Infrastructural power: dealing with abuse, crime, and control in the Tor anonymity network

Abstract

This chapter reports on the first empirical criminological research on the Tor Project, the organisation which develops the Tor anonymity network. There has been little focus as yet by cybercrime researchers on the human factors shaping the platforms and infrastructures on which the Internet depends. These are emerging as powerful technologies of control and profound sites of resistance in contemporary societies, increasingly taking on responsibility for enormous user communities and the crime and abuse which comes with them. Of these, I focus on Tor, an international anonymity infrastructure which offers its users extremely strong protections against online surveillance and censorship. Tor has become a particularly important subject of criminological research on online crime. However, there is as yet no criminological research which deals with how the people who develop and maintain Tor understand these issues. Through interviews and archival research, I study how this community perceive Tor's use for crime and harm and how they navigate these issues in practice, identifying three distinct sites at which Tor deals with crime, and three concomitant ways of making sense of Tor's crime problem (conceptualised as 'social worlds' of Tor). I explore how Tor has developed from a disruptive character to an increasingly governmental one, and the implications of this for understanding the role of platforms and infrastructures in the governance of online crime more broadly.

Keywords

CYBERCRIME, SOCIAL WORLDS, DARKNET, TOR, INFRASTRUCTURE, PLATFORMS, ABUSE, PRIVACY

Introduction: power, crime, and control online

Within criminological research, the infrastructures of the Internet often fade into the background to become part of the context of social action rather than dynamic social subjects of their own. In fact, the decisions which go into the creation of these technologies and the values of the people who make and support them shape in important ways how the Internet's infrastructures and platforms become implicated in crime, power and control (Pinch 2010). In this chapter, I draw on interviews and archival research with the community which develops and maintains Tor, an online anonymity infrastructure, to explore how they understand the use of their network for crime. I begin by setting out issues of platform governance more generally, moving to an overview of Tor's history, how it works, and some of the problems it faces in practice. Next, I set out my use of social worlds theory to unpick the complex values of the Tor community around crime, abuse, and control (Star, 1989) and the research methods through which I explored these values. Moving on to the results of this research, I characterise individually three separate worlds of discourse in Tor and how each frames and tackles abuse, crime, and harm in practice. Finally, I discuss the implications of my findings for the study of Internet infrastructures in criminological research.

Context and review of the literature

Platforms, privacy, and abuse

The platform model of the modern Internet, through which online services are increasingly administered by a small number of large international corporations like Facebook and Google, poses challenges for traditional modes of governance. These Internet services developed by private companies and non-profits are increasingly moving from models of disruptive innovation to exercising governmental power of their own over their colossal userbases, which can number in the billions (Gillespie 2018; West, 2017; Zuboff 2015;). The continuing problem of online crime puts these providers in a rather contentious position. As they are increasingly called on to take responsibility for the harms facilitated by their platforms, so too have the main mechanisms available to them – which largely rely on technological surveillance and censorship – been the subject of substantial public backlash (Lyon, 2014; Nissenbaum, 2009). Balancing these is particularly challenging for the software engineers and business professionals who make these decisions, who by necessity approach them differently from police or civil servants (Gillespie, 2018; Sandvig, 2014). Although they lack the formal powers of government, they arguably collect more personal data about a wider range of individuals than any government in human history ever has (Lyon, 2014).

The privacy properties of Internet technologies have therefore become battlegrounds over freedom, control, and power, as the technologies of control to which we are subjected revolve around increasingly authoritarian forms of technological surveillance (Bauman and Lyon, 2013; Kohl, 2013; Lyon 2014). These mechanisms of state and corporate online power are resisted by a range of organisations and social movements which envision different futures for the Internet, collectively known as the Internet freedom movement. These groups engage in these struggles not