



Article

Tor, what is it good for? Political repression and the use of online anonymity- granting technologies

new media & society
2018, Vol. 20(2) 435–452
© The Author(s) 2016
Reprints and permissions:
sagepub.co.uk/journalsPermissions.nav
DOI: 10.1177/1461444816639976
journals.sagepub.com/home/nms



Eric Jardine

Centre for International Governance Innovation (CIGI), Canada

Abstract

Why do people use anonymity-granting technologies when surfing the Internet? Anecdotal evidence suggests that people often resort to using online anonymity services, like the Tor network, because they are concerned about the possibility of their government infringing their civil and political rights, especially in highly repressive regimes. This claim has yet to be subject to rigorous cross-national, over time testing. In this article, econometric analysis of newly compiled data on Tor network usage from 2011 to 2013 shows that the relationship between political repression and the use of the Tor network is U-shaped. Political repression drives usage of Tor the most in both highly repressive and highly liberal contexts. The shape of this relationship plausibly emerges as a function of people's opportunity to use Tor and their need to use anonymity-granting technologies to express their basic political rights in highly repressive regimes.

Keywords

Dark Web, online anonymity-granting technologies, privacy, anonymity, political repression, quantitative methods, Tor

People often want to remain anonymous on the Internet. Yet, simply firing up a browser like Google Chrome or Mozilla and surfing the web does not provide any real anonymity. Internet service providers (ISPs), malicious programs, and state agencies can all pinpoint

Corresponding author:

Eric Jardine, Centre for International Governance Innovation (CIGI), 67 Erb Street West, Waterloo, ON N2L 6C2, Canada.

Email: ejardine@cigionline.org

who is doing what online. Anonymity-granting technologies exist that can counter this trend by masking a person's identity while they surf the web, host websites, and express their political views. Such services are functionally neutral. People can use them to either cover their tracks as they commit illegal activities or to avoid state/corporate censorship and surveillance (or some combination of both). Activists often state that such tools are used by political dissidents who stand up to repressive regimes, but these claims lack a broad, cross-national empirical basis (Goodin, 2014). Without a solid empirical foundation, it is unclear if anonymity networks are used consistently by political dissidents in highly repressive contexts. All of this begs the question, "does the level of political repression in a country incite the use of online anonymity-granting technologies?"

Data on use of the Tor network from 2011 to 2013 suggest that political repression does drive usage of anonymity-granting technologies. The results indicate that both very high and very low levels of political repression tend to drive use of Tor the most. In other words, the relationship between a country's level of repression and the rate of individual usage of anonymity-granting technologies is U-shaped.

A single cost-benefit logic underpins these results. The nature of a political area shapes the incentives that actors face. In highly repressive regimes, people need to protect their identities online to circumvent censorship and to avoid surveillance. As a result, people in such contexts tend to use a lot of anonymity-granting technologies. In highly liberal regimes, people have a lot of opportunity to use anonymity-granting technologies. As a result, individuals in such countries can freely use a lot of anonymity-granting technologies. People in regimes with a mixed amount of political repression lack both the severe need and the clear opportunity to use anonymity-granting technologies, stifling usage rates.

The remainder of this article is structured as follows. The section "Seeking anonymity online: political need and opportunity" elaborates upon a theory explaining the relationship between political repression and the use of online anonymity-granting technologies. The section "The data" details the operationalization of the variables and provides descriptive statistics. The section "Empirical Analysis" presents the results of the statistical analysis and discusses their fit with the proposed theory. The last section "Conclusions" concludes with a discussion of policy recommendations and areas of future research.

Seeking anonymity online: political need and opportunity

People's behaviors are often structured by the political arena within which they find themselves. As shown below, people's choice to use anonymity-granting technologies is a function of two structural factors that are, themselves, conditioned by the level of political repression in a country, namely, *political need* and *opportunity*. The plausible relationship between repression and anonymity-granting technologies is U-shaped.

In some ways, the current question about the connection between political repression and the use of anonymity-granting technologies online is similar to research on the relationship between a country's regime type and the state's use of political repression. Both involve the idea that the nature of the political arena in which people operate affects the behavior of both states and dissidents. Some early studies found a linear relationship between political repression and regime type (Henderson, 1991). Most recent studies

qualified the result, finding that the relationship between these variables is shaped like an inverted U (Fein, 1995; Muller, 1985; Regan and Henderson, 2002). Dissent and political repression are lowest in both highly authoritarian and highly democratic contexts.

The U-shaped relationship emerges because of the incentive toward political protest provided by different political contexts. In highly repressive regimes, the organization of political dissent is prohibitively costly for individuals and also likely to fail, so dissidents do not even attempt collective action. Under these conditions, the state is rarely compelled to repress protests. In highly liberal regimes, the ability to organize is legally protected so individual participation in collective action is fairly likely. At the same time, there are numerous democratic, nonviolent pathways to political change, so dissent rarely reaches a level where political repression is used against protestors (Muller, 1985).

Mixed regimes, in contrast, cannot deter people from collective action through constant surveillance and the promise of severe collective punishment because they are too internally divided between authoritarian and democratic elites to act decisively one way or the other. Such indecision reduces the individual-level costs of organization. At the same time, mixed regimes lack the institutional channels to effectively respond to political demands in a nonviolent way, causing protests to turn violent and making repression a more likely option.

