



The potential harms of the Tor anonymity network cluster disproportionately in free countries

Eric Jardine^{a,1,2} , Andrew M. Lindner^{b,1} , and Gareth Owenson^{c,1} 

^aDepartment of Political Science, Virginia Tech, Blacksburg, VA 24061; ^bDepartment of Sociology, Skidmore College, Saratoga Springs, NY 12866; and ^cCyber Espion Ltd, Portsmouth PO2 0TP, United Kingdom

Edited by Douglas S. Massey, Princeton University, Princeton, NJ, and approved October 23, 2020 (received for review June 10, 2020)

The Tor anonymity network allows users to protect their privacy and circumvent censorship restrictions but also shields those distributing child abuse content, selling or buying illicit drugs, or sharing malware online. Using data collected from Tor entry nodes, we provide an estimation of the proportion of Tor network users that likely employ the network in putatively good or bad ways. Overall, on an average country/day, ~6.7% of Tor network users connect to Onion/Hidden Services that are disproportionately used for illicit purposes. We also show that the likely balance of beneficial and malicious use of Tor is unevenly spread globally and systematically varies based upon a country's political conditions. In particular, using Freedom House's coding and terminological classifications, the proportion of often illicit Onion/Hidden Services use is more prevalent (~7.8%) in "free" countries than in either "partially free" (~6.7%) or "not free" regimes (~4.8%).

Dark Web | political freedom | political rights | cryptomarkets | child abuse

Debate rages about the social utility of an anonymous portion of the global Internet accessible via the Tor network and colloquially known as the Dark Web (1).^{*} Although other similar tools exist, The Onion Router (Tor) is currently the largest anonymity network. Tor users can act as publishers of content by using the network to anonymously administer Onion/Hidden Services for the use of others. They can also use the Tor browser to anonymously read either these Onion/Hidden Services (i.e., sites with rendezvous points located internal to the Tor network) or to access Clear Web sites (1–5). With these diverse supply-side and demand-side functions (6), many point to the socially harmful uses of Tor as an anonymous platform for child abuse imagery sites (7, 8), illicit drug markets (9–13), gun sales (14, 15), and potential extremist content that has shifted to the Dark Web after extensive Clear Web content moderation efforts (16). Others emphasize its socially beneficial potential as a privacy-enhancing tool and censorship circumvention technology (17–22).

Both sides of the debate illustrate genuine uses of the technology. Like any tool that is inherently dual use, questions abound about whether its benefits are worth the costs. Such questions have both net (i.e., do costs or benefits predominate) and distributional (i.e., how are the harms/benefits spread out) dimensions. Overall, a technology like the Tor anonymity network might do more harm than good. It may also be more harmful in some locations than others. Ultimately, these are empirical questions.

In the case of the Tor anonymity network, our data provide clear, if probabilistic, answers to these questions. Our data show that in net terms, only a small fraction of Tor users employ the anonymity system for likely malicious purposes. On an average day during our sample period, for example, about 6.7% of Tor network clients globally use the network to connect to "Onion/Hidden Services" that are predominantly used for illicit and illegal activities, such as buying drugs, distributing malware, or consuming and sharing child abuse imagery content. To be sure, there some socially beneficial content on Onion/Hidden Services and plenty of troubling content on the Clear Web.

However, substantial evidence has shown that the preponderance of Onion/Hidden Services traffic connects to illicit sites (7). With this important caveat in mind, our data also show that the distribution of potentially harmful and beneficial uses is uneven, clustering predominantly in politically free regimes. In particular, the average rate of likely malicious use of Tor in our data for countries coded by Freedom House as "not free" is just 4.8%. In countries coded as "free," the percentage of users visiting Onion/Hidden Services as a proportion of total daily Tor use is nearly twice as much or ~7.8%. These findings are robust to a different measure of political freedom and the inclusion of a variety of statistical controls. They also give rise to a number of important public policy challenges.

Data, Material, and Methods

Our data were collected by running 1 percent of entry (Guard) nodes in the Tor network from December 31, 2018, to August 18, 2019, with a short interruption to data collection from May 4, 2019, to May 13, 2019. Tor clients (users) randomly choose an entry node from the set of available nodes in the network (weighted by available bandwidth). By running 1 percent of Guard nodes, we observe a random sample of all Tor relay users, although our data do not include those who employ Tor bridges to access the network. By analyzing unique signatures in the traffic (e.g., directory lookups), we can distinguish whether clients are using Tor to visit either the Clear Web (e.g.,

Significance

Measuring the proportion of Tor anonymity network users who employ the system for malicious purposes is important as this technology can facilitate child abuse, the sale of illicit drugs, and the distribution of malware. We show that only a small fraction of users globally (~6.7%) likely use Tor for malicious purposes on an average day. However, this proportion clusters unevenly across countries, with more potentially malicious Tor users in "free" countries (~7.8%) than in "not free" regimes (~4.8%). These results suggest that the countries which host most of the infrastructure of the network and house the Tor Project plausibly experience a disproportional amount of harm from the Tor anonymity network.

