

© The International Institute for Strategic Studies

This content may be used for research and private study purposes. All rights reserved. Any substantial or systematic reproduction, redistribution, reselling, loan, sub-licensing, systematic supply, or distribution in any form to anyone is expressly forbidden.

Full terms and conditions of use: <http://www.iiss.org/terms-and-conditions>

SCROLL DOWN FOR DOWNLOADED CONTENT

Cryptopolitik and the Darknet

Daniel Moore and Thomas Rid

Encryption policy is becoming a crucial test of the values of liberal democracy in the twenty-first century. The trigger is a dilemma: the power of ciphers protects citizens when they read, bank and shop online – and the power of ciphers protects foreign spies, terrorists and criminals when they pry, plot and steal. Encryption bears directly on today's two top threats, militant extremism and computer-network breaches – yet it enables prosperity and privacy. Should the state limit and regulate the fast-growing use of cryptography? If so, how?

In September 2013, the *New York Times* and the *Guardian* jointly revealed *Bullrun*, a \$250-million-per-year programme to make encrypted internet traffic accessible to the United States' and United Kingdom's intelligence agencies. A few weeks later, another story broke: the US National Security Agency (NSA) had successfully intercepted Google traffic; data had been securely encrypted between Google and its users, but sent in clear text between the company's data centres. On a now-famous yellow Post-it note, one NSA spy outlined how to trick Google at the spot where the public internet meets Google's cloud, cheekily drawing a smiley face and scribbling that encryption was 'added and removed here!'. When the *Washington Post* showed the drawing to two engineers close to Google, they 'exploded in profanity'.¹ Even worse, just before Christmas that year, Reuters reported that the NSA had worked with the pioneering security company RSA to

Daniel Moore is a cyber-threat intelligence engineer and a PhD candidate in the Department of War Studies at King's College London. **Thomas Rid** is a professor at King's. His new book, *Rise of the Machines*, comes out in June with W.W. Norton and Scribe.

undermine the standard for a random-number generator, an engine that powers encryption.² Cryptographers were shocked, and trust in the US government evaporated.³ The Crypto Wars of the early 1990s, it became clear, had never ended; the fight simply entered the next round, with stakes raised and gloves off.

Is more encryption better?

On one side of the renewed rift are those who argue that more cryptography can only be beneficial. Bring up the strength of encryption, the number of applications using it by default and the volume of encrypted traffic, and the benefits are universal: to democracy, commerce, privacy, human rights and even cyber security. This belief unites the Electronic Frontier Foundation, the American Civil Liberties Union and a range of Silicon Valley companies and privacy start-ups. 'Pretty Good Privacy means pretty good society', wrote Kevin Kelly, one of the founders of *Wired* magazine, in 1993.⁴ Indeed, as envisioned then, many of the biggest websites today are fully encrypting their traffic, displayed to the user through a small green lock symbol in the browser's address bar. The leaks of NSA materials by former contractor Edward Snowden revealed that encryption is still secure; the problem was bad system architecture. The NSA's cheeky smiley face nudged the big internet companies to fix the flawed architecture and to encrypt internal traffic. More encryption, in short, is better.

On the other side of the argument are governments, law-enforcement agencies and intelligence services. They argue that more encryption that they cannot break can only be worse. 'Terrorists, criminals and hostile states increasingly use encryption to protect their communications', the UK Parliament's Intelligence and Security Committee observed in a major report published in March 2015.⁵ Prime Minister David Cameron had already set the tone in January, saying, 'In extremis, it has been possible to read someone's letter, to listen to someone's call, to listen in on mobile communications ... The question remains: are we going to allow a means of communications where it simply is not possible to do that? My answer to that question is: no, we must not.'⁶

The rise of the Islamic State (also known as ISIS or ISIL), and especially the 13 November 2015 attacks in Paris, have prompted more calls to deny cryptographically protected virtual safe spaces to terrorists. 'There's no doubt that use of encryption is part of terrorist tradecraft now', FBI Director James Comey told the US Senate Judiciary Committee on 9 December 2015. 'Increasingly, we are unable to see what they say, which gives them a tremendous advantage against us.'⁷ More encryption than already implemented, in short, is worse.

Who is right? Is more encryption in the interest of liberal democracy, or not? Both positions are overstated and flawed. To see how, and to overcome this stalemate, three hard questions need to be answered. Firstly, is there a difference between liberal and illiberal cryptographic architectures? Secondly, if more 'crypto' is not always better, and if less 'crypt'⁸ is not always worse, an even harder question follows immediately: where should the line be drawn between desirable and undesirable properties of crypto systems? In short, what is the difference between good encryption and bad encryption? And what should be done as a consequence?

These questions are not primarily moral or normative. Liberal societies have already made their defining moral choices in the past 250 years. Free speech is protected in all liberal democracies. Extreme forms of militancy and terrorism, incitement of violence, the abuse of children, child pornography, fraud, money laundering and unrestricted trade in drugs and arms are illegal in all liberal jurisdictions. The internet and associated technologies do not mean entirely new choices are necessary – they require a reaffirmation of established moral choices in a new technical reality. The question is therefore primarily an empirical one: do specific cryptographic architectures encourage more illegitimate than legitimate behaviour?⁹

We argue that more encryption is better most of the time, but not all of the time. Not all applications of cryptography are beneficial to a liberal order. Crypto systems are not politically neutral; they embody political choices. In some cases, the costs incurred by establishing a specific cryptographically enabled service may outweigh the benefits. We test and establish this argument by critically assessing one of the most sophisticated and controversial encryption platforms today: the Tor Project. If there is a line that demarcates

liberal from illiberal cryptographic architectures, it runs right through Tor. To be more precise, it runs right through hidden services.

Five properties of encryption

The cut-throat crypto confrontation between the US government and non-governmental privacy advocates dates back to the discovery of public-key encryption in 1976 (or perhaps even to its secret discovery in the late 1960s). The face-off escalated in the early 1990s. To appreciate the causes and the depth of the dispute, it is crucial to understand the remarkable way in which public-key cryptography was discovered. The story started when modern telecommunication clashed with the so-called key-distribution problem. For many centuries, this puzzle had hampered efforts to send messages securely on an insecure channel. To illustrate the problem, cryptographers usually introduce a fictional Alice and Bob as the two parties to a conversation, and prying Eve as the eavesdropper. One of the most basic problems of anybody trying to encrypt a message is how first to share a secret. Bob, the recipient, needs Alice's secret key to decrypt her message. How do Alice and Bob share the secret key without Eve intercepting it? This is the problem of key distribution.

As the development of telecommunication technology led to an increase in the number of participants in conversations, the key-distribution problem became more pressing. By the 1960s, it had begun to affect military use of tactical radio. 'The management of vast quantities of key material needed for secure communication was a headache for the armed forces', recalled James Ellis, a pioneer at the UK's Government Communications Headquarters (GCHQ) in Cheltenham, who first articulated a novel possibility to solve the problem in the late 1960s, in secret.¹⁰ A few years later, Stanford University's Whitfield Diffie and Martin Hellman foresaw the impending key-distribution quandary, not in the context of tactical radio, but of networked computers. They predicted, correctly, that 'telecommunications' would soon replace 'most mail and many excursions'.¹¹ This meant that physically sharing unique keys among a vast number of pairs was unrealistic. 'The cost and delay imposed by this key distribution problem is a major barrier to the transfer of business communications to large teleprocessing networks', they wrote in 1976.¹²

These pioneers in Cheltenham and Stanford had the same solution: there was no reason why only the sender of a message could encrypt it; the recipient could take part in encrypting the message. This was counter-intuitive, even revolutionary. But the mathematics were sound. Once a suitable one-way function was found, an algorithm could generate a key pair instead of just one key. The deciphering key was kept private, while the enciphering key was revealed to the recipient – even more, it could be ‘made public by placing it in a public directory along with the user’s name and address’, the pioneers wrote.¹³

After the theorem of existence was formulated – by Ellis in secret, and by Diffie and Hellman in public – an implementation was needed. This implementation was, again, first discovered in secret at GCHQ and then by US academic cryptographers in public a few years later. The desired one-way function was the so-called ‘factorisation’ problem, the mathematical curiosity that multiplying two very large prime numbers is easy, but inferring those two large prime numbers from their far larger product is computationally extremely hard. This provided the necessary one-way connection between the public-and-private key pair: the private key could be derived from the two large prime numbers, and the public key from their hard-to-reverse product.