Intuitively, how political dissidents behave in repressive, mixed or free political contexts might be similar both online and offline. In highly repressive regimes, the danger of getting caught using the Internet to circumvent censorship or avoid surveillance is so high that people might choose to completely forgo expressing dissent, which means they will not use anonymity-granting technologies. In highly liberal contexts, people should be able to express their political views openly without fear of reprisal, so dissent should be expressed but not in tandem with the use of effectively superfluous anonymity-granting technologies. In mixed regimes, individuals should want to express political dissent and do so via anonymity-granting technologies because political opposition is a potentially costly endeavor. As a result, expressing dissent without taking steps to protect one's identity could be personally costly and the use of anonymity-granting technologies should be highest in mixed political contexts. In other words, there are some reasons to think that the relationship between a repressive context and the use of online anonymity-granting technologies is also shaped like an inverted U.

However, the literature on regime type and political repression is a poor fit as an explanation of the connection between political repression and individual usage of anonymity-granting technologies for two reasons. First, the regime type/repression theory explains the state's use of political repression, which is essentially a reaction (at times proactive) to organized political dissent. While the mechanisms underlying the theory do tentatively suggest when people are likely to organize for collective action and when they are not, the ultimate outcome to be explained is state repression, not political protest. In this sense, the theory of regime type and repressive violence explains the wrong thing (state repression, not the use of online anonymity-granting technologies) and explains it at the wrong level (state actions, not the actions of individuals).

Second, the individual-level mechanism driving people to protest in particular ways is not well specified in the regime type/political repression model. Nominally, the costs and benefits of political protest drive people to act. It is not clear, however, what drives

people to undertake one type of action compared to another. In this sense, the regime type/repression model cannot draw a distinction between when people will express political opinions offline, online, or online with the cover of anonymity-granting technologies. These two deficiencies of the regime type/repression theory suggest that it is in fact ill-suited to explain the occurrence of anonymous online activity.

It is far more plausible that individuals choose to use online anonymity-granting technologies due to the interaction of two structural-level factors: *political need* and *opportunity*. These factors interact to determine the level of anonymity-granting technologies that a person will use. *Political need* indicates the benefits that anonymity-granting technologies convey onto an individual. Political need is positively related to the level of political repression in a country. In highly repressive contexts, where people are very likely to suffer severe personal costs for voicing a political opinion or viewing censored content online, the increased privacy that anonymity-granting technologies provide bestows significant benefits onto individual users. This intuition suggests that as political repression rises, so does the level of political need and this should lead to more use of anonymity-granting technologies.

Opportunity is the chance that people have to use anonymity-granting technologies free from state sanction. Opportunity has two components, namely, access to the technology and the forms that state counteraction take in response to its use. Both reflect the cost of using anonymity-granting technologies. Opportunity should vary inversely with the repressiveness of a political context. Repressive regimes provide less of an opportunity to use anonymity-granting technologies by restricting both their effectiveness and the ability of people to access them without sanction. For example, China makes a concerted effort to block the use of the Tor anonymity network (*MIT Technology Review*, 2012). Likewise, Russia recently offered a US\$110,000 bounty to anyone who could crack the anonymity of the Tor network (*BBC World News*, 2014). Some authoritarian regimes even view possessing any encryption technology with deep suspicion, treating it as either an illegal offense or as an action that requires special licensing (Princeton's Office of Information Technology, 2014). Liberal political contexts, in contrast, typically allow the use of all sorts of encryption and anonymity-granting technologies, ranging from the Tor network to encrypted messaging services such as Wickr. Indeed, many such technologies are developed and hosted in the Western world. The largest single funder of the Tor Project, for example, is the US government. Since there are only the most minimal of costs associated with the use of anonymity-granting technologies in liberal contexts, the use of such technologies should be more widespread, everything else being equal.

In the case of both political need and opportunity, a clear distinction needs to be made between the way in which the political arena shapes behavior and the precise motive of individuals who choose to use anonymity-granting technologies. The arena has a population-wide effect on the incentive to use anonymity-granting technologies. The motives that drive people to actually use the technology are variable, however, with different people undertaking the same action (using Tor, for instance) for different reasons.

The high level of opportunity found in highly liberal contexts frees people to use anonymity-granting technologies for reasons related to their own personal preferences. Some people in highly liberal contexts might, for instance, use anonymity-granting technologies because they want to undertake illegal activities on the Dark Web. As an

example, a recent study by Gareth Owen and Nick Savage (2015) finds the disturbing result that over 80% of the site visits on the Tor hidden services—that is, Dark Web websites that are only accessible by use of the anonymity-granting Tor browser—go to child abuse sites.

Other users in liberal democracies might want to use anonymity-granting technologies in order to cover up for socially unacceptable behavior, particularly if those activities crossed heteronormative lines. Others might care about online privacy and could use systems like Tor because they no longer trust that the Internet is a private place in the wake of Edward Snowden's revelations (Hampson and Jardine, 2016). The motives of others might even be altruistic. Those viewing anonymity as a right could use the network to create more traffic, effectively making anonymity-granting technologies more effective for those that really need them.

The low opportunity (high cost) and high political need (benefit) to using anonymity-granting technologies in highly repressive countries again incentivizes the use of Tor but again does not reveal the precise motives of the people that turn to the network. Some people in a repressive country might use the technology to cover up for their illegal activity, be it politically oriented or otherwise. For many others in repressive regimes, imposed restrictions on access to online content and pervasive surveillance could lead people to routinely use anonymity-granting technologies to circumvent state censorship and to avoid the prying eyes of governments.

While political need might be high, the opportunity to use anonymity-granting technologies in repressive regimes should stifle usage rates to some extent. In terms of censorship circumvention, for example, one problem with programs like Tor is that they are slow and have limited browser plug-in capabilities, limiting their functionality. In terms of dodging surveillance, uptake of anonymity-granting technologies might be truncated in counties that actively police the Internet. These programs often use distinct encryption, so in some cases even using the technology at all could attract the unwanted attention of the state.

