5% in Model 1. Not surprisingly, previous levels of Tor usage explain much of the current usage, as indicated by the large coefficient of the temporal lag and the high R^2.

In Models 3 and 4, we distinguish country-days with a single protest ("one protest") from larger protest movements with more than one event ("multiple protests"). In these models, the occurrence of more than one protest increases total Tor usage by 1.3%, and this effect is significant at the 1% level (Model 3). Including country × year to weekday fixed effects, the positive and significant association still holds (Model 4) and remains comparable in magnitude to the previous model. In both models, the occurrence of only one protest, however, shows no significant effect. We can conclude that the overall magnitude of protest activity matters; multiple protests in a country in one day, for example, in various cities, are necessary for users to feel the need to use anonymity-granting browsers when accessing the Internet.[7]

According to our theoretical discussion, we also expect to see more pronounced changes in Tor usage if protest events are attended by a higher number of participants, and if the government responds with physical intervention. In

Table 1. Effect of protest occurrence (Models 1 and 2) and number of protests divided into one and multiple protests (Models 3 and 4) on total Tor usage, estimated with country, year and weekday FEs (Models 1 and 3) and country × year and weekday FEs (Models 2 and 4).

	Dependent variable			
	Total Tor usage (log)			
	(1)	(2)	(3)	(4)
Total Tor usage (log) (t-1)	0.910** (0.001)	0.799** (0.001)	0.910** (0.001)	0.799** (0.001)
Protest (t-1)	0.005** (0.002)	0.003* (0.002)		
One protest (t-1)			0.003 (0.002)	0.001 (0.002)
Multiple protests (t-1)			0.013** (0.003)	0.011** (0.003)
Constant	0.298** (0.004)	0.623** (0.012)	0.298** (0.004)	0.623** (0.012)
Weekday FEs?	Yes	Yes	Yes	Yes
Country FEs?	Yes	No	Yes	No
Year FEs?	Yes	No	Yes	No
Country × year FEs?	No	Yes	No	Yes
Observations	179,739	179,739	179,739	179,739
R^2	0.984	0.985	0.984	0.985
Adjusted R^2	0.984	0.985	0.984	0.985

Note: *p < 0.05; **p< 0.01.

Table 2. Effect of protest occurrence on total Tor usage, interacted with Internet filtering in general (Model 5) and social media censorship (Model 6), and estimated with weekday FEs.

	Dependent variable	
	Total Tor usage (log)	
	(5)	(6)
Total Tor usage (log) (t-1)	0.991** (0.0003)	0.991** (0.0003)
Protest (t-1)	0.005 (0.004)	−0.002 (0.004)
Internet filtering	0.001** (0.0002)	
Social media censorship		0.0002 (0.0002)
Protest × internet filtering (t-1)	0.001 (0.001)	
Protest × social media censorship (t-1)		0.004** (0.001)
Constant	0.020** (0.001)	0.023** (0.001)
Weekday FEs?	Yes	Yes
Observations	173,896	173,896
R^2	0.983	0.983
Adjusted R^2	0.983	0.983

Note: *p < 0.05; **p< 0.01

Table A2 in the Appendix, we recode our protest variables such that they distinguish between days with a few number and several participants, and days when the government intervenes physically or non-physically. For both distinctions, however, we see no substantive results on a change in Tor usage[8].

The role of digital freedom

We also expect the effect of protest activity on Tor usage to vary with the overall level of digital control exerted by the government. In particular, the effect should be stronger in those countries where digital freedom is more severely restricted, and where citizens therefore are more likely to be pressured into the use of anonymity-granting technologies. To test this claim, we estimate the effect of the binary protest occurrence variable again, interacted with the *Internet filtering* and *social media censorship* variables from the Digital Society Project.[9]

Table 2 shows the results of these interaction models, and Figure 2 visualizes the estimated coefficients for Tor usage across the range of the two moderator variables. In line with