Author contributions: E.J., A.M.L., and G.O. designed research; E.J., A.M.L., and G.O. performed research; E.J., A.M.L., and G.O. analyzed data; E.J., A.M.L., and G.O. wrote the paper; and G.O. provided the data collection.

The authors declare no competing interest.

This article is a PNAS Direct Submission.

This open access article is distributed under [Creative Commons Attribution-NonCommercial-NoDerivatives License 4.0 \(CC BY-NC-ND\)](https://creativecommons.org/licenses/by-nc-nd/4.0/).

¹E.J., A.M.L., and G.O. contributed equally to this work.

²To whom correspondence may be addressed. Email: ejardine@vt.edu.

First published November 30, 2020.

^{*}A note on terminology. The Dark Web is an often hotly contested term. In the interest of technical accuracy, we use the term "Tor anonymity network." The Tor anonymity network both hosts a Dark Web (e.g. Onion/Hidden services running standard web technologies) and includes a special routing browser that allows users to engage with content on both the Dark and the Clear Webs while masking their IP addresses. We use the term Clear Web to describe sites like CNN.com, which are accessible on the regular Internet via the Tor browser. For a more detailed discussion of the conceptual boundaries of the Dark Web, see ref. 1.

CNN.com) or a Tor Onion/Hidden Service (e.g., xyz.onion). This process does not reveal anything about the precise content a user is querying. Additionally, we geolocate the user's incoming IP address to a country of origin and aggregate these data into 1) a count of all Tor network users per country per day and 2) a count of Onion/Hidden Services users per country per day.

We merge these aggregate Tor network data with measures of country-level political freedom, taken from both Freedom House's annual *Freedom in the World* reports (23, 24) and the "PolityV Political Regime Characteristics and Transitions" dataset (25); the most recent available country-level indicators for wealth, Internet penetration, and population size from the World Bank (26); and an estimate of per capita Darknet cryptomarket activity at a country level in the years immediately preceding our study period (27). Table 1 provides a descriptive summary of the data.

Initial ethics approval for the data collection infrastructure was granted by the University of Portsmouth Ethics Committee (ETHICS-GO1). An additional "not human subjects research" determination was made by Virginia Tech's Institutional Review Board (20-576), as all data for this project were analyzed in aggregate at a country level and user IP addresses were not included in any dataset used for the analysis. IP addresses were translated at source using Maxmind's Geolocation Service.

To approximate the potential harms and benefits of Tor, we measure the percentage of a country's total Tor clients that go to Tor Onion/Hidden Services on the average day in the observed time range (hereafter %HS).

Interpreting this variable as a country-level measure of the potential illicit use of Tor requires an assumption about the nature of Tor traffic. For our purposes, we assume that a larger share of Tor Clear Web users likely employ the Tor anonymity network for rights-based uses (i.e., privacy protection or censorship circumvention) than Onion/Hidden Services users. Conversely, we assume that a larger share of Tor clients heading to Onion/Hidden Services are more likely to be engaged in predominately illicit activities, such as visiting child abuse sites, drug markets, or hacker forums than is common among Tor Clear Web users.

While this assumed pattern is not universally true (e.g., Facebook and the New York Times have commonly used Onion/Hidden Services and many Clear Web sites host extremist content or other illicit material), there are four reasons why this assumption might be generally, that is to say probabilistically, correct.

First, for users who employ Tor to access Clear Web content, the destination sites have known/knowable administrators and web services providers. These features of Clear Web content allow for greater transparency and content moderation and can minimize (although not eliminate) the possibility that users are engaging with content that is widely considered malicious or illegal (16, 28, 29).

In contrast, while the majority of hosted Onion/Hidden Services are often not per se illegal (28, 30), user demand for this content tends to suggest that illicit use of Onion/Hidden Services predominates (6). First, one 2015 study investigating site visits to Onion/Hidden Services found that roughly 82% of requests over a 6-month observation period went to child abuse imagery sites, although this period also corresponded with a major FBI investigation into the Playpen child abuse imagery site (7). Second, a number of empirical investigations have documented the rapid growth of drug cryptomarkets, even while the average global number of Tor users remains relatively constant over time (2, 9, 31). Since Snowden's disclosures of National Security Agency surveillance in 2013, for example, Tor network use has stabilized to

around 2 to 2.5 million users per day (2, 32). During this same period, cryptomarket vendor counts alone have increased from reportedly 3,877 unique vendor accounts on Silk Road at the time of its closure in October 2013 to ~40,000 vendors on AlphaBay in 2017, just four short years later (33, 34). Third, cryptomarkets originally had a significant political dimension to them, often coupling the sale of drugs with discussion of libertarianism and politics (1, 35, 36). The Silk Road site administrator, Ross Ulbricht (also known as the Dread Pirate Roberts), even hosted a political book club via the site. These political dimensions, however, have declined significantly over time, leaving largely just drug exchange (37).