In short, while individual motives certainly vary, it is plausible that the various motives can be lumped into two categories: opportunity-based motives and political need-based motives. Opportunity-based motives, as broad as they are, should be most common in highly liberal regimes, while political need-based motives should be most common in highly repressive countries.

Together, the general expectation is that these two variables (opportunity and political need) should interact to determine the level of anonymity-granting technologies used in a given country. *In highly repressive contexts*, we would expect that individuals will have a high need to use anonymity-granting technologies online. At the same time, their opportunity to use such technologies should be fairly low because the government may try to block usage and prosecute severely those that are found using the technology. *In a context with low levels of political repression*, the need to use anonymity-granting technologies should be low, since governments are restrained by laws, norms and checks and balance and are unlikely to commit wholesale abuses of their citizenry. In contrast, the opportunity to use such technologies in liberal contexts should be very high because the cost of obtaining and effectively using such technologies is so low. *In contexts with medium levels of repression*, both need and opportunity

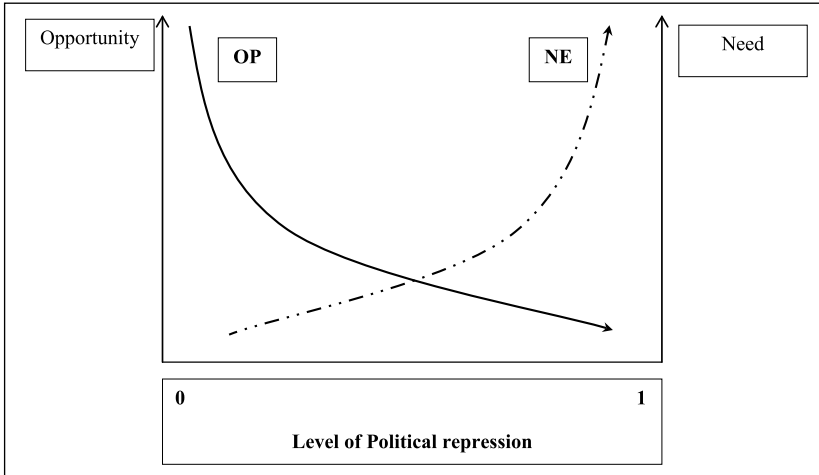


Figure 1. The relationship between political repression, need, and opportunity.

should have a low-to-medium value, falling somewhere in between the values on the two extremes. In other words, the relationship between political repression and anonymity-granting technology usage should be shaped like a U, with opportunity being the overarching driver for use in liberal countries and political need being the main driver of usage in repressive countries.

Figure 1 outlines how the theory predicts low, medium, and high levels of political repression should condition the opportunity and the political need to use anonymity-granting technologies. The level of political repression runs along the x-axis. It ranges from very low at origin, 0, to very high, 1, at the far right-hand side of the x-axis. The primary y-axis plots the level of opportunity that people have to use anonymity-granting technologies. The secondary y-axis plots the level of political need that people have to use anonymity-granting technologies. The opportunity curve (OP) starts out high and declines as political repression increases. The political need curve (NE) starts out low and increases as the level of political repression rises.

The proposed theory suggests that we would expect two things empirically. First, the relationship between political repression and use of anonymity-granting technologies should be shaped like a U, driving usage the most in both highly liberal and highly illiberal contexts. Second, the underlying mechanism leading people to use Tor should shift from the nearly costless ability (opportunity) to do so in liberal democracies to the need to do so in order to exercise basic political rights in repressive regimes. The remaining sections of the article demonstrates empirically that the first of these two expectations is valid. The other requires further investigation with different data sources.

The data

This section details the operationalization of the dependent variable (use of anonymity-granting technologies). It also discusses the core explanatory variable—political repression—and

outlines the relevant controls. The data are structured into an unbalanced panel, with country year observations from 2011 to 2013 for 157 countries.

Dependent variable: anonymity-granting technologies

Data on Tor anonymity network usage are a good proxy for the level of individual usage of online anonymity-granting technologies. The Tor network, which stands for The Onion Router, is one of the most prominent system for obtaining anonymity online, routinely reaching over 2 million daily users in 2014 (Tor Project, 2014b). Using data on actual Tor usage rates are also valuable as it reveals the number of times that computers connected to the network and thus does not rely upon self-reporting, which could be biased if people thought admitting to using Tor might suggest they were up to something legally or normatively illicit.

The Tor network is based upon a bank of volunteered computers. Anyone can volunteer their system to be a part of the Tor network. The initial funding for the Tor Project came from the US government, which wanted to develop an online anonymity-granting technology for use by dissidents in repressive regimes. According to the Tor website, users of the network range from the military to political activists and normal Internet users, although it is not clear how exactly Tor knows who is using the technology without relying upon after-the-fact disclosures (Tor Project, 2015).

The programming underlying the Tor network is dynamic and supported by an active team of people who fundamentally believe that online anonymity is a basic right. At the same time, regimes (liberal regimes included) often try to “break” the Tor network. This creates a cycle of action and reaction, with moves to unmask Tor users being met with efforts to provide continued anonymity (Jardine, 2015). This cycle of competition actually tends to make the system more robust over the long term, while also allowing law enforcement to apprehend criminals in the short term.