Public-key encryption was one of the most pivotal inventions of the twentieth century. Over the next quarter-century, it would enable the recreation – and improvement – in electronic form of five fundamental, age-old properties of human communication.

The first property was the most obvious: privacy, the equivalent of an envelope for a letter, a way to protect messages in transit against unauthorised access. In the case of postal mail, message security was accomplished through a sealed envelope. The security was not absolute, but it was good enough for everyday use. Using cryptography in this way was not new, of course; cryptography had been used to secure messages for thousands of years. But unlike envelopes, cryptographic security could not be scaled up to a large number of users, until public-key encryption overcame the key-distribution problem.

The second property that public-key encryption recreated electronically was authentication, the equivalent of a handwritten signature on a letter:

a way to prove that a message comes from a specific sender. In the case of postal mail, a signature in ink guaranteed with a reasonable degree of certainty that the right person signed a letter, a method that worked at scale, for a very large number of ‘users’ of signed sheets of paper. Ronald Rivest, Adi Shamir and Leonard Adleman, the MIT pioneers who suggested factorisation, were clear: ‘The era of “electronic mail” may soon be upon us’, they predicted in 1978. ‘We must ensure that two important properties of the current “paper mail” system are preserved: messages are private, and messages can be signed.’¹⁴

Some academic cryptographers have dismissed the secret GCHQ pioneers because the intelligence agencies did not develop ‘non-secret encryption’ into an algorithm or actual product. ‘You did a lot more with it

than we did’, Ellis famously told Diffie in a pub in Cheltenham.¹⁵ But this dismissal is short-sighted.

Cryptography was becoming hip

The real significance of Ellis’s insight was not technical, but political. When *Scientific American* columnist Martin Gardner revealed the cipher magic of prime numbers to the world in summer

1977,¹⁶ the NSA was shocked. At Fort Meade this was not an exciting new discovery – it felt like a bunch of pesky academics had stolen their precious secret. The US government would therefore attempt to stop the spread of encryption, even with desperate means, including threats, lawsuits, funding, censorship, export controls and treating algorithms as weapons.¹⁷ Meanwhile, cryptography was becoming a hip thing to study, with many long-haired and bearded geeks talking about mathematical one-way functions and trap doors. The line between research and activism began to blur. Soon, this new crop of cryptographers discovered a third and fourth property needed for all things digital.¹⁸

The third property was anonymity, the equivalent of not writing one’s sender address on a letter in the first place, as a way of hiding one’s identity from the recipient and other possible observers. Cryptographers speak of ‘traffic analysis’ when they sift communications and related material for patterns that could reveal users’ identities, or more. In the case of electronic mail, traffic analysis could be made harder without encryption, simply by

stripping the sender address from an email, at least at first. Early privacy activists, known as ‘cypherpunks’, achieved this through what they called ‘remailers’. Remailers were specialised email servers: when they received a message, the server kept the recipient address but replaced the sender address with an unusable one, and then forwarded the anonymised email to the intended recipient. It was as if somebody at the post office had taken a pair of scissors and snipped off the sender’s address.

In 1981, David Chaum, a 26-year-old mathematician at the University of California, Berkeley, formulated a cryptographic method to defeat traffic analysis, and thus to create anonymity, not by snipping off the address field, but by hiding the letter in layers of encrypted envelopes and passing it through computers that he called ‘mixes’.¹⁹ Chaum’s idea was the inspiration for ‘onion routing’, most prominently implemented in the Tor Project.

A fourth privacy-preserving property of paper was missing, and that was cash. Banknotes, like letters without a sender’s address, were anonymous. Using cash, a punter could walk into a liquor store or gun shop, say, and buy a product without leaving a trace. As credit-card payments became more and more common, offline and later online, transactions became alarmingly traceable. ‘The foundation is being laid for a dossier society’, wrote Chaum in 1985. Computers, the cryptographer feared, ‘could be used to infer individuals’ life-styles, habits, whereabouts, and associations from data collected in ordinary consumer transactions’. To avoid that, Chaum suggested a payment-transaction system based on ‘blind signatures’.²⁰ He went on to start a company to develop an actual digital currency, Digicash, and in 1994 made the world’s first cryptographic payment.²¹

So far, public-key cryptography had a conservative promise: it could preserve the cherished anonymity afforded by established, non-electronic social interaction, such as envelopes, signatures, anonymity and cash. The magic of large prime numbers did not just recreate these properties of human communication, it improved them, making them safer, more scalable and more efficient. But there was no reason to stop there.

The fifth property had the most explosive potential: hidden exchanges. Suddenly, it became possible to create online exchange and marketplaces in which transactions were secure, authenticated and anonymous. But these

advantages did not just apply to the buyer, as with a regular cash transaction in a brick-and-mortar shop with a street address. Anonymity and cash already served that purpose. Now, it became possible for the shop itself to remain hidden, both from the buyer and from any other third party, such as a service provider, tax authorities or law enforcement.

As with the first four properties, brick-and-mortar equivalents had long existed for hidden exchanges, where sellers and service providers opted to remain anonymous. There have long been some such exchanges, with two-way anonymity, that are non-illicit, including swap meets, some street-food vendors, small farmers' markets and irregular flea markets that do not require application and registration, as well as some bazaars in developing countries. But particularly in developed countries, non-digital hidden markets tend to be predominantly illicit: public spaces, back alleys and hotel rooms in which to buy drugs, weapons, sex, and other goods and services, such as gambling, fraudulent material, money laundering and undocumented labour.

Cryptography recreated a space for such hidden markets, both licit and illicit, except it did so in an improved form. Seller and buyer would never see each other's faces, and the risk of being physically assaulted or apprehended, either by gangs or by the police, would be significantly reduced. Some argued that cryptography promised to take violence itself out of crime.²²

The first vision of such a cryptographically protected hidden exchange was BlackNet, a half-prank, half-serious announcement that led to an FBI investigation. An 'Introduction to BlackNet' was delivered via email to the cypherpunks list, through a remailer, signed nowhere@cyberspace.nil. It started ominously:

Your name has come to our attention. We have reason to believe you may be interested in the products and services our new organization, BlackNet, has to offer. BlackNet is in the business of buying, selling, trading, and otherwise dealing with *information* in all its many forms.²³

The anonymous author made clear that BlackNet would use public-key crypto systems to guarantee perfect security for customers. The marketplace

'in cyberspace', the message made clear, would have no way of identifying its own customers, 'unless you tell us who you are (please don't!)'. BlackNet's suggested messaging system was ingenious: 'we can be contacted (preferably through a chain of anonymous remailers) by encrypting a message to our public key (contained below)'. BlackNet, it said, monitored 'several locations in cyberspace', such as the *alt.extropias* Usenet newsgroup and the cypher-punks list. Posting an encrypted message to a monitored public space meant that only the single intended holder of the private key could read it.

BlackNet was 'nominally' non-ideological, but the email's author made clear that he considered nation-states, export laws, patent laws and national security to be 'relics of the pre-cyberspace era'. These things simply served the nefarious purpose of expanding the state's power, furthering what BlackNet called 'imperialist, colonialist state fascism'. The curious email pamphlet then made clear that BlackNet would build up its 'information inventory', and that it was interested in acquiring a range of commercial secrets. 'Any other juicy stuff is always welcome', the anonymous voice added. 'Join us in this revolutionary – and profitable – venture.' BlackNet, the then-notorious half-prank, predated the actual darknet by almost ten years – yet, in a highly revealing twist, the actual inspiration for the darknet was rather different.