In combination, past findings such as these suggest that users of Onion/Hidden Services content tends to cluster heavily toward malicious/illicit/illegal uses, although many benign or even beneficial Onion/Hidden Services sites do exist (1, 28, 36). In contrast, use of Tor to access Clear Web content implies that users are browsing sites where the operators of the accessed content are known/knowable, and so comparatively less likely to be hosting content that is widely agreed to be illegal or malicious, although illicit or illegal activity certainly still does occur on the Clear Web (16, 28). In sum, the current study relies upon a well-documented if probabilistic pattern: A larger share of Onion/Hidden Services users are likely to be engaged in illicit activity than Tor Clear Web users.

Net Dimensions. Our data suggest that over time and across all countries, users of the Tor anonymity network predominately employ the system to venture to Clear Web content. Aggregating daily observations into a country's daily average and then tabulating these values suggests that, overall, just 6.7% of Tor users during the study period in 2019 visited Onion/Hidden Services sites. Framed differently, only about 1 in 20 Tor users on an average day may be employing the system for potentially illicit purposes. The remainder of the users employed the Tor network to view Clear Web sites, suggesting that approximately upwards of 93% of Tor users globally go to websites on the Clear Web that are not administered anonymously and so comparatively less likely to be hosting malicious or illegal content. In net terms, these data suggest that the bulk of Tor users are, on an average daily basis, doing comparatively licit things with the Tor anonymity network and are not viewing Onion/Hidden Services content.

Extrapolating our average %HS estimate onto the total daily number of Tor network clients at a global level reveals the scale of the potentially benign and malicious use of Tor. The Tor Project publishes aggregate daily client counts per country, collected by observing requests for Tor Directory sites and mirrors (2). Like our own data, these numbers do not reveal the number of unique Tor users, only the number of distinct connection requests. For the full observation period of our data excluding the 9-d period of interrupted data collection, the average daily number of Tor relay clients is 2,231,334 connections globally. Using a strict version of our simplifying assumption and combining our estimates of the overall %HS with this total count indicates that on an average day in 2019 about 149,499 Tor network clients (nonunique) were potentially using the network to engage in possibly illicit activity on Onion/Hidden Services. Inversely, of course, this also implies that on an average day in 2019 upwards of 2,081,834 Tor users visited likely benign Clear Web content.

Distributional Dimensions. Our data suggest that likely malicious/benign users of the Tor anonymity network are not evenly spread globally and, in fact, vary

Table 1. Descriptive statistics

Variables	<i>n</i>	Mean	Median	SD	Min	Max
Mean % HS	195.00	6.70	6.65	4.40	0.00	27.04
Political freedom (Freedom House)*	195.00	8.19	9.00	4.05	1.00	16.00
Polity2 (autocratic to democratic scale)	164.00	4.18	7.00	6.14	-10.00	10.00
GDP per capita (in \$10,000s)	195.00	1.48	0.54	2.41	0.03	17.28
Population (in 100,000s)	195.00	38.89	8.84	143.67	0.01	1,392.73
% Net penetration	195.00	54.78	58.77	29.08	0.10	99.65
No. of cryptomarket sales (per 100,000)†	195.00	1.18	0.00	5.51	0.00	61.81

*All political freedom measures have been coded so that higher levels indicate greater freedom.

†The cryptomarket sales volume per 100,000 people variable is a country-level measure of items listed as "ships from" on four markets that operated shortly before our data collection period (scraped 2017–2018; data available here, ref. 27). These cryptomarkets include: Dream, Traderoute, Berlusconi, and Valhalla. The country-level sales volume estimations are normalized around the population of each country, providing the resulting measure of per capita cryptomarket activity. The top 10 countries in terms of such activity are Netherlands, United Kingdom, Australia, Germany, Finland, Luxembourg, Canada, United States, Hong Kong, and Belgium. Countries without any items sold were coded as a zero.

systematically based upon prevalent political conditions within a country. These results are in line with the predictions of the need/opportunity framework, which posits that the opportunity to use Tor should be the predominant driver of use in free regimes and that the political need to use anonymity-enhancing technologies should be the primary driver of use in not free regimes (19, 21). The analysis in this section is broken into three parts: 1) cross-sectional and longitudinal summary findings by aggregate political conditions; 2) cross-sectional trends by disaggregated political conditions; and 3) tests of the robustness of the main findings once controlling for confounders such as wealth, population size, estimated country-level cryptomarket activity, and Internet penetration rates, as well as an alternative measure of political freedom.