Tor works by routing a person’s web queries to a website via a series of intermediate computers, which limits the ability of governments and private companies to determine a person’s Internet Protocol address and by extension their physical identity and location. The Tor network is composed of two elements: relays and bridges. The Tor network passes an Internet signal through at least three relays before it dumps the data onto the Internet. The first two computers act as intermediate relays between the origin and the destination of the query. The final computer is an exit relay, which deposits the query onto the actual Internet. The Internet Protocol addresses of exit relays are visible to websites, companies, and governments. All relays points are publicly known (although what route a particular packet takes is randomized). Most users directly connect to relays in the network.

Bridges are a way to access the Tor network via a hidden route. Bridges are especially important in regimes that attempt to prevent people from using Tor services by blocking known Tor relays. Bridges are a better censorship circumvention tool than relays because they are not published publicly, making it harder for a regime to prevent people from using them (Tor Challenge, 2014).

The two separate ways of accessing Tor—directly via relays or indirectly via bridges—suggest three potential ways to measure how people use the network. The first measure

is the rate of Tor bridge usage per country per year per 100,000 Internet users. The second measure is the rate of Tor relay usage per country per year per 100,000 Internet users. The final measure is the sum of both Tor bridge and Tor relay usage per country per 100,000 Internet users per year. Importantly, a single person can use the network multiple times in a year and, as such, the rate of usage per 100,000 people can be greater than 100,000.

Data on all these measures are available from the Tor Project (2014b) on a daily basis. Aggregating the data into a yearly structure for each country makes the measure of the dependant variable compatible with the structure of the explanatory and control variables. Obviously, aggregations of this sort smooth over short run changes in Tor usage rates that could be driven by transient political clampdowns, so some information on the interaction of political repression and usage of the Tor network is lost (Tor Project, 2014b). At the same time, the yearly data structure provides a good sense of the degree to which the general structural conditions that are present in a country drive people to use online anonymity-granting technologies. Normalizing the number of Tor bridge and Tor relay users around the Internet using population of the country accommodates for the fact that there should be, everything else being equal, more users of anonymity-granting technologies in a country with more Internet users.

Tor bridges are the most effective way to circumvent access restrictions imposed by repressive regimes, so we should expect the relationship between political repression and Tor bridge usage rates to be the clearest indicator of how a country's political context drives use of the Tor network. *Tor relays* are publicly known and can be blocked by governments. As a result, we would expect that the relationship between repression and Tor usage to be slightly more muted using this indicator. As the number of clients accessing the Tor network via relays dwarfs the number of people using bridges (the relay maximum is 165,365.8 while the bridge maximum is 4692.979), the relationship between the overarching political context and all Tor usage should be driven largely by relays.

The explanatory variable

The theory posits that the usage rate of anonymity-granting technologies within a country is a function of the interaction between individual-level opportunity and political need. Unfortunately, data do not currently exist to directly measure these variables. However, the theory also predicts that these factors are conditioned in largely predictable ways by the level of political repression in the country within which people engage in online activity. Measures of political repression can, therefore, substitute for separate measures of political need and opportunity.

Political repression is an aggregate index of two of Freedom House's rights-based measures. The first measure is Freedom House's political rights measure. The second is Freedom House's civil liberties measure. Both measures range from 1 to 7, with 1 being liberal regimes such as Canada and 7 being highly repressive regimes such as China (Freedom House, 2014). The political repression variable is the summed total of the two indicators, ranging from 2 (Canada) to 14 (Uzbekistan).

Each component measure captures slightly different elements of political repression. Political rights involve rights such as free and fair elections, while civil liberties captures

a population's right to free assembly and expression. Separately, each measure captures a portion of the political environment within which people decide to use Tor. Together, they provide a strong indicator of the repressiveness of a particular context. Moreover, since the pairwise correlation between the two measures is .93, they cannot be included separately in the same regression and need to be combined, although aggregating an index in this way can be perilous.

Control variables

To correctly model the relationship between political repression and use of the Tor network, the models need to include a series of control variables that also likely affect the use of online anonymity-granting technologies. A country's Internet penetration rates, the robustness of its intellectual property (IP) rights regime, its wealth, its secondary education levels, and its openness to foreign influences all plausibly have an effect on how often people use Tor. Controlling for each provides a more accurate estimation of the effect of political repression on the use of online anonymity-granting technologies.

Controlling for Internet penetration rates is important because it is not possible for an individual to use online anonymity-granting technologies if they cannot first access the Internet. Internet penetration rates, therefore, are a necessary, but not necessarily sufficient, cause of Tor network usage. The operational measure of this variable is taken from the World Bank Indicators and is a count of the number of Internet users per 100 people. Since the ability to use the Internet is a prerequisite for the use of anonymity-granting technologies online, the expectation is that higher Internet penetration rates should be correlated with higher levels of Tor usage.

Another important control is the intellectual property (IP) rights regime in a country. Anonymity-granting technologies are used to mask an individual's identity online. What an individual decides to do with that new found anonymity is a more open question. One thing that the Tor network allow individuals to do is to host websites without having the site linked to the poster and to download illegal content or purchase illegal goods and services online. One illicit activity that is greatly facilitated by anonymity-granting technologies, particularly commercial Virtual Private Networks (VPNs) but plausibly Tor as well, is the downloading of illegal movies, songs, and other digital content. IP rights regimes criminalize the theft of IP like movies and songs. As a result, rigorous IP regimes should both discourage the theft of IP through deterrence (Husted, 2000) and encourage the use of anonymity-granting technologies that reduce the chances that a person will get caught for breaking the law. These data are from the World Economic Forum's (2015) *Global Competitiveness Index* and ranges from 1 (worse) to 7 (best). The basic expectation here is that as the robustness of a country's IP regime goes up, the rate at which individuals use anonymity-granting technologies should also rise since individuals are trying to avoid being caught breaking the law online.