Hidden services

'Darknet', colloquially, refers to a distinct network supporting cryptographically hidden sites. Four Microsoft researchers popularised the term in a seminal 2002 article. Their case study was digital piracy and the gradual development of secure peer-to-peer networks as alternatives to internet sites that could be easily censored.²⁴ Piracy thrived in the dark, as hidden hosts were harder to sue. Modern darknets use unique software to allow use of the distributed network. The most notable examples today are Tor, I2P and Freenet. The fluid architecture of these networks makes estimating their size difficult, but it appears that Tor is the largest, with I2P a distant second.²⁵ Others are significantly smaller in scope and popularity.²⁶

The Tor architecture provides two services – anonymous browsing (property 3), and hosting of anonymous information exchanges (property

5) – through one piece of software, the so-called ‘Tor Browser’. Although distinct, both services employ roughly the same protocols and rely on the same distributed infrastructure. But that is where their mutual dependency ends. There is no technical requirement for anonymous browsing and anonymous hosting to be bundled. Indeed, browsing is overwhelmingly more popular than hosting. Most Tor users have never visited any hidden website at a *.onion address; hidden services account for around 3–6% of overall Tor traffic.²⁷ Most users instead use the software merely to browse the internet’s conventional address space more securely or anonymously. An analogy illustrates the significance of anonymous browsing. Alice, who lives in a

*Anonymous
browsing is
legitimate and
laudable*

small town, wants to buy a pregnancy test, but doesn’t want to be seen doing so by the shop owner, Bob, a friend of Alice’s father. Rather than simply going to the store, Alice wears a mask, walks a detour, and pays in cash. Bob will not be able to identify her or trace her. Alice’s privacy and anonymity are assured. Anonymous browsing is not part of the ‘dark web’; it is a legitimate and laudable service that Tor provides.

Tor in its entirety originated as a collaborative project between the US Naval Research Laboratory and the non-profit organisation Free Haven Project. The underlying purpose was to create a distributed, anonymous, easily deployable and encrypted network to be used by those who needed it.²⁸ Specifically, it was offered as a free service to promote unfettered access to the internet in locations where online censorship was heavily enforced or where the threat of persecution for those who sought access to locally illegal information was prohibitive.²⁹

Onion routing, inspired by Chaum, was the solution. In order for a user to securely access a website without being identified or traced, he or she would instead be routed through a series of intermediary servers. The resulting pathways between servers were labelled ‘circuits’, in Tor jargon.³⁰ Each packet of information to be relayed over the network would be encased in multiple layers of encryption, each to be sequentially peeled away only by the subsequent node in the circuit. Consequently, intermediary nodes could only decrypt one layer of the encryption, preventing access to the

underlying data and its originator. The final such hop – or exit node – would reveal the original packet and proceed to deliver it to the desired destination, thus protecting the sender's identity. As a result, intercepting and decoding the information along its path would be significantly harder – albeit not impossible – to accomplish.

Over time, civilian researchers and government agencies successfully de-anonymised some users, through methods ranging from planting compromised exit nodes that recorded traffic to employing malicious code within websites to covertly force users to access a public internet address controlled by the attacker, thereby revealing their true IP address.³¹ But if a user employs even a fairly rudimentary set of cautionary procedures (such as keeping the browser up to date), the Tor core architecture remains relatively secure. Efforts to identify vulnerabilities in the Tor platform and repair them are continuous, and the architecture is gradually improving over time.³²

Tor's anonymisation function has received widespread support and praise. Google, Human Rights Watch and the Electronic Frontier Foundation advocate Tor browsing and recommend its use by dissidents in circumventing repressive government measures.³³ A prime example of this usage became evident in 2011 in Egypt, as thousands successfully used the Tor browser to communicate and disseminate information in spite of a severe clampdown on the internet instigated by the Mubarak regime.³⁴ A second noteworthy instance is the use of Tor browsing by the rebels in strife-stricken Syria, as they scrambled to release digital evidence of atrocities committed by the regime by way of the internet without exposing the disseminators to what likely were feverish efforts to uncover their identities.³⁵

Tor, however, does not stop there. The network enables a far more controversial property as well.³⁶ This capability, called a hidden service, allows anybody to create a virtually untraceable server hosted within the Tor network, simply by adding two short lines of code to a short configuration file.³⁷ This allows circumvention of all known forms of content restrictions or surveillance. Neither the Internet Service Providers (ISPs) that route the traffic, nor law-enforcement agencies, nor even the developers of the Tor project itself have visibility into the hosted service's location, or the identity of its operator.³⁸

An analogy is again useful. Alice's job in finance is very stressful and she wishes to buy a bag of cocaine. Bob has one to sell, but understands the risks. Alice and Bob agree through an intermediary on a trunk sale at a rendezvous point in a little-used back alley at 11pm. Both arrive with their number plates covered, their faces hidden in the dark. Alice buys the cocaine, pays in cash and disappears quickly into the night. Both Alice and Bob have remained anonymous and safe; neither knows the other's identity, and they have left no trace.

Hidden services were designed to be untraceable. The *.onion sites do not correspond to an IP address that could be located anywhere. To access a website within Tor, Alice must first know its unique address (always concluding with the hallmark *.onion suffix). Once the connection is initiated, both the server and the user establish encrypted circuits to a neutral node within the Tor network – known as a rendezvous point – through which all communication of a particular session will take place. Hence, connecting to a hidden service ensures the anonymity and privacy of both users – and, remarkably, the service provider – unless, of course, either party decides to reveal his or her identity.

Into the dark

In order to show how Tor hidden services are used in practice, as opposed to how they fuel the fears and hopes of players in the renewed crypto wars, we carried out an in-depth scan of hidden-services websites within the Tor network. We collected data through a website crawler (an automated scanning software, which hops from website to website by following links) specifically tailored to crawl Tor web-based hidden services. We assessed the results in two steps: firstly, we manually classified an initial batch of websites into categories and used them to train a Support Vector Machine document classifier, a statistical classification algorithm often used in machine learning to categorise content. Secondly, we then used this automated classifier to complete the categories for the rest of the web pages.

The Tor darknet is designed to avoid a central stable repository of existing sites. In contrast to the conventional internet, there are no easy website registries where one might look up information on who is managing what

website and where they are registered as doing so. As previous research has shown, however, HTTP-based hidden services (as opposed to less numerous email providers or chat rooms) consistently number in the low thousands, and many websites are unstable. Methods other researchers have used include operating directory servers,³⁹ conducting traffic analysis to determine popularity,⁴⁰ and performing a narrow web-crawl.⁴¹ We chose a different approach: an in-depth, lengthy web-crawl of every web-based hidden service reasonably accessible by an individual seeking content within the darknet. Consequently, our resulting dataset includes the vast majority of websites one might have reasonably encountered during the given period, although there is no way to index every existing hidden site. This methodology is bound to have missed a small number of very closely held sites, but the number of sites we mapped roughly matches the estimated number of Tor websites at any given moment.⁴² Most hidden services, indeed, seem to welcome visibility. The beauty of the system architecture, after all, is that a public site will not reveal its location or operator.

Tor hidden services even have an unofficial, limited search functionality. Several indexes provide a significant, if incomplete, roster of websites hosted within the Tor darknet. We took website lists from two such services, onion.city and ahmia.fi. The two sites yielded 5,615 total unique *.onion addresses, enough for an initial seed for a more detailed crawl.⁴³ The entire seed list of websites was crawled during a two-month period between January and March 2015. Our crawler used standard link traversal (following links like a human user would), specifically tailored to only follow *.onion addresses to limit the results to hidden services. Our method went beyond previous Tor-mapping research, by going up to five hops 'deep' into a targeted site and scanning a maximum of 100 pages within each site (instead of thousands from a few prolific sites, thus preventing bias while maintaining statistical diversity). We thus retrieved more data from more pages.

In order to avoid illegal material, such as media files of child pornography or publications by terrorist organisations, only textual content was harvested automatically. Any other material was either filtered out or immediately discarded. Our methodology relied exclusively on text input, but it is unlikely that a classification based on image or media files would

have yielded divergent results, so this automatic pre-selection should not affect the significance of our findings.