Aggregate Political Freedom Levels and the Malicious Use of Tor. A country's political conditions systematically predict the degree to which users within that jurisdiction employ the Tor anonymity network for potentially licit or illicit purposes. Fig. 1, for example, presents the %HS variable for countries coded as "free" ($n = 86$), "partially free" ($n = 59$), or "not free" ($n = 50$) by Freedom House for which we have Tor usage data. The violin plot shows that the density of average daily %HS in "not free" regimes is significantly lower than the same proportion of Tor users in "free" countries. More precisely, a significantly greater share of Tor users in free countries go to Onion/Hidden Services (free %HS = 7.8%) than the global average of %HS = 6.7%. Nearly all of the countries with %HS of greater than 15% are categorized as "free." Users in "partially free" regimes have average daily %HS users (partially free %HS = 6.7%), comparable to the global average. Onion/Hidden Services have significantly fewer users of Tor in "not free" regimes such as China, Russia, or Algeria as a proportion of local connections (not free %HS = 4.8%).

Fig. 2 plots the average %HS variable by political conditions over the full observation period. The ordinal ranking of "free," "partially free," and "not free" remains relatively constant throughout this period, suggesting that the core findings are not driven by outlying, short-run blocks of time. Interestingly, the results also suggest that from early-to-late 2019 there was a fairly steady increase in the proportion of users visiting Onion/Hidden Service in all regime types. There was a large spike in the %HS variable in late August 2019. The reasons for the increase are unclear but do not detract from this study's claims regarding the global distribution of Onion/Hidden Services use since the ordinal ranking of the %HS by political conditions remains the same as predicted even during this period of anomalous volume.

Disaggregated Political Freedom Measures and the Potential Illicit Use of Tor. Political freedoms/restrictions take many forms, some of which might be more relevant to the potential illicit use of Tor than others. Freedom House's "free," "partially free," and "not free" categories are aggregates of a country's level of political freedom and civil liberties. Each of these two categories, in turn, breaks down into additional, more refined subcategories. Political Rights (range 0–12) is composed of: 1) the freedom of a

country's electoral processes, 2) its levels of political pluralism, and 3) the functioning of its government. Civil Liberties (range 0–16) is composed of: 1) freedom of expression and belief, 2) associational and organizational rights, 3) rule of law and 4) personal autonomy and individual rights (23, 24). In each case, higher scores correspond with higher levels of freedom and rights protection.

Fig. 3 shows the association between a country's various subcategory political conditions and its %HS. Each subcategory of political freedom exhibits a positive correlation with a higher %HS. Correlation coefficients suggest that the "function of government" and "individual rights" variables are a bit more strongly associated with higher %HS, but generally the pattern exhibited by every subcategory of political rights and civil liberties is the same. Increases in political freedom correlate with higher %HS or, essentially, a greater share of potentially illicit users of Tor.

What of Potential Confounders or Other Measures of Political Conditions? A number of other factors might also correlate with Tor anonymity network usage and act as potential confounders. Tor might, for example, be a tool of the wealthy, be functionally useful only in larger populations due to how it produces anonymity, require a certain amount of aggregate Internet penetration, or cluster in countries that are high users of cryptomarkets (19, 38). These potential confounders suggest that the correlation between country-level political conditions and %HS might be spurious. Additionally, while Freedom House provides a widely used assessment of political conditions around the world, no measure is without its potential blind spots. The results presented above might, therefore, be driven by some unknown feature of Freedom House's coding schema and may not be robust to alternative measures of country-level political conditions.

The regression models presented in Table 2 test the robustness of the core finding to both the inclusion of a number of controls (i.e., country wealth, population size, Internet penetration rates, and estimated cryptomarket activity) and alternative dependent variable specifications. Across all models, increasing political freedom, regardless of the political freedom measure used, correlates positively with a higher %HS. In the full models (models 3 and 6), only country political conditions and Internet penetration rates correlate significantly with the %HS variable. Neither wealth, population size, nor a cryptomarket sales volume activity measure correlate significantly with %HS.

The effect size of the political freedom variables is also substantively large. Moving from the lowest to highest observed levels of political freedom in the two full models (models 3 and 6), for example, increases the predicted %HS by 2.39 percentage points using the Freedom House scoring and by 3.42 percentage points using the Polity2 measure. In short, the association between country-level political conditions and %HS is highly consistent, surviving both the inclusion of other country-level variables and alternative measures of political freedom.