Rich countries are also more likely to see the use of anonymity-granting technologies. In order to operate effectively, programs like Tor require a robust network capacity, which is far more likely in countries that are economically developed. Wealth might also be a good measure for the level of economic development in a country, which could also drive use of Tor as a more open economic environment could provide people with the

Table 1. Descriptive statistics.

	Observations	Minimum	Mean	Maximum	Standard deviation
Tor bridge usage per country per year per 100,000 Internet users	461	0	93.28108	4692.979	287.0104
Tor relay usage per country per year per 100,000 Internet users	466	23.97272	15,609.53	165,365.8	22,514.58
All Tor usage per country per year per 100,000 Internet users	462	74	2,205,177	9.04e+07	7,294,889
Political repression	471	2	7.046709	14	3.810672
Internet penetration	470	0.9	38.05739	95.0534	28.56824
IP regime	403	1.574762	3.691783	6.278549	1.109872
Log GDP	460	6.349718	9.1399	11.80647	1.236329
Education	316	14.68277	81.21254	165.5813	28.95737
Openness	442	23.71042	92.99999	376.1456	48.64676

opportunity to find and employ anonymity-granting technologies. The log of a country's gross domestic product (GDP) per capita is a good measure of its level of wealth. Data on GDP per capita is taken from the World Bank Indicators. The general idea would be that as wealth increases, the use of Tor should go up as well, as the network in the country is more robust and economic competition is greater.

The Tor network's website provides step-by-step instructions on how to download and use the Tor browser (Tor Project, 2014a). Nevertheless, the use of relatively sophisticated programs like Tor requires a certain level of technological competence. As such, it is reasonable to expect that only people with a certain capacity are likely to use the Tor network. Secondary education rates within a country proxy well for the basic numeracy and literacy that people would likely need in order to operate Tor. Secondary education might enable people to use Tor either because of the things learned in school or via a sort of diffusion effect, where people become familiar with Tor through social networks that are centered upon the school. Data on secondary education are taken from the World Bank Indicators. The general expectation is that as a country's secondary education level rises, so should its usage of the Tor network.

Some regimes are more open to outside ideas, technology, and influence than others. Numerous studies have found that the openness of a country to international influences, whether measured as foreign direct investment (Apodaca, 2001) or imports and exports as a percentage of GDP (Greenhill, 2010; Hafner-Burton and Tsutsui, 2007), is negatively related with a regime's abuse of its citizenry. One mechanism to account for this trend is the idea that elites are exposed to norms that socialize them into better behavior (Greenhill, 2010). However, regime openness might also result in the diffusion of both knowledge of programs like Tor and the technical knowhow to use the network among ordinary users. Clearly, the Internet's ability to transcend geographical borders with relative ease limits the extent to which accessing and using Tor is contingent upon a country's physical openness to the external world, but people still need to know of a technology

before they can use it. The sum of a country's imports and exports as a percentage of its GDP is a good measure of a regime's openness to the external world, as it captures the volume of goods and services flowing in and out of the country. The general expectation here is that regime openness should encourage the diffusion of knowledge of Tor, leading to more usage of the technology, everything else being equal. Table 1 outlines the descriptive statistics for all the variables.

Empirical analysis

This section presents the results of the econometric analysis into the question of whether political repression drives usage of online anonymity-granting technologies. The results show that, controlling for other relevant factors, political repression does drive usage of the Tor network. As predicted, the relationship is U-shaped, confirming one empirical expectation of the opportunity and political need framework suggested above. Political repression and Tor bridge use have the strongest and most consistent association. The relationship between a regime's political context and both the use of Tor relays and the summed measure for all Tor use is still statistically significant and U-shaped in nearly every model specification.

Given the panel structure of the data, a series of linear regression models with random effects are the best means of analyzing the data. A series of Hausman tests indicate that the random effects estimator is a good fit for the data. While random effects help to control for the effects of country-specific factors and hence provides a better estimation of the size of the effect, it is also useful to cluster the standard errors on country code to correct for heterogeneity across countries and the fact that some of the factors driving Tor usage are idiosyncratic to each particular nation state (Lambert, 2003).

Random effects models make a number of assumptions about the structure of the data. Many of these are hard to find in reality. As a result, clustering standard errors, even in random effects models, is useful because, as Kurt Schmidheiny (2015) puts it, "in practice, we can rarely be sure about equicorrelated errors [an assumption of the model] and better always use cluster-robust standard errors for the RE [random effects] estimator" (p. 7).

Table 2 presents the results of the regression analysis. The overall story to come out of the findings is that political repression does indeed drive people to use the Tor network. The precise relationship between political repression and Tor usage is consistently U-shaped, with repression driving usage of Tor the most in both the highly liberal and highly repressive contexts and the least in partly free countries.

IP: intellectual property; GDP: gross domestic product.

Political repression is most strongly related to Tor bridge usage rates, which fits with the idea that they are a purposefully designed censorship circumvention technology and should be most widely used in regimes that try to block access to the Tor network via normal relays. Nevertheless, political repression remains a significant driver of Tor usage when it comes to both relays and the sum of relays and bridge users, suggesting even more strongly that the political context within which people find themselves matters a lot for the use of anonymity-granting technologies.