We divided analysis into two consecutive steps. The first phase was primarily a necessary precursor to the second, and included manual categorisation of websites. We randomly sampled the harvested data and assembled a collection of web pages that was then used to train the automatic Support Vector Machine content classifier. This produced a taxonomy of 12 high-level categories (see Figure 1, below). In the second phase, the trained classifier assigned the rest of the documents to these 12 categories, the results of which were again randomly inspected for potential errors, to ensure the fidelity and accuracy of the classification. Each hidden website's category was deduced by the aggregate classification of all web pages retrieved under that domain, rather than just the home page as commonly done in other studies.⁴⁴

Figure 1. **Taxonomy**

Category	Details
Arms	Trading of firearms and weapons
Drugs	Trade or manufacture of illegal drugs, including illegally obtained prescription medicine
Extremism	Content espousing extremist ideologies, including ideological texts, expressions of support for terrorist violence, militant how-to guides and extremist community forums
Finance	Money laundering, counterfeit bills, trade in stolen credit cards or accounts
Hacking	Hackers for hire, trade or distribution of malware or DDoS ⁴⁵ capabilities
Illegitimate pornography	Pornographic material involving children, violence, animals or materials obtained without participants' consent
Nexus	Websites primarily focused on linking to other illicit websites and resources within the darknet
Other illicit	Materials that did not easily fit into the other categories but remain problematic, such as trade of other illegal goods and fake passports or IDs
Social	Online communities for sharing illicit material in the form of forums, social networks and other message boards
Violence	Hitmen for hire, and instructional material on conducting violent attacks
Other	Non-illicit content, such as ideological or political content, secure drop sites, information repositories, legitimate services
None	Websites which were either completely inaccessible or otherwise had no visible content, including websites which hosted only placeholder text, indicating that their operator had yet to generate indicative content

Results

The results, as shown in Figure 2, were clear. The database analysis brought to light the overwhelming presence of illicit content on the Tor darknet. The scans returned a total of 5,205 live websites, out of which 2,723 were

successfully classified into the above categories with a high degree of confidence. ('None' indicates a lack of content, and was thus counted as neither licit or illicit, while 'Unknown' means that content was there but we were unable to determine its nature, as it was too sparse or illegible.) The outcome of the automatic classification was again manually inspected to ensure accuracy, with good results. Overall, our crawler accessed roughly 300,000 addresses within the Tor hidden-services network, yielding a highly diverse and significant corpus of data in the form of 205,000 unique pages.

Figure 2. **Classification**

Category	Websites
None	2,482
Other	1,021
Drugs	423
Finance	327
Other illicit	198
Unknown	155
Extremism	140
Illegitimate pornography	122
Nexus	118
Hacking	96
Social	64
Arms	42
Violence	17
Total	5,205
Total active	2,723
Total illicit	1,547

The results suggest that the most common uses for websites on Tor hidden services are criminal, including drugs, illicit finance and pornography involving violence, children and animals. One noteworthy finding was our confirmation of the near-absence of Islamic extremism on Tor hidden services, with fewer than a handful of active sites. Jihadis tend to use the internet for at least two general purposes: public-facing activities (propaganda, recruitment and sharing advice) and non-public-facing activities (internal communication, and command and control). The darknet's propaganda reach is starkly limited, not least because novices may be deterred by taking an 'illicit' step early on, as opposed to simple, curious Googling. Hidden services, secondly, are often not stable or accessible enough for efficient communication; other platforms seem to meet communication needs more

elegantly. Islamic militants do commonly use the Tor browser on the open internet, however, for added anonymity.⁴⁶

The financial category comprised primarily three prominent sub-categories: Bitcoin-based methods for money laundering, trade in illegally obtained credit cards and stolen accounts, and trade in counterfeit currency. As is often the case with unidentified vendors, the quality of the services offered was very difficult to ascertain. The sellers tried to sound professional, as in this fairly common boast about the quality of counterfeit notes:

Our notes are made with the highest quality cotton fibre, all security features are included: watermarks, security thread, microprint, magnetic ink, color shifting ink, etc.⁴⁷

Various websites offering cloned credit cards or financial information obtained via malware are already commonplace on the internet, and now find a willing home on the darknet, where sellers can more easily evade detection and monitoring. Online criminal communities under pressure from law enforcement now frequently retreat to Tor, where security and anonymity are greater.⁴⁸ Websites in this sub-category range from wholesale trade in the details of compromised credit cards to the outright distribution of the malicious software used to acquire them.

Many of the sites we examined offered services for laundering money through Bitcoin. Bitcoin is the most common currency employed in all Tor hidden-services trade, often via reliance on third-party escrow services to alleviate concerns stemming from anonymous, unverifiable transactions between two unscrupulous parties. As Bitcoin transactions can be monitored even if not easily de-anonymised, however, services to blur the trail of Bitcoins have proliferated as well, for a nominal transaction fee.⁴⁹ One such site promised that

When you use CleanCoin to mix your Bitcoins, you will receive Bitcoins that originate from lots and lots of different transactions and wallet addresses, making it almost impossible for someone to track your wallet activity.⁵⁰

A wide variety of drugs – both pharmaceutical and otherwise – constituted the single most common commodity within the Tor darknet. A multitude of vendors exist, ranging from communal marketplaces, such as the fairly notorious Agora, to single-page vendors offering a limited selection of their own making. The drugs on offer range from marijuana, cocaine, methamphetamines and various forms of acid, to more discerning markets such as those trading in anabolic steroids and Viagra-type medication. One dubious example: ‘Kamagra is the preferred alternative to Viagra for customers wishing to use the generic version of this popular treatment for impotence and erectile dysfunction.’⁵¹ Many of the vendors appear to be fraudulent, with customers frequently complaining of being swindled by sellers, or ‘ripped’ in the community jargon.

The pornographic content was perhaps the most distressing. Websites dedicated to providing links to videos purporting to depict rape, bestiality and paedophilia were abundant. One such post at a supposedly non-affiliated content-sharing website offered a link to a video of ‘a 12 year old girl ... getting raped at school by 4 boys’.⁵² Other examples include a service that sold online video access to the vendor’s own family members:

My two stepsisters ... will be pleased to show you their little secrets. Well, they are rather forced to show them, but at least that’s what they are used to.⁵³

Several communities geared towards discussing and sharing illegitimate fetishes were readily available, and appeared to be active. Under the shroud of anonymity, various users appeared to seek vindication of their desires, providing words of support and comfort for one another in solidarity against what was seen as society’s unjust discrimination against non-mainstream sexual practices. Users exchanged experiences and preferences, and even traded content. One notable example from a website called Pedo List included a commenter freely stating that he would ‘Trade child porn. Have pics of my daughter.’⁵⁴ There appears to be no fear of retribution or prosecution in these illicit communities, and as such users apparently feel comfortable enough to share personal stories about their otherwise stifled tendencies.

Other commonly available illegal services included procurement of fake identification documents; stolen or otherwise illegally obtained equipment, firearms and associated peripherals; and malicious software used to harvest personal accounts or perform denial-of-service attacks against websites. Several websites supposedly offered assassinations for a hefty fee. For example:

We are a team of 3 contract killers working in the US (+Canada) and in the EU. Once you made a “purchase” we will reply to you within 1-2 days, contract will be completed within 1-3 weeks depending on target. Only rules: no children under 16 and no top 10 politicians.⁵⁵

Unlike other commodities and services, however, there was no indication that anyone had successfully contracted the supposed contract killers, and their credibility remained largely unconfirmed.

Many websites classified as ‘Other’ were in fact about the darknet itself. These ‘meta-sites’ included hosting services and tutorials on using Tor. Alongside these meta-sites, the ‘Other’ category included personal blogs, several journalist drop sites and other innocuous services.

The classifier most frequently assigned the category ‘None’. The abundance of failed websites and those with no content is a stark indication of how unreliable the hidden-service platform is for most users. An overwhelming number of websites either displayed an empty page or a server-generated placeholder, returned a server error, or did not respond at all to repeated requests. While they were technically active, they were assigned to the ‘None’ category as they held no meaningful content.

Many websites were inaccessible beyond their home page because they required users to log in. These reclusive sites were not classified automatically, and were only assigned a category if such was readily apparent from examining the home page. Otherwise, they were consigned to the ‘Other’ category, given the benefit of the doubt as non-illicit if their purpose seemed ambiguous.

Lastly, illicit and non-illicit Tor websites differed in one stark and important way: legitimate sites almost always chose to identify their operators,

while illicit hidden services almost always chose not to do so. Examples of self-identification included hidden services run by newspapers, blogs, search engines and even Facebook. In other words, the most significant variable distinguishing illicit from non-illicit service providers is whether they use onion sites to hide behind anonymity, or to take advantage of some of the platform's other security benefits. This notable difference suggests a potential way forward.