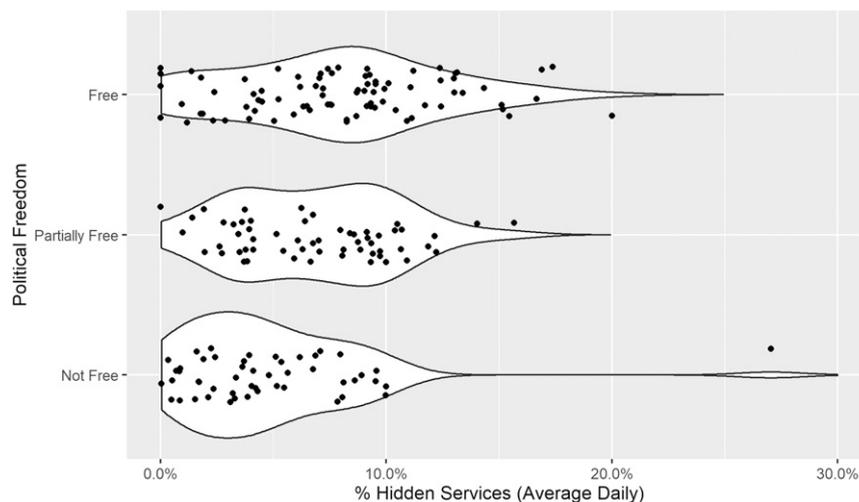


Fig. 1. More politically "free" countries have higher proportions of Hidden Services traffic than is present in either "partially free" or "not free" nations ($n = 195$ countries). Each point indicates the average daily %HS for a given country. The white regions represent the kernel density distributions for each ordinal category of political freedom ("free," "partially free," and "not free").

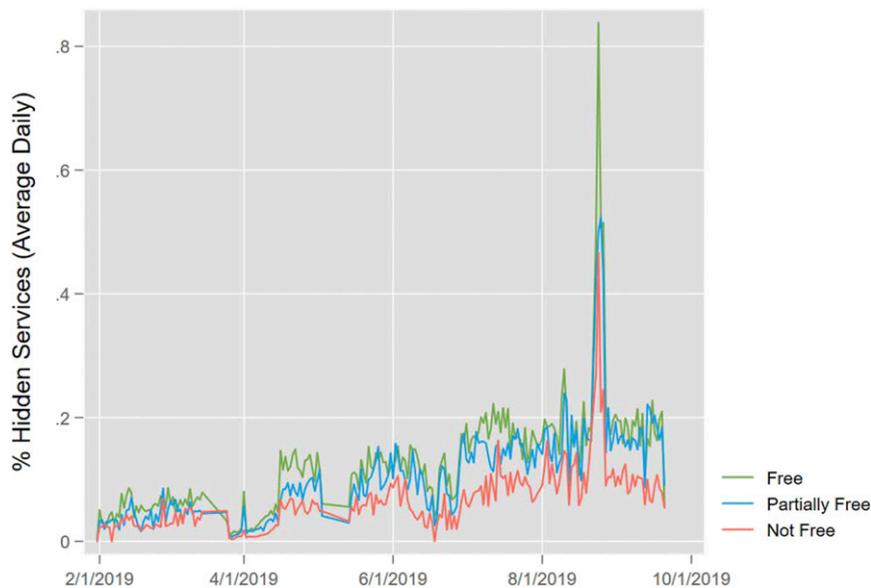


Fig. 2. Throughout the observed time period (12/31/2018–8/18/2019), “free” countries typically had a higher proportion of Hidden Services traffic than either “partially free” or “not free” nations on a daily basis. The green, cyan, and mauve lines indicate the average daily %HS for “free,” “partially free,” and “not free” countries, respectively ($n = 37,922$ country-days).

Discussion

Our results make two primary contributions to knowledge and have a number of implications.

First, while other studies have measured and parsed malicious Tor network traffic (39), our study provides a viable estimates of how users, as distinct from the traffic they generate, employ the



Fig. 3. Data show a positive association between Freedom House’s political freedom subcategories and average daily %HS. Each graph represents a different aspect of political freedom (the freedom of a country’s “electoral process,” its levels of “political pluralism and participation,” the “functioning of government,” the extent of “freedom of expression and belief,” “associational and organizational rights,” “rule of law,” and “personal autonomy/individual rights”). The same 195 countries (each represented by a point) appear in all seven graph. The blue line in each graph is a linear fitted line for the association between each subcategory of political freedom and %HS.