Many of the control variables are highly collinear, so including all the variables in a single regression is problematic. Logged GDP per capita, Internet penetration rates, and

Table 2. (Continued)

	Model 10 (bridges)	Model 11 (relays)	Model 12 (All Tor)	Model 13 (bridges)	Model 14 (relays)	Model 15 (All Tor)	Model 16 (bridges)	Model 17 (relays)	Model 18 (All Tor)
Secondary education				0.80 (0.55)	264.35*** (65.53)	266.64*** (65.92)			
Constant	-146.82* (82.96)	-3034.57 (9011.25)	-2775.25 (9291.85)	204.82* (104.12)		17,071.38 (11,099.74)	236.05*** (87.76)	32,851.59*** (7447.89)	32,961.24*** (7533.2)
R ²	.15	.15	.14	.6	.14	.14	.07	.15	.15
Observations	427	431	427	261	261	261	400	400	400
	Model 19 (bridges)	Model 20 (relays)	Model 21 (All Tor)	Model 22 (bridges)	Model 23 (relays)	Model 24 (All Tor)			
Political repression	-44.85*** (14.61)	-4856.15*** (1724.65)	-4879.38*** (1735.64)	-38.87** (18.59)	-4355.70* (2482.94)	-4388.47* (2495.32)			
Political repression ²	2.94*** (1.01)	274.82** (123.22)	276.64** (124.00)	2.39* (1.34)	275.08 (179.76)	277.59 (180.70)			
Openness	0.08 (0.30)	36.46 (23.96)	36.46 (24.21)	0.01 (0.48)	30.48 (31.39)	30.43 (31.74)			
IP regime	-8.67 (9.89)	-4141.43*** (1391.12)	-4159.95*** (1397.83)	-14.81 (15.03)	-4585.42*** (1760.57)	-4609.74*** (1771.97)			
Log GDP per capita	63.58** (27.58)	464.80 (2892.13)	466.44 (2909.33)	107.13** (46.57)	-8587.53 (5476.76)	-8568.18 (5486.41)			
Internet	-1.68* (0.97)	251.58** (129.47)	253.83** (129.97)	-2.74* (1.49)	325.30 (202.23)	325.26 (202.82)			
Secondary education				-0.48 (0.66)	349.42*** (112.09)	351.38*** (112.25)			
Constant	-285.12** (161.81)	30,037.08 (20,352.77)	30,126.51 (20,436.33)	-575.28 (303.93)	79,216.6 (34,614.92)	79,125.72 (34,663.54)			
R ²	.10	.15	.15	.10	.16	.16			
Observations	379	379	379	259	259	259			

IP: intellectual property; GDP: gross domestic product.
 Standard errors in parentheses.
 *** = 99%; ** = 95%; * = 90%.

secondary education rates are all highly collinear with each other (above a pairwise correlation of .6 or greater). IP regime strength is collinear with logged GDP per capita and Internet penetration rates but not secondary education levels. The collinearity between these variables means that including all of the controls in the same regression can bias the coefficients on the respective variables, while not affecting the political repression term (PRT), which is not collinear with any other terms.

To test the robustness of the core finding, models 1–18 estimate the effect of political repression on Tor usage, including only unproblematic controls and a single collinear variable. With the exceptions of models 8 and 9 (for Tor relay and all Tor usage), the political repression variables emerge as significant and the relationship between a country's political context and use of the Tor network is consistently U-shaped.

Models 19–24 include different combinations of all the control variables to further see if the central result is robust even to the presence of a larger set of factors. Political repression is again consistently associated with Tor bridge usage rates in every model specification, while Tor relay usage and usage of the whole network is consistently driven by the overarching political context in all but two model specifications.

In all cases, the relationship is U-shaped as expected by the opportunity and political need framework. The negative sign on the coefficient for the PRT and the positive sign on the coefficient for the political repression squared term (PRT²) across all the models demonstrate that the relationship between political repression and Tor usage always forms a U-shaped pattern. As an indicative depiction of the relationship, Figure 2 shows the U-shaped pattern for the political repression variable found in Model 1.

The point at which political repression contributes the least to Tor usage rates varies across the different structures of the dependent variable. For Tor bridge usage, the minimum is around 8 on the political repression scale, which is equivalent to a country like Honduras. Up to this point, worsening political repression tends to reduce the extent to which the political context drives usage rates. Beyond this level of repression, the worsening political context actually starts to drive up Tor bridge usage rates. For Tor relay usage and total Tor network usage, the minimum is around 10 on the political repression scale, which is equivalent to a country like Venezuela.

Moving along the political repression scale has substantively large effects on Tor usage rates. Table 3 specifies the effect of movement along the political repression scale for models 1–3. These three models capture each structure of the dependent variable—bridges, relays, and all Tor usage—and give a good sense of the potential effects that a changing political context has upon people's decision to use the Tor network.

In each case, the models allow us to estimate the contribution of the overarching political context to a country's Tor usage rate. In the case of Tor bridges, moving from Canada (political repression equals 2) to Guatemala (political repression equals 7) results in a decrease of 166.14 Tor bridge users per 100,000 Internet users per year. In the case of Tor relays, moving from a country like the United States (political repression equals 2) to a country like Togo (political repression equals 10) tends to result in a loss of about 23,744.96 Tor relay users per 100,000 Internet users per year. A similar sized reduction occurs with usage of the entire Tor network.