Political technology

Mathematics may be pure, but implementations embody moral values and political choices, and these choices can either advance or undermine liberty. Developers bear responsibility for their creations. In this sense, the history of hidden services is curiously upside down. In the Crypto Wars of the 1990s, Timothy May, a gun-loving anarchist from Santa Cruz, California, foresaw the illicit future of cryptographically hidden exchanges, yet the mathematicians and computer scientists at the US Naval Research Laboratory who developed Tor hidden services a decade later chose to neglect these known risks.

In 1997, when one user asked about potential abuse, one of Tor's core developers shrugged it off: 'That is a political question', he wrote. 'We have tried to only deal with the technological issues instead of the political ones. As soon as we start dealing with political issues, this thing will fall apart.'⁵⁶ As late as 2007, with hidden services up and running, the platform developers were still oblivious, and still shunning political considerations. 'Simple technical mechanisms can remove the ability to abuse anonymously without undermining the ability to communicate anonymously',⁵⁷ two of them wrote, in a remarkably optimistic overstatement.

By contrast, May, who had come up with the BlackNet vision and later claimed credit for it, tackled the political quandaries head-on. 'Just because some folks mis-use free speech is no reason to ban free speech', he wrote in defence of encryption in 1994. 'And just because some will mis-use encryption – in the eyes of government – is not a good reason to ban encryption.' Recalling the US constitution's second amendment, which protects the right to bear arms, he coined the memorable phrase: 'Crypto

= Guns'.⁵⁸ May articulated a pre-emptive defence of his bold idea back in February 1994, long before the first such hidden marketplace appeared in the wild:

They'll cite the 'unpopular' uses: kiddie porn nets, espionage, selling of trade secrets (especially to 'foreigners'), the bootlegging of copyrighted material, 'digital fences' for stolen information, liquid markets in liquidations.⁵⁹

He was right. Governments would indeed cite all these 'unpopular' uses, along with terrorism. And he already had a canned libertarian answer for them, akin to free speech and, indeed, gun ownership in the US: 'Just because some people mis-use camcorders to film naked children is no reason to ban cameras, camcorders, and VCRs.'

May's foresight was politically savvy, but it was technically primitive, even naive. The cypherpunks lumped crypto into one box, not anticipating that different architectures would serve different purposes, as they do today. The debate has evolved in the 20-odd years since the activists were known to boast that 'you can have my crypto keys when you pry them from my cold, dead hands'.⁶⁰

All five cryptographically recreated properties – security, authentication, anonymity, digital currencies (to be more precise, 'blockchain' technology) and hidden exchanges – can be used or abused. Most forms of encryption have become a bedrock of the modern internet and the ubiquitous Public Key Infrastructure, or PKI. Encryption has been widely acknowledged as crucial to the protection of free speech, privacy and commerce. Even GCHQ supports the use of encryption to improve security and protect business.⁶¹ The UK government encrypts almost all its websites, and the US government is planning to do so in the near future. The fourth function, crypto currencies, is more controversial, but is quickly becoming mainstream. Most liberal democratic governments no longer wish to pry crypto keys from anybody's cold, dead hands, despite a sometimes shrill discussion about 'back doors'.⁶² For economic, security and ethical reasons, most liberal democracies have accepted some widespread implementations of encryption as a beneficial and necessary development.⁶³

The implementation of hidden services, by contrast, was technically sophisticated but politically naive. Onion routing was originally designed not for what it was then actually used for – it was designed to provide availability against denial-of-service attacks, as well as against physical attacks. None of the original designers ever mentioned illicit marketplaces. Paul Syverson, one of the key developers, said his inspiration was ‘The Eternity Service’, a paper by Ross Anderson, a Cambridge University security engineer.⁶⁴ At an October 1996 conference in Prague, Anderson had suggested an innovative way to use crypto to protect the availability of information against both electronic and physical attacks, not just the confidentiality or integrity of information, which, he lamented, had been the misguided focus of much computer-science security research for years.⁶⁵ Another inspiration was David Chaum, who articulated the abstract principle of onion routing in 1981: ‘Our onion routers are based on mixes’, the developers of onion routing wrote in 1997, after acknowledging Chaum’s pioneering work.⁶⁶

In the 2010s, however, defending a network’s eternal availability is no longer Tor’s priority. The priority instead is to hide illegal exchanges. As our research has demonstrated, the vast majority of non-self-identified service providers peddle goods and services that are illegal not just in the most restrictive jurisdictions, but even in the most liberal jurisdiction imaginable. The future of Tor hidden services was May’s BlackNet, not Anderson’s Eternity Service. The Tor developers wanted to enlighten, but created darkness instead.

Saving crypto from itself

Yet we should not rush to judgement too quickly. Tor hidden services, in their present form, combine two separate features. The first feature, hiding the physical location of all parties that are communicating, is a technical consequence of the onion-routing protocol – the trunk-sale effect. But the second feature, hiding the identity of the host, is a choice on the part of each individual service provider. Identities can be revealed, naturally, without losing the platform’s security features, as one of Tor’s most significant pioneers argues.⁶⁷ The trunk sale, in short, doesn’t have to happen in the dark without number plates.

For users of the DuckDuckGo search engine, of ProPublica's hidden site, of the various SecureDrop sites (offered by outlets such as *Wired*, the *New Yorker*, the *Sun*, the *Guardian* and the *Washington Post*)⁶⁸ or indeed of Facebook, *.onion sites offer two noteworthy security benefits that take us back to the two core features that gave rise to public-key encryption in the first place: better privacy and better authentication. Hidden services increase privacy because there is no longer an exit node: both parties to the transaction are within the Tor network. Secondly, the protocol improves authentication, by making a so-called man-in-the-middle attack significantly more difficult. Facebook, for instance, at <https://facebookcorewwi.onion>, is not just revealing the social network's identity on the dark web; it is guaranteeing it while protecting the user. The engineer who developed Facebook's Tor branch pointed out that authentication in the Tor network is even better than on the more widely used internet protocols.⁶⁹ The same applies, for example, to the *New Yorker's* SecureDrop site, called 'Project Strongbox'. Anybody visiting <http://strngbxhwyuu37a3.onion> in the Tor browser has a very high certainty that the site is actually affiliated with the *New Yorker*.⁷⁰

Proponents of hidden services argue that the cryptographic protocols that power the internet today were at the fringe of software development and considered a threat as late as 1995. Hidden services, they argue, are what https was 20 years prior: the future of security, not a threat to security. These arguments are strong, and cannot be dismissed; the technology may well mature and move into the mainstream in the future. But the crypto purists, Tor's developers among them, often fail to acknowledge an even more fundamental point, one that is deeply rooted in the recent history of cryptography: enhanced privacy, enhanced authentication and enhanced user anonymity are not tied to the service or content provider remaining anonymous and unregistered. Our first four properties – security, authentication, user anonymity and cash (or blockchains) – are entirely disconnected from the fifth: unidentified hidden exchanges. These issues are conceptually, politically and technically distinct. Facebook's own use of hidden services is the best example.