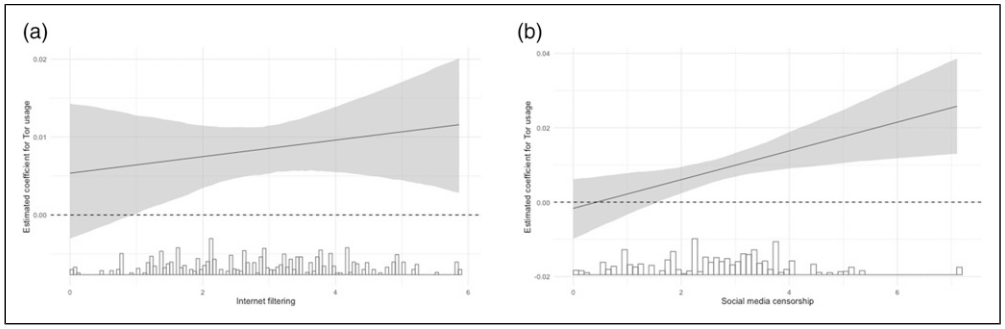

Figure 2. Effect of protest occurrence on total Tor usage across the range of Internet filtering (left) and social media censorship (right) at the 95% confidence interval. The rugs show the distribution of the mediator variable.

our expectations, the effect of protest on Tor usage is generally higher in countries with high levels of digital control. According to Figure 2, the effect of protest on Tor usage doubles when comparing countries with low and high levels of Internet filtering (left panel), and increases by a similar order of magnitude when comparing countries with low and high social media censorship (right panel). These findings support our initial expectation that users are more likely to switch to the Tor network as a result of protest if they live in countries with high levels of digital control and censorship. This effect is more pronounced for social media censorship as compared to Internet filtering in general, even though we observe a positive tendency also for the latter.

Robustness checks and placebo tests

We conduct a number of robustness checks and placebo tests to support our main results. Detailed results are provided in the online Appendix. With regard to robustness checks, we find that our results hold when we use a reduced sample, in which we address the extreme peak in Tor usage and missing values (Table A9), or when we use alternative protest data (Table A10). We also check for reverse causality by switching the independent and the dependent variable, which shows no effect (Table A11). Furthermore, we include a set of placebo tests. First, we show that the timing of *anti-regime* protest matters when it comes to an increased notion of Tor usage. If we run a Monte-Carlo analysis that shifts the occurrence of these protests to other dates in the same year, we observe no such increase. As shown in Figure A4 in the Appendix, the coefficient values we estimate from these Monte–Carlo dataset are centered around zero, such that our true model coefficients end up in the upper 1% of these distributions and are thus unlikely to be the result of chance. We further show that events that should be unrelated to an increase in Tor—*pro-regime protests* (Table A12) and natural disasters (Table A13)—exhibit no positive and significant relationship to Tor usage.

Conclusion

The use of anonymity-granting technology on the Internet is often discussed in the relation to the Darknet and criminal cyber-activities. In this article, we show that Tor usage patterns are also partly determined by political events on the ground. Using daily observations of protest activity and Tor usage, we show that protests are positively related to the use of anonymity-granting technology. Although the effect is small and only visible for larger protest episodes across the country, it appears that major contentious events make users more likely to employ this technology, allowing them to access blocked content, bypass censorship, or avoid digital surveillance. We also examine how the effect plays out in different countries, depending on the level of digital censorship employed by the government. Our results show an increase in Tor usage in the aftermath of protest events in particular in those countries where social media content is censored, lending further support to our conjecture.

Although our analysis is the first cross-national, fine-grained exploration of the political motivations for Tor usage, it faces several theoretical and technical problems. Tor Metrics allows us to obtain data on network usage but no information on who uses Tor primarily. This means that we cannot attribute Tor usage to particular groups of users (e.g., citizen activists afraid of repression). This type of attribution problem is common in the research of digital interference (Keremoğlu and Weidmann, 2020), making it difficult to infer which actor uses a particular service or technology. For similar reasons, our analysis is limited to the total volume of traffic per country and day, making it impossible to infer whether an increase was due to a higher number of users, or due to existing users generating more traffic.

The modest size of the effect also shows that a technology such as Tor has not become a mass phenomenon. Our results are consistent with increases in Tor usage driven by a small number of activists treading carefully in their online communications. At the same time, the majority of

citizens currently does not adopt this technology, leaving them exposed to governmental interference. This is not because users are in principle not able to change the way in which they navigate the Internet and switch to anonymity-preserving technology, as existing research has shown (Hobbs and Roberts, 2018). Rather, one explanation could be that the vast majority of citizens might be unaware of the degree of governmental interference, which is likely to affect their motivation to adopt this technology in the first place. Still, even if mass adoption is not a likely outcome, the demand among political activists for anonymity-preserving technology in autocracies may increase. This mirrors similar trends during contentious episodes in democracies. Overall, our study shows that digital adaptation happens, and that citizens are able to fight back against governments interfering with digital communication—even if their numbers remain small.