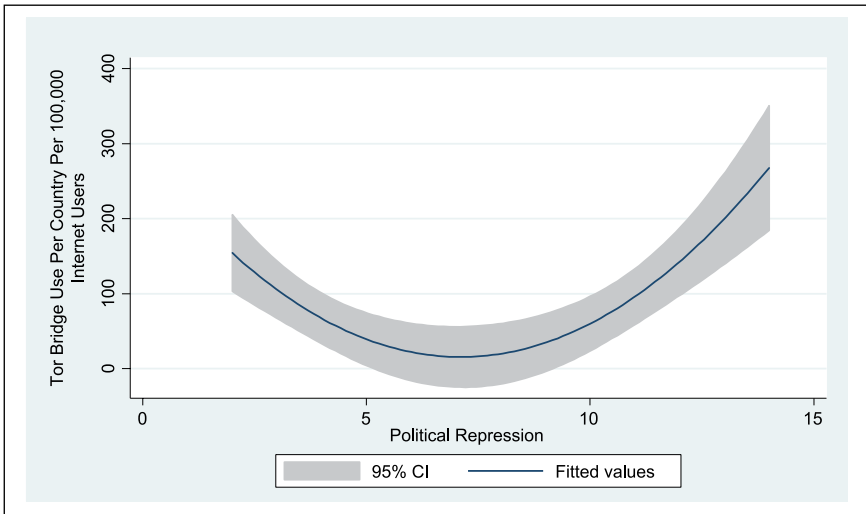


Figure 2. Political repression and Tor bridge usage.

Table 3. Changing political repression and use of the Tor network.

Change in Political Repression Scale	Tor Bridge Usage Rates per 100,000 Internet Users per Year	Tor Relay Usage Rates per 100,000 Internet Users per Year	All Tor Network Usage Rates per 100,000 Internet Users per Year
2 to 3	-53.99	-5588.29	-5562.03
3 to 4	-43.47	-4839.67	-4816.37
4 to 5	-32.95	-4091.05	-4070.71
5 to 6	-22.43	-3342.43	-3325.05
6 to 7	-11.91	-2593.81	-2579.39
7 to 8	-1.39	-1845.19	-1833.73
8 to 9	9.13	-1096.57	-1088.07
9 to 10	19.65	-347.95	-342.41
10 to 11	30.17	400.67	403.25
11 to 12	40.69	1149.29	1148.91
12 to 13	51.21	1897.91	1894.57
13 to 14	61.73	2646.53	2640.23

Beyond these minimum levels, worsening political repression starts to increase people’s use of the Tor network. In the case of Tor bridges, moving from a country like Burkina Faso (political repression equals 8) to a country like Uzbekistan (political repression equals 14) results in an increase of around 212.58 Tor bridge users per 100,000 Internet users per year. Likewise, moving from a country like Venezuela (political repression equals 10) to a country like Uzbekistan results in an increase of 6094.44Tor relay users per 100,000 Internet users per year.

Both movements can also be cast in absolute terms. Imagine as a most extreme example that China, with its roughly 620 million Internet users in 2013, moved from 8 on the political repression scale to 14 (China is actually 13). Such a move would result in a total increase of 1,317,996 Tor bridge users per year within China. A similar movement in repression would also result in an increase of 37,785,528 Tor relay users per year. Both are substantively large effects.

Interestingly, the small R^2 in most of the models suggests that individual Tor usage patterns remain highly idiosyncratic. While political context does matter, the use of anonymity-granting technologies is probably equifinal. In highly liberal regimes, a mixture of factors ranging from a normative desire to protect one's privacy through to trying to cover one's tracks while conducting socially unacceptable behaviors likely drive individuals to use Tor. In repressive countries, there is also likely a range of motivators driving use of Tor, but it is plausible that the underlying mechanism driving use shifts from the cost free opportunity to do so to the real need to avoid the eyes of the state and to use technological means to access restricted content.

Political repression is not the only significant variable in the model. Wealth emerges as a significant predictor in many specifications. This finding suggests that as a country becomes richer, it is more likely that its population will use the Tor network. This could be because the higher network capacity within wealthy countries makes Tor more efficient, faster, and easier to use. It could also be that the increased levels of economic competition that generate greater levels of wealth also encourage Tor usage. At the same time, the significance of log GDP per capita could, however, be a residual effect of wealthy nations having higher levels of Internet penetration, as the two variables are correlated at the .88 level.

Internet penetration rates are also significant predictors of use of the Tor network in a lot of the models. However, the relationship tends to emerge most strongly when it comes to accessing Tor relays rather than Tor bridges. Internet access rates are likely a good predictor of a country's network capacity, so it is reasonable to interpret the results as suggesting that, in most cases, better, more widespread Internet access increase use of the Tor network. Rather than being a real association, the negative and statistically significant relationship between Internet penetration and Tor bridge usage in models 19 and 22 is probably due to the inclusion of other controls, such as log GDP and secondary education levels that are collinear with Internet penetration rates. The lack of significance of the Internet penetration variable when focusing on Tor bridge usage in two of the models suggests that using this form of the technology is done under such a niche (likely politically driven) set of circumstances that overall network access rates do not really matter.

The effect of a country's IP regime is mixed. It is a significant predictor of the use of Tor relays but has little effect on Tor bridge use. Interestingly, the results suggest that a strong IP regime actually discourages the use of Tor, which is contrary to expectations. This finding could be because of the slow download speed and limited plug-in capabilities on the Tor network. Commercial VPNs, which are usually faster, might be a better tool to violate intellectual property laws in a country.

Conclusion

Does the level of political repression in a country incite the use of online anonymity-granting technologies? Data on Tor usage rates from 2011 to 2013 provide the start of an

answer to this question. Even after controlling for Internet penetration rates, IP regime strength, wealth, secondary education levels, and the openness of the regime, political repression emerges as a significant predictor of Tor network usage. The precise relationship to emerge is consistently U-shaped, with political repression driving Tor network usage most in both highly liberal and highly repressive regimes.

The interaction of the opportunity to use Tor and the need for people to use anonymity-granting technologies to exercise their fundamental political rights plausibly accounts for this U-shaped pattern. Political need increases as political repression worsens because people need to take additional steps to protect their identities online or risk severe repercussions. The opportunity to use anonymity-granting technologies, in contrast, is highest in liberal democratic states and lowest in countries with high levels of political repression. The implication of this framework is that the underlying bundle of motivations driving use of anonymity-granting technologies varies between highly liberal and highly repressive regimes.