Despite their unquestionable technical benefits – or rather because of their technical benefits – Tor hidden services present a formidable political

risk to cryptography itself. The founders are aware of this risk. Even Roger Dingledine, one of the original Tor developers, has expressed ambivalence. ‘Why not scrap hidden services?’, an anonymous user asked him in late December 2014. ‘We do think about that option periodically’, Dingledine responded.⁷¹ Tor pioneer Paul Syverson has expressed similar doubt, and highlighted a better use of *.onion sites: authentication, a property ‘largely orthogonal to hiding server location’.⁷²

The original developers see hidden services as a woefully underdeveloped technology, still in search of its proper uses.⁷³ With hidden services, the proverbial dark alleyways, crypto idealism overshot its target, even though May prominently articulated BlackNet a decade before the Tor darknet came online. The warning was clear, yet system architects ignored it, blinded by the idealistic view that more sophisticated cryptographic implementations could only be a good thing, always, in all forms and shapes. If crypto equals guns, in May’s provocative phrase, then hidden services are the equivalent of shoot first, ask questions later. Too many activists treat cryptography as if it were a godlike force for good. This naivety is embodied in a remarkable line from Edward Snowden: ‘Let us speak no more of faith in man, but bind him down from mischief by the chains of cryptography.’ Snowden wrote these pompous words to his initial journalist contacts,⁷⁴ repurposing a well-known line by Thomas Jefferson, from November 1798: ‘In questions of power, then, let no more be heard of confidence in man, but bind him down from mischief by the chains of the Constitution.’⁷⁵ Man is fallible, Snowden implied; maths is not. ‘The universe believes in encryption’, wrote WikiLeaks founder Julian Assange in *Cypherpunks*, a rambling 2012 book that celebrated the cult of crypto.⁷⁶ The political consequences of such idealism are grave. The widespread and highly visible abuse of unidentified Tor hidden services provides an easy target for any critic of encryption. Even the Islamic State has a propaganda site provided as a Tor hidden service, launched in November 2015.⁷⁷

*Crypto idealism
overshot its target*

If the debate is of such disappointing quality even in the United States and United Kingdom, the very countries that have been developing and

discussing crypto systems for 40 years, then we should expect much worse in fast-developing authoritarian countries over the coming decades. China's government is already blocking Tor.⁷⁸ Russia's Safe Internet League, a powerful censorship organisation, has described all of Tor as an 'invisible Internet' that would enable criminals 'to hide their actions from the authorities in order to commit crimes: trading drugs and weapons, distributing child pornography, engaging in human trafficking and leading political campaigns'. Vadim Ampelonsky, the press secretary of Roskomnadzor, the Kremlin-controlled media regulator, called Tor a marketplace 'where all ghouls gather'.⁷⁹ The ugly truth is that these statements are not entirely wrong. Hidden services give not just Tor, but encryption more generally, a bad name, for very little benefit. China's and Russia's policies on encryption cannot easily be dismissed – they are shaping the online reality of 1.5 billion people, around double the population of the US and Europe combined. These two powerful countries are also setting de facto standards for a large number of non-democratic but fast-developing countries. The global future of encryption is wide open, and hotly contested.

* * *

Liberal democracies are still leading the way; it is therefore of paramount importance that we get it right. Encryption is too important to be left to true believers. The future design of crypto systems should be informed by hard-nosed political and technical considerations. A principled, yet realistic, assessment of encryption and technology more broadly is needed, informed by empirical facts, by actual user behaviour and by shrewd statecraft – not by cypherpunk cults, an ideology of technical purity and dreams of artificial utopias. Pragmatism in political decision-making has long been known as *realpolitik*.⁸⁰ Too often, technology policy has been the exception. It is high time for *cryptopolitik*.

Several hard questions deserve serious discussion, and should be on the minds of pragmatic and politically astute developers as they work on rolling out new services to a larger user base. Can we envision an architecture that guarantees the first four properties recreated by encryption without ena-

bling the fifth? Can a platform be developed that improves anonymity, authentication and availability, but doesn't incentivise illiberal and illegal behaviour at the same time? How can Tor hidden services be modified to get more self-identifying service providers like Facebook in, while pushing more criminals out? And which crypto systems should be fair game for intelligence agencies in democracies?

We are only at the beginning of a serious and difficult debate about cryptopolitik. It is therefore crucial to get straight some initial assumptions on which liberal democracies should be able to agree. Encryption is already implemented as a technology of freedom. A trustworthy public-key infrastructure is a crucial ingredient for any free political order in the twenty-first century. Cryptographic applications should be used (and optimised) to recreate and improve on the prized advantages of centuries-old paper-based systems: confidentiality, integrity, availability, authentication and anonymity, as well as innovative public ledgers and cash powered by highly promising blockchain technology.⁸¹ Yet at least two quandaries stand out.

The first concerns end-to-end encryption, a specific form of encryption that protects users' messages from being read by any third party or platform provider. Any of the basic properties may be further protected in this way. This kind of user-to-user end-to-end encryption can be provided by crypto systems that are commercially developed and maintained by communication-service providers, such as some Apple, Google and Facebook products; or by crypto systems that are publicly distributed and open source, such as PGP, Signal and, indeed, the Tor hidden-service protocol.⁸² Service providers can be presented with a warrant coercing them to remove end-to-end encryption for a given account (or to maintain the capability to do so), but an open-source protocol cannot be so coerced. End-to-end encryption will therefore always be available to a determined, capable user. Moreover, at present, the powerful dynamics of open markets for communication services do not favour end-to-end encryption among individuals at a large scale, thus limiting the technology's wider appeal and uptake. Any attempt to systematically undermine end-to-end encryption – through legislation requiring service providers to retain the option of removing encryption

for any given user – will likely strengthen more secure implementations by creating more demand for them, and thus help criminals and militants. We believe it should be a political no-go area for democratically elected governments to pursue such a path.

The White House seems to have recognised that attempting to systematically undermine end-to-end encryption would be politically fraught, damage security for legitimate users, bring little net security benefit, disadvantage American companies and set a useful precedent for authoritarian regimes.⁸³ The UK government has yet to come to a settled opinion. GCHQ, however, has clearly recognised both the benefits and the challenges of encryption post-Snowden. As a result, among the various techniques on offer, bulk interception is losing significance in relative terms, and computer-network exploitation at the end points of electronic communications – or hacking – is becoming more important.⁸⁴

The other quandary is how to deal with darknets. Hidden services have already damaged Tor, and trust in the internet as a whole. To save Tor – and certainly to save Tor's reputation – it may be necessary to kill hidden services, at least in their present form. Were the Tor Project to discontinue hidden services voluntarily, perhaps to improve the reputation of Tor browsing, other darknets would become more popular. But these Tor alternatives would lack something precious: a large user base. In today's anonymisation networks, the security of a single user is a direct function of the number of overall users. Small darknets are easier to attack, and easier to de-anonymise. The Tor founders, though exceedingly idealistic in other ways, clearly appreciate this reality: a better reputation leads to better security.⁸⁵ They therefore understand that the popularity of Tor browsing is making the bundled-in, and predominantly illicit, hidden services more secure than they could be on their own. Darknets are not illegal in free countries and they probably should not be. Yet these widely abused platforms – in sharp contrast to the wider public-key infrastructure – are and should be fair game for the most aggressive intelligence and law-enforcement techniques, as well as for invasive academic research. Indeed, having such clearly cordoned-off, free-fire zones is perhaps even useful for the state, because, conversely, a bad reputation leads to bad security. Either way, Tor's ugly example should loom

large in technology debates. Refusing to confront tough, inevitable political choices is simply irresponsible. The line between utopia and dystopia can be disturbingly thin.

Acknowledgements

The authors would like to thank Ian Brown, Ben Buchanan, Alec Muffett and David Omand for their comments on earlier drafts of this article.

Notes

- 1 Barton Gellman and Ashkan Soltani, 'NSA Infiltrates Links to Yahoo, Google Data Centers Worldwide, Snowden Documents Say', *Washington Post*, 30 October 2013, https://www.washingtonpost.com/world/national-security/nsa-infiltrates-links-to-yahoo-google-data-centers-worldwide-snowden-documents-say/2013/10/30/e51d661e-4166-11e3-8b74-d89d714ca4dd_story.html.
- 2 Joseph Menn, 'Secret Contract Tied NSA and Security Industry Pioneer', *Reuters*, 20 December 2013, <http://www.reuters.com/article/us-usa-security-rsa-idUSBRE9BJ1C220131221>.
- 3 Bruce Schneier, 'NSA Spying: Whom Do You Believe?', *Schneier on Security*, 23 December 2013, https://www.schneier.com/blog/archives/2013/12/nsa_spying_who.html.
- 4 Kevin Kelly, 'Cypherpunks, E-Money and the Technologies of Disconnection', *Whole Earth Review*, no. 79, Summer 1993, p. 48.
- 5 Intelligence and Security Committee, 'Privacy and Security', UK Parliament, HC 1075, 12 March 2015, p. 67.
- 6 Alex Hern, 'How Has David Cameron Caused a Storm Over Encryption?', *Guardian*, 15 January 2015, <http://www.theguardian.com/technology/2015/jan/15/david-cameron-encryption-anti-terror-laws>.
- 7 Pierre Thomas, 'Feds Challenged by Encrypted Devices of San Bernardino Attackers', *ABC News*, 9 December 2015, <http://abcnews.go.com/US/feds-challenged-encrypted-devices-san-bernardino-attackers/story?id=35680875>.
- 8 Terminology reflects differences in culture: privacy enthusiasts and activists often refer to 'crypto' whereas intelligence agencies and law enforcement speak of 'crypt'.
- 9 Our baseline for 'legitimate behaviour' is simply what is considered legal in the most liberal democratic jurisdiction on a specific issue. For free speech, for instance, this would be the United States; for recreational drugs, the Netherlands; for LGBT rights, probably Denmark.
- 10 James Ellis, 'The Story of Non-Secret Encryption', *Communications-Electronics Security Group*, 1987, available at <https://cryptome.org/jya/>