Acknowledgments

We would like thank the two anonymous reviewers, Clara Neupert-Wentz and the participants at the "Empirical Peace and Conflict Research Workshop" 2020 at the Hertie School of Governance and at the "Digital Democracy Workshop" 2020 at the University of Zurich for comments and suggestions.

Author's Note

The study was preregistered on February 17, 2020 at https://osf.io/84ys9. The interaction effects tested in Models 5 and 6 were not included in the pre-registration and were added later.

Declaration of conflicting interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author(s) disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: This study was funded by the Deutsche Forschungsgemeinschaft (DFG), project number 402127652.

ORCID iD

Rebecca Strauch https://orcid.org/0000-0001-8142-3694

Supplemental Material

Supplemental material for this article is available online.

Notes

1. Autocracies are identified as countries where an executive has come to power through undemocratic methods, which includes anything except direct, fair and competitive elections. This could also mean that competition in elections was limited or prevented by the military.
2. In the pre-registration, the temporal coverage of the Tor data was reported to start in March 2011. This initial measurement was simply a short test run, and comprehensive coverage is not available until September 2011. Hence, the period of analysis has been adjusted accordingly.
3. In addition to *total Tor usage*, we also run separate analyses with *relay* and *bridge usage*, since *bridge usage* is less vulnerable to governmental surveillance and may therefore be more affected by contentious events (see Appendix).
4. The evolution of *relay* and *bridge usage* is documented separately in Figure A2 in the Appendix.
5. With c being Tor usage as provided in our data, we use $log10(c + 1)$ as the dependent variable.
6. We assess Durbin–Watson tests and ACF plots in Sec. 2.2 in the Appendix to show that we effectively remove serial correlation at first and several orders by introducing the lagged dependent variable.
7. Results on the effect of protest occurrence and number of protests on relay and bridge usage are recorded separately in the Appendix (for reference see Tables A3 and A5). While total Tor usage mainly consists of relay users, bridge usage is comparatively small. In the pre-registration, we expected that protest would have a stronger effect on *bridge usage* in comparison to *relay usage* due to the structure of the network but we failed to find such an effect.
8. Tables A4 and A6 in the Appendix show the results for the effect of protests with many participants and those which were physically repressed protests separately for *relay* and *bridge usage*.
9. Since these indicators are available only at an annual resolution and exhibit almost no variation over time, we drop the country and year fixed effects from these models but keep the weekday fixed effects.

References

Coppedge M, et al. (2019) *V-Dem Codebook v9: Varieties of Democracy (V-Dem) Project*.

Deibert RJ, Palfrey JG, Rohozinski R, et al. (2010) *Access Controlled: The Shaping of Power, Rights, and Rule in Cyberspace*. MIT Press.

Geddes B, Wright J and Frantz E (2014) Autocratic breakdown and regime transitions: A new data set. *Perspectives on Politics* 12(2): 313–331.

Hobbs WR and Roberts ME (2018) How sudden censorship can increase access to information. *American Political Science Review* 112(3): 621–636.

Jardine E (2018) Tor, what is it good for? Political repression and the use of online anonymity-granting technologies. *New Media & Society* 20(2): 435–452.

Keremoğlu E and Weidmann NB (2020) How dictators control the Internet: A review essay. *Comparative Political Studies* 53(10–11): 1690–1703.

Keremoğlu E, Hellmeier S and Weidmann NB (2020) *Coding Instructions for the Mass Mobilization in Autocracies Database, version 3.0*. University of Konstanz.

King G, Pan J and Roberts ME (2017) How the Chinese government fabricates social media posts for strategic distraction, not engaged argument. *American Political Science Review* 111(3): 484–501.

Mechkova V, Pemstein D, Seim B, et al. (2019) *Measuring Internet Politics: Introducing the Digital Society Project*. http://digitalsocietyproject.org/data/

Roberts ME (2018) *Censored: Distraction and Diversion inside China's Great Firewall*. Princeton University Press.

Tor Project (2020a) Serious network overload. https://blog.torproject.org/tor-weekly-news-september-4th-2013 (Online; accessed 16-October-2020.

Tor Project (2020b) Tor Metrics. https://metrics.torproject.org (Online; accessed 13-February-2020.

Weidmann NB and Rød EG (2019) *Oxford Studies in Digital Politics*. New York: Oxford University Press.The Internet and political protest in autocracies.