The results suggest that the technology of Tor is useful for political dissidents and those trying to exercise their basic political rights. They also suggest that the underlying rationale for use likely varies between countries. In relative terms, the results suggest that the Tor network is probably more prone to abuse in liberal countries where opportunity is the underlying driver of use than in repressive regimes where people might only turn to the network because they need to do so. The ancillary expectation here is that the social costs and benefits to the Tor network are not evenly distributed globally (Jardine, 2015). Liberal countries plausibly have to deal with relatively more of the negative implications (i.e. crime and child abuse imagery) of the technology of Tor than repressive societies. In contrast, the social benefits of Tor likely cluster disproportionately in repressive regimes.

Acknowledgements

This paper has been helped immensely by the numerous eyes that have fallen upon it. In no particular order, the paper has benefited tremendously from the insights of Fen Hampson, Gordon Smith, Dane Rowlands, Simon Palamar, Laura DeNardis and Leanna Ireland. I am also grateful to Andy Greenberg for suggesting the current structure of the dependent variable. Despite all these eyes, errors might remain. Any problems that remain are my fault and my fault alone

Funding

The author(s) received no financial support for the research, authorship, and/or publication of this article.

References

- Apodaca C (2001) Global economic patterns and personal integrity rights after the cold war. *International Studies Quarterly* 45(4): 587–602.
- BBC World News (2014) Russia offers \$110,000 to crack Tor anonymous network. *BBC World News*. Available at: <http://www.bbc.com/news/technology-28526021> (accessed 7 November 2014).
- Fein H (1995) More murder in the middle: life-integrity violations and democracy in the world, 1987. *Human Rights Quarterly* 17(1): 170–191.
- Freedom House (2014) Available at: <https://freedomhouse.org/>

- Goodin D (2014) Active attack on Tor network tried to decloak users for five months. *ArcTechnica*. Available at: <http://arstechnica.com/security/2014/07/active-attack-on-tor-network-tried-to-decloak-users-for-five-months/> (accessed 29 September 2014).
- Greenhill B (2010) The company you keep: international socialization and the diffusion of human rights norms. *International Studies Quarterly* 54(1): 127–145.
- Hafner-Burton EM and Tsutsui K (2007) Justice lost! the failure of international human rights law to matter where needed most. *Journal of Peace Research* 44(4): 407–425.
- Hampson F and Jardine E (2016) *Look Who's Watching: Why the World is Losing Faith in the Internet*. Waterloo: Centre for International Governance Innovation.
- Henderson CW (1991) Conditions affecting the use of political repression. *Journal of Conflict Resolution* 35(1): 120–142.
- Husted BW (2000) The impact of national culture on software piracy. *Journal of Business Ethics* 26(3): 197–211.
- Jardine E (2015) *The Dark Web Dilemma: Tor, Anonymity and Online Policing*. Global Commission on Internet Governance paper series no. 21. Available at: http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2667711
- Lambert T (2003) When you need to correct for clustering. *Deltoid*. Available at: <http://scienceblogs.com/deltoid/2003/09/10/cluster/> (accessed 7 November 2014).
- MIT Technology Review (2012) How China blocks the Tor anonymity network. *MIT Technology Review*. Available at: <http://www.technologyreview.com/view/427413/how-china-blocks-the-tor-anonymity-network/> (accessed 8 November 2012).
- Muller E (1985) Income inequality, regime repressiveness, and political violence. *American Sociological Review* 50(1): 47–61.
- Owen G and Savage N (2015) *The Tor Darknet*. Global Commission on Internet Governance paper series no. 20. Available at: <https://ourinternet.org/publication/the-tor-dark-net/>
- Princeton's Office of Information Technology (2014) Encryption and international travel. Available at: <http://www.princeton.edu/itsecurity/encryption/encryption-and-international/>
- Regan P and Henderson E (2002) Democracy, threats and political repression in developing countries: are democracies internally less violent? *Third World Quarterly* 23(1): 119–136.
- Schmidheiny K (2015) Panel data: fixed and random effects. Available at: <http://kurt.schmidheiny.name/teaching/panel.pdf>
- Tor Challenge (2014) What is Tor? Available at: <https://www.eff.org/torchallenge/what-is-tor.html>
- Tor Project (2014a) Tor. Available at: <https://www.torproject.org/index.html.en>
- Tor Project (2014b) Tor metrics: users. Available at: <https://metrics.torproject.org/users.html?graph=userstats-relay-country&start=2014-01-01&end=2014-07-24&country=all&events=off#userstats-relay-country>
- Tor Project (2015) Tor: overview. Available at: <https://www.torproject.org/about/overview>
- World Economic Forum (2015) *Global Competitiveness Index*. Geneva: World Economic Forum.

Author biography

Eric Jardine is a research fellow in the Global Security & Politics Program at Centre for International Governance Innovation (CIGI). As a part of this position, he contributes to CIGI's work on Internet governance as a member of the Secretariat for the Global Commission on Internet Governance. He has authored numerous scholarly articles on trends in cybercrime, the uses and abuses of the Dark Web, and contention in Internet governance. He is also the author, with Fen Hampson, of the book, *Look Who's Watching: Why the World is Losing Faith in the Internet*. He regularly appears as a commentator on cybersecurity issues in television, radio, and print media. He holds a PhD in International Affairs from the Norman Paterson School of International Affairs, Carleton University, Canada.