- ellisdoc.htm.
- 11 Whitfield Diffie and Martin Hellman, 'New Directions in Cryptography', *IEEE Transactions on Information Theory*, vol. 22, no. 6, November 1976, p. 644.
 - 12 *Ibid.*, p. 644.
 - 13 *Ibid.*, p. 648.
 - 14 Ronald L. Rivest, Adi Shamir and Leonard Adleman, 'A Method for Obtaining Digital Signatures and Public-Key Cryptosystems', *Communications of the ACM*, vol. 21, no. 2, 1978, p. 120.
 - 15 For the reasons why, see Stephen Levy, 'The Open Secret', *Wired*, vol. 7, no. 4, April 1999, pp. 1–6.
 - 16 Martin Gardner, 'A New Kind of Cipher That Would Take Millions of Years to Break', *Scientific American*, vol. 237, August 1977, pp. 120–4.
 - 17 See Stephen Levy, *Crypto* (New York: Penguin, 2000).
 - 18 For a detailed account, see Thomas Rid, *Rise of the Machines: A Cybernetic History* (New York: W.W. Norton, 2016), p. 255.
 - 19 David L. Chaum, 'Untraceable Electronic Mail, Return Addresses, and Digital Pseudonyms', *Communications of the ACM*, vol. 24, no. 2, February 1981, p. 85.
 - 20 David L. Chaum, 'Security Without Identification: Transaction Systems to Make Big Brother Obsolete', *Communications of the ACM*, vol. 28, no. 10, October 1985, p. 1,030.
 - 21 David L. Chaum, 'World's First Electronic Cash Payment Over Computer Networks', Digicash, press release, 26 May 1994.
 - 22 Judith Aldridge and David Décarv-Héty, 'Not an "Ebay for Drugs": The Cryptomarket "Silk Road" as a Paradigm Shifting Criminal Innovation', 13 May 2014, <http://dx.doi.org/10.2139/ssrn.2436643>.
 - 23 'Introduction to BlackNet', email to cypherpunks@toad.com, 18 August 1993.
 - 24 Peter Biddle, Paul England, Marcus Peinado and Bryan Willman, 'The Darknet and the Future of Content Distribution', ACM Workshop on Digital Rights Management, 18 November 2002, p. 54, <https://crypto.stanford.edu/DRM2002/prog.html>.
 - 25 It is immensely difficult to accurately gauge darknet sizes, as metrics vary greatly, including traffic, services offered, users, relays and others.
 - 26 Cath Everett, 'Moving Across to the Dark Side', *Network Security*, no. 9, 2009, pp. 10–12.
 - 27 Recent assessments by the Tor Project gauge hidden-services traffic to constitute 3–6% of the overall traffic in the Tor network. For a technical breakdown, see 'asn', 'Some Statistics about Onions', Tor Project, 26 February 2015, <https://blog.torproject.org/blog/33> for a technical breakdown.
 - 28 Roger Dingledine, Nick Mathewson and Paul Syverson, 'Tor: The Second-Generation Onion Router', in *Proceedings of the 13th USENIX Security Symposium*, 2004, <https://www.usenix.org/legacy/event/seco4/tech/dingledine.html>.
 - 29 Abdelberi Chaabane, Pere Manils and Mohamed Ali Kaafar, 'Digging into Anonymous Traffic: A Deep Analysis of the Tor Anonymizing Network', 2010 4th International Conference on Network and System Security (NSS), 1 September 2010, p. 167, <http://planete>.

- inrialpes.fr/papers/TorTraffic-NSS10.pdf.
- ³⁰ Dingledine et al., 'Tor: The Second-Generation Onion Router', p. 1.
- ³¹ A compromise successfully employed by the FBI to decloak users accessing illicit services in 2013. See Kevin Poulsen, 'FBI Admits It Controlled Tor Servers Behind Mass Malware Attack', *Wired*, 13 September 2013, <http://www.wired.com/2013/09/freedom-hosting-fbi/>.
- ³² The Tor Project's commitment to security is exemplified by their extensive release notes for the Tor Browser Bundle, and Tor's core. See <https://blog.torproject.org/category/tags/release>.
- ³³ Keith D. Watson, 'The Tor Network: A Global Inquiry into the Legal Status of Anonymity Networks', *Washington University Global Studies Law Review*, no. 11, 2012, pp. 718, 723.
- ³⁴ *Ibid.*, pp. 719–20.
- ³⁵ SecDev Foundation, 'Syrian Regime Tightens Access to Secure Online Communications', 27 October 2015, <http://new.secdev-foundation.org/syrian-regime-tightens-access-to-secure-online-communications>.
- ³⁶ See Gareth Owen, 'Tor: Hidden Services and Deanonymisation', presentation at Chaos Computer Club Conference 2004, <https://www.youtube.com/watch?v=oTEoLB-ses>; Alex Biryukov, Ivan Pustogarov and R. Weinmann, 'Trawling for Tor Hidden Services: Detection, Measurement, Deanonymization', 2013 IEEE Symposium on Security and Privacy, 19 May 2013, pp. 80–94, <https://doi.org/10.1109/SP.2013.15>; Clement Guitton, 'A Review of the Available Content on Tor Hidden Services: The Case Against Further Development', *Computers in Human Behavior*, vol. 29, no. 6, 2013, pp. 2,805–15.
- ³⁷ Dingledine et al., 'Tor: The Second-Generation Onion Router', p. 9.
- ³⁸ 'Some Statistics About Onions', Tor Project, 26 February 2015, <https://blog.torproject.org/blog/some-statistics-about-onions>.
- ³⁹ Biryukov et al., 'Trawling for Tor Hidden Services', pp. 80–94.
- ⁴⁰ Alex Biryukov et al., 'Content and Popularity Analysis of Tor Hidden Services', Distributed Computing Systems Workshop, 2014 IEEE 34th International, Madrid, 30 June 2014, pp. 188–93, <https://doi.org/10.1109/ICDCSW.2014.20>.
- ⁴¹ Guitton, 'A Review of the Available Content on Tor Hidden Services'.
- ⁴² See Biryukov et al., 'Content and Popularity Analysis of Tor Hidden Services'; and the hidden-service repository Ahmia.fi, available at <http://ahmia.fi>, which registers 5,197 active hidden services as of 10 January 2016. These numbers approximately match our own.
- ⁴³ Our findings are comparable in scale to those using the previous approaches mentioned.
- ⁴⁴ See, for example, Biryukov et al., 'Content and Popularity Analysis of Tor Hidden Services'.
- ⁴⁵ DDoS stands for Distributed Denial of Service, a technique in which the attacker floods the victim service with more traffic than it can handle, thereby preventing legitimate users from accessing it.
- ⁴⁶ See, for example, a recent ISIS Tor guide, 'Full and Detailed Archive of Information Security and Browsing