Table 2. Ordinary Least Squares regression on two measures of country-level political conditions

	Mean % HS					
	Model 1	Model 2	Model 3	Model 4	Model 5	Model 6
Political freedom (Freedom House)	0.325*** (0.075)	0.300*** (0.077)	0.171* (0.081)			
Polity2 (autocratic to democratic scale)				0.232*** (0.048)	0.216*** (0.049)	0.171*** (0.048)
Number of cryptomarket sales (per 100,000)		0.077 (0.056)	0.036 (0.056)		0.086 (0.050)	0.045 (0.053)
GDP per capita (in \$10,000s)			0.024 (0.156)			-0.114 (0.221)
Population (in 100,000s)			-0.003 (0.002)			-0.003 (0.002)
% Net penetration			0.044*** (0.013)			0.047*** (0.013)
Constant	4.038*** (0.681)	4.152*** (0.685)	2.921*** (0.800)	5.711*** (0.359)	5.664*** (0.358)	3.672*** (0.634)
N	195	195	195	164	164	164
R ²	0.090	0.098	0.176	0.124	0.140	0.236

*P < 0.05, **P < 0.01, ***P < 0.001.

Tor anonymity network. On this front, our results show that most users on an average day are using Tor to view and engage with Clear Web content, not Onion/Hidden Services. Based upon a strict reading of our simplifying assumption, this finding suggests that most users of Tor are likely engaged in predominately licit or even rights-based activities. Debates about the future of encryption and anonymity-granting technologies should weigh carefully such preponderance of potential use.

Second, as predicted elsewhere (19, 21), the findings contribute to our understanding of the global spread of the potential harms/benefits of Tor. Our results suggest that likely licit and illicit use of the Tor anonymity network is highly uneven and varies systematically by political context. Potentially harmful uses of Tor tend to cluster predominately in “free” regimes and users who are more likely leveraging the network for rights-based purposes tend to cluster in “not free” countries. These trends befit the dual-use nature of the Tor anonymity network. The distribution of this pattern is relatively stable over time and common across all subcategories of political rights and civil liberties. It is, finally, political context that is more strongly associated with %HS levels, not other potential confounders like wealth, cryptomarket usage, or population size—although Internet penetration rates also matter.

These results have a number of consequences for research and policy. First, the results suggest that anonymity-granting technologies such as Tor present a clear public policy challenge and include clear political context and geographical components. This policy challenge is referred to in the literature as the “Dark Web dilemma” (21). At the root of the dilemma is the so-called “harm principle” proposed in *On Liberty* by John Stuart Mill (40). In this principle, it is morally permissible to undertake any action so long as it does not cause someone else harm. The challenge of the Tor anonymity network, as intimated by its dual use nature, is that maximal policy solutions all promise to cause harm to some party. Leaving the Tor network up and free from law enforcement investigation is likely to lead to direct and indirect harms that results from the system being used by those engaged in child exploitation (7, 8), drug exchange (9–13), and the sale of firearms (14, 15), although these harms are of course highly heterogeneous in terms of their potential negative social impacts and some, such as personal drug use, might also have predominantly individual costs in some cases. Conversely, simply working to shut down Tor would cause harm to dissidents and human rights activists, particularly, our results suggest, in more repressive, less politically

free regimes where technological protections are often needed the most (19, 20).

Our results showing the uneven distribution of likely licit and illicit users of Tor across countries also suggest that there may be a looming public policy conflagration on the horizon. The Tor network, for example, runs on ~6,000–6,500 volunteer nodes. While these nodes are distributed across a number of countries, it is plausible that many of these infrastructural points cluster in politically free liberal democratic countries. Additionally, the Tor Project, which manages the code behind the network, is an incorporated not for profit in the United States and traces both its intellectual origins and a large portion of its financial resources to the US government (1, 41, 42). In other words, much of the physical and protocol infrastructure of the Tor anonymity network is clustered disproportionately in free regimes, especially the United States. Linking this trend with a strict interpretation of our current results suggests that the harms from the Tor anonymity network cluster in free countries hosting the infrastructure of Tor and that the benefits cluster in disproportionately highly repressive regimes.

Like recent debates over encryption developed by predominately Western technology companies (43, 44), it is plausible that the uneven spread of harms and benefits, when combined with a sense of who can exert jurisdictional control over the infrastructure of the Tor anonymity network, might lead to public policy debates about the future of the technology. Already, survey work suggests that a majority of the global public support shutting down the “Darknet”—a task that is easier said than done, rife with ethical, legal, and technological challenges and sensitive to respondents attitudes toward privacy, censorship, and their previous exposure to online crime (20). The uneven spread of harms/benefits from the Tor anonymity network might feed into this debate and fuel it anew.

Our results suggest two things to this potential debate. First, it is likely that there is some cost redistribution occurring as a result of the dual-use functionalities of Tor anonymity network. In other words, “free” countries are likely bearing an increased social cost of some size (via the harms from hosted child abuse content, illicit drug markets, etc.) so that those in not free regimes might have access to a robust anonymity-enhancing tool. Determining if these increased costs are an acceptable burden to pay so that others might exercise basic political rights is an important normative debate to which the present study supplies some modest empirical results. Second, in the context of any debate about the future of Tor, it is important to bear in mind that our simplifying assumption is merely probabilistic. There are

a lot of rights-oriented uses of Onion/Hidden Services so resolving the tensions surrounding Tor by simply closing these sites would not be an unambiguously effective policy approach.