- Security and the Use of the Tor Network Provided to Supporters of #Ath_Alasalamah', 4 October 2015, <https://archive.org/details/ISIL-tor-guide>.
- 47 From the home page of a small vendor website named 'Unique Notes'.
- 48 In a recent example, the shuttering of the Darkode underground market and subsequent arrests have led to the site's resurfacing as a hidden service. See Kim Zetter, 'Dozens Nabbed in Takedown of Cybercrime Forum Darkode', *Wired*, 15 July 2015, <http://www.wired.com/2015/07/dozens-nabbed-takedown-cyber-crime-forum-darkode/>; and Liat Clark, 'Hacker Forum Darkode is Back and Safer than Ever', *Wired*, 28 July 2015, <http://www.wired.co.uk/news/archive/2015-07/28/darkode-back-and-more-secure>.
- 49 Often called 'Bitcoin mixers', these services generate a stream of transactions to turn investigations into a highly complex procedure.
- 50 Website named 'CleanCoin – Low-Fee Bitcoin Mixing/Tumbling/Laundry Service', URL withheld.
- 51 Website named 'Dark Tor Drugs', URL withheld.
- 52 Website named 'paste.lolz', a darknet-based clone of the popular textual posting site Pastebin (<http://pastebin.com>), URL withheld.
- 53 The front page of a website named 'Meet My Sisters', URL withheld.
- 54 Website named 'Pedo List', discussion named '13 Year-Olds Have Sex', URL withheld.
- 55 Website named 'Hitman Network – Hire Real Killers With Bitcoin, the Only True Hitman Site on the Deep Web', URL withheld.
- 56 Michael G. Reed, 'Re: Onion Routing', email to Wei Dai, 4 March 1997, <http://www.onion-router.net/Archives/onions-1997.txt>.
- 57 Roger Dingledine, Nick Mathewson and Paul Syverson, 'Deploying Low-Latency Anonymity: Design Challenges and Social Factors', *IEEE Security & Privacy*, vol. 5, no. 5, September/October 2007, pp. 83–7.
- 58 Timothy May, 'Re: Encryption and the 2nd Amendment', email to cypherpunks@toad.com, 20 January 1996.
- 59 Timothy May, 'Re: Blacknet Worries', email to cypherpunks@toad.com, 20 February 1994.
- 60 The line was commonly repeated among cypherpunks. See, for example, Loren James Rittle, 'The Clipper Chip Proposal', email to the Vice President of the United States, copied to cypherpunks@toad.com, 21 July 1994. See also Stephen Levy, 'Crypto Rebels', *Wired*, May/June 1993, pp. 54–61.
- 61 'We advocate encryption', said Robert Hannigan, the director of the spy agency, in November 2015. See Alexander J. Martin, 'GCHQ Director Blasts Free Market, Says UK Must Be a "Sovereign Cryptographic Nation"', *Register*, 15 November 2015, http://www.theregister.co.uk/2015/11/10/gchq_director_speech/.
- 62 This is a counterproductive term. The controversial draft UK Investigatory Powers Bill, for example, does not call for back doors; instead, the government may demand that communication-service providers remove end-to-end encryption from specific user

- accounts if presented with a warrant. Maintaining such a capability is not technically a back door.
- ⁶³ Nicole Perlroth and David Sanger, 'Obama Won't Seek Access to Encrypted User Data', *New York Times*, 11 October 2015, p. A24. Even before the White House's decision, the Home Office had decided that the Investigatory Powers Bill would probably not change the legislative status quo on encryption in the UK. Thomas Rid, conversation with senior Home Office officials, 9 October 2015.
- ⁶⁴ See, for example, Lasse Øverlier and Paul Syverson, 'Locating Hidden Servers', 2006 IEEE Symposium on Security and Privacy, May 2006, pp. 100–14, <https://doi.org/10.1109/SP.2006.24>; and Lasse Øverlier and Paul Syverson, 'Valet Services: Improving Hidden Servers with a Personal Touch', *Privacy Enhancing Technologies*, 1 January 2006, pp. 223–44.
- ⁶⁵ Ross Anderson, 'The Eternity Service', *Proceedings of PRAGOCRYPT*, October 1996, pp. 242–52.
- ⁶⁶ Michael G. Reed, Paul F. Syverson and David M. Goldschlag, 'Anonymous Connections and Onion Routing', *IEEE Journal on Selected Areas in Communications*, vol. 16, no. 4, 1998, pp. 482–94, <https://doi.org/10.1109/49.668972>. See also David M. Goldschlag, Michael G. Reed and Paul F. Syverson, 'Hiding Routing Information', *Information Hiding*, 1996, pp. 137–50; and Dingleline et al., 'Tor: The Second-Generation Onion Router'.
- ⁶⁷ Paul Syverson and Griffin Boyce, 'Genuine Onion: Simple, Fast, Flexible, and Cheap Website Authentication', IEEE Workshop on Web 2.0 Security and Privacy, May 2015, <http://www.nrl.navy.mil/itd/chacs/sites/www.nrl.navy.mil/itd/chacs/files/pdfs/15-1231-0478.pdf>.
- ⁶⁸ For an overview, see 'The Official SecureDrop Directory', <http://secdrop5wyphb5x.onion/directory>. DuckDuckGo is at <http://3g2upl4pq6kufc4m.onion/>; ProPublica, the Darknet's first major news site, is at <http://www.propub3r6espa33w.onion/>. The first news site was possibly Kavkaz Center, a Chechen jihadi site launched in 2013.
- ⁶⁹ Alec Muffett, conversation with authors, 1 April 2015.
- ⁷⁰ For an introduction, see Kevin Poulsen, 'Strongbox and Aaron Swartz', *New Yorker*, 14 May 2013.
- ⁷¹ See 'Tor: 80 Percent of ??? Percent of 1–2 Percent Abusive', Tor Project, 31 December 2014, especially Roger Dingleline's posts, as 'arma', <https://blog.torproject.org/blog/tor-80-percent-percent-1-2-percent-abusive>.
- ⁷² Syverson and Boyce, 'Genuine Onion'.
- ⁷³ 'Hidden Services Need Some Love', Tor Project, 22 April 2013, <https://blog.torproject.org/blog/hidden-services-need-some-love>.
- ⁷⁴ Quoted in Glenn Greenwald, *No Place to Hide* (New York: Macmillan, 2014), p. 24.
- ⁷⁵ Thomas Jefferson, 'Jefferson's Draft', in *The Papers of Thomas Jefferson, Volume 30: 1 January 1798 to 31 January 1799* (Princeton, NJ: Princeton University Press, 2003), pp. 536–43, available at <https://jeffersonpapers.princeton.edu/selected-documents/jefferson%E2%80%99s-draft>.
- ⁷⁶ Julian Assange, *Cyberpunks* (New

- York: OR, 2012), p. 4.
- ⁷⁷ Joseph Cox, 'ISIS Now Has a Propaganda Site on the Dark Web', *Motherboard*, 16 November 2015, <http://motherboard.vice.com/read/isis-now-has-a-propaganda-site-on-the-dark-web>.
- ⁷⁸ For details, see 'Learning More About the GFW's Active Probing System', Tor Project, 14 September 2015.
- ⁷⁹ Kari Paul, 'Russia Wants to Block Tor, but it Probably Can't', *Motherboard*, 18 February 2015, <http://motherboard.vice.com/read/russia-wants-to-block-tor-but-it-probably-cant>. See also Oren Dotan, 'Russia Moves to Ban Tor and Anonymous Web Surfing', *Vocativ*, 10 February 2015, <http://www.vocativ.com/tech/internet/dark-net-russia/>.
- ⁸⁰ For an introduction, see John Bew, *Realpolitik: A History* (New York: Oxford University Press, 2016).
- ⁸¹ See 'The Trust Machine', *The Economist*, 31 October 2015, <http://www.economist.com/news/leaders/21677198-technology-behind-bitcoin-could-transform-how-economy-works-trust-machine>.
- ⁸² One caveat that does not affect our argument: some commercial products, such as Symantec Encryption Desktop, use open-source code; and some non-profit crypto systems are not open-source, such as Wickr, an end-to-end encrypted app. Some large communication-service providers also argue that their system architecture may not allow them to remove encryption. Author conversations with three global end-to-end encryption providers, November–December 2015. For recent precedent, see Vindu Goel and Vinod Sreeharsha, 'Brazil Restores WhatsApp Service After Brief Ban', *New York Times*, 18 December 2015, p. A10.
- ⁸³ Perlroth and Sanger, 'Obama Won't Seek Access to Encrypted User Data'.
- ⁸⁴ See Ciaran Martin [director-general for cyber security at GCHQ], witness statement to Investigatory Powers Tribunal, IPT/14/85/CH, 16 November 2015, pp. 4–8.
- ⁸⁵ See Dingleline, Mathewson and Syverson, 'Deploying Low-Latency Anonymity'.