Conclusions

The Tor anonymity network can be used for both licit and illicit purposes. Our results provide a clear, if probabilistic, estimation of the extent to which users of Tor engage in either form of activity. Generally, users of Tor in politically “free” countries are significantly more likely to be using the network in likely illicit ways. A host of additional questions remain, given the anonymous nature of Tor and other similar systems such as I2P and Freenet. Our results narrowly suggest, however, users of Tor in more repressive “not free” regimes tend to be far more likely to venture via the Tor network to Clear Web content and so are

comparatively less likely to be engaged in activities that would be widely deemed malicious.

Data Availability. CSV data for this project, “The potential harms of the Tor anonymity network cluster disproportionately in free countries,” have been deposited in Open Science Foundation and were last accessed on 13 November 2020. The data and code for analysis can be accessed at: https://osf.io/svrdz/?view_only=d676216acb714521b35759238f69c731.

ACKNOWLEDGMENTS. We thank the editors and reviewers of PNAS for helpful comments and guidance. Paul Avey and Leanna Ireland also provided valuable feedback on earlier drafts. Portions of this work were financially supported by a grant from the Institute for Society, Culture and Environment at Virginia Tech. We are also grateful to the Department of Political Science at Virginia Tech for facilitating open access publication of this article.

1. R. W. Gehl, *Weaving the Dark Web: Legitimacy on Freenet, Tor, and I2P* (MIT Press, Cambridge, 2018), pp. 276.
2. Project T, Users. <https://metrics.torproject.org/userstats-relay-country.html>. Accessed 13 November 2020.
3. Project T, Onion services. <https://metrics.torproject.org/hidserv-dir-onions-seen.html>. Accessed 13 November 2020.
4. R. Dingleline, N. Mathewson, P. Syverson, “Tor: The second-generation onion router” in *Proceedings of the 13th conference on USENIX Security Symposium* (USENIX Association, San Diego, CA, 2004), pp. 21, Vol. 13.
5. R. Graham, B. Pitman, Freedom in the wilderness: A study of a darknet space. *Convergence* **26**, 593–619 (2020).
6. E. Jardine, The trouble with (supply-side) counts: The potential and limitations of counting sites, vendors or products as a metric for threat trends on the dark web. *Intell. Natl. Secur.* **34**, 95–111 (2019).
7. G. Owen, N. Savage, The Tor Dark Net. *Global Commission on Internet Governance Paper Series* (20), (2015), pp. 1–20.
8. B. R. da Cunha et al., Assessing police topological efficiency in a major sting operation on the dark web. *Sci. Rep.* **10**, 73 (2020).
9. N. Christin, “Traveling the silk road: A measurement analysis of a large anonymous online marketplace” in *Proceedings of the 22nd international conference on World Wide Web* (ACM, Rio de Janeiro, Brazil, 2013), pp. 213–224.
10. K. Soska, N. Christin, “Measuring the longitudinal evolution of the online anonymous marketplace ecosystem” in *24th USENIX Security Symposium* (2015), pp. 33–48.
11. J. Martin, Lost on the Silk Road: Online drug distribution and the ‘cryptomarket’. *Criminol. Crim. Justice* **14**, 351–367 (2014).
12. J. Martin, *Drugs on the Dark Net: How Cryptomarkets are Transforming the Global Trade in Illicit Drugs* (Palgrave Macmillan, Basingstoke, 2014).
13. E. Jardine, A. M. Lindner, The Dark Web and cannabis use in the United States: Evidence from a big data research design. *Int. J. Drug Policy* **76**, 102627 (2020).
14. C. Copeland, M. Wallin, T. J. Holt, Assessing the practices and products of Darkweb Firearm vendors. *Deviant Behav.* **41**, 949–968 (2020).
15. R. Liggett, J. R. Lee, A. L. Roddy, M. A. Wallin, “The dark web as a platform for crime: An exploration of illicit drug, firearm, CSAM, and cybercrime markets” in *The Palgrave Handbook of International Cybercrime and Cyberdeviance*, T. J. Holt, A. M. Bossler, Eds. (Palgrave Macmillan, 2020), pp. 91–116.
16. E. Jardine, Online content moderation and the dark web: Policy responses to radicalizing hate speech and malicious content on the darknet. *First Monday* **24**, 12 (2019).
17. A. M. Lindner, T. Xiao, Subverting surveillance or accessing the Dark Web? Interest in the Tor anonymity network in U. S. States, 2006–2015. *Soc. Curr.*, **7**, 352–370 (2020).
18. Project T, Tor: Overview. <https://2019.www.torproject.org/about/overview>. Accessed 13 November 2020.
19. E. Jardine, Tor, what is it good for? Political repression and the use of online anonymity-granting technologies. *New Media Soc.* **20**, 435–452 (2018).
20. E. Jardine, Privacy, censorship, data breaches and Internet freedom: The drivers of support and opposition to Dark Web technologies. *New Media Soc.* **20**, 2824–2843 (2018).
21. E. Jardine, The Dark Web Dilemma: Tor, Anonymity and Online Policing. *Global Commission on Internet Governance Paper Series* (21), (2015), pp. 1–24.
22. D. A. Makin, L. Ireland, The secret life of PETs A cross-sectional analysis of interest in privacy enhancing technologies. *Policing*, 121–136 (2019).
23. F. House, *Freedom in the world comparative and historical data: Country and territory ratings and statuses, 1973–2019* (Freedom House, Washington, DC, 2019).
24. F. House, “Freedom in the world: Country and territory ratings and statuses, 1973–2019” (Freedom House, Washington, DC, 2019).
25. Peace Cfs, Polity V: Political regime characteristics and transitions, 1800–2018. <https://www.systemicpeace.org/inscrdata.html>. Accessed 13 November 2020.
26. The World Bank, World Bank Data | Indicators. <https://data.worldbank.org/indicator>. Accessed 12 November 2020.
27. N. Christin, Dream, traderoute, Berlusconi and Valhalla marketplaces, 2017–2018: Anonymized datasets. https://www.impactcybertrust.org/dataset_view?idDataset=1200. Accessed 14 October 2020.
28. D. Moore, T. Rid, Cryptopolitik and the Darknet. *Survival (Lond.)* **58**, 7–38 (2016).
29. T. Gillespie, *Custodians of the Internet Platforms, Content Moderation, and the Hidden Decisions that Shape Social Media* (Yale University Press, New Haven, 2018).
30. M. Faizan, R. A. Khan, Exploring and analyzing the dark Web: A new alchemy. *First Monday* **24**, <https://doi.org/10.5210/fm.v24i5.9473> (2019).
31. N. Christin, S. Egelman, T. Vidas, J. Grossklags, “It’s all about the benjamins: An empirical study on incentivizing users to ignore security advice” in *Financial Cryptography and Data Security*, G. Danezis, Ed. (Springer, Berlin Heidelberg, 2012), pp. 16–30.
32. F. O. Hampson, E. Jardine, *Look Who’s Watching: Surveillance, Treachery and Trust Online* (CIGI Press/McGill-Queen’s University Press, Waterloo, ed. 1, 2016), p. 340.
33. PBS News Hour, The Web’s biggest illegal drug marketplace is shut down by the FBI. <https://www.pbs.org/newshour/nation/the-webs-biggest-illegal-drug-marketplace-is-shut-down-by-the-fbi>. Accessed 13 November 2020.
34. FBI, Darknet takedown. *Authorities Shutter Online Criminal Market AlphaBay*, <https://www.fbi.gov/news/stories/alphabay-takedown>. Accessed 13 November 2020.
35. A. Maddox, M. J. Barratt, M. Allen, S. Lenton, Constructive activism in the dark web: Cryptomarkets and illicit drugs in the digital ‘demimonde’. *Inf. Commun. Soc.* **19**, 111–126 (2016).
36. R. W. Gehl, Power/freedom on the dark web: A digital ethnography of the dark web social network. *New Media Soc.* **18**, 1219–1235 (2016).
37. R. Munksgaard, J. Demant, Mixing politics and crime—The prevalence and decline of political discourse on the cryptomarket. *Int. J. Drug Policy* **35**, 77–83 (2016).
38. J. Demant, R. Munksgaard, D. Décarry-Héту, J. Aldridge, Going local on a global platform: A critical analysis of the transformative potential of cryptomarkets for organized illicit drug crime. *Int. Crim. Justice Rev.* **28**, 255–274 (2018).
39. M. Prince, The trouble with tor. <https://blog.cloudflare.com/the-trouble-with-tor/>. Accessed 13 November 2020.
40. J. S. Mill, *On liberty*, D. Bromwich, G. Kateb, Eds. (Yale University Press, New Haven, 2003), p. 80.
41. N. Marechal, “Use signal, use tor: The political economy of digital rights technology,” PhD dissertation, University of Southern California (2018).
42. Y. Levine, *Surveillance Valley: The Secret Military History of the Internet* (PublicAffairs, 2018).
43. A. Brantly, Banning encryption to stop terrorists: A worse than futile exercise. *CTC Sentinel* **10**, 29–33 (2017).
44. H. Abelson et al., Keys under doormats: Mandating insecurity by requiring government access to all data and communications. *J. Cybersecurity* **1**, 69–79 (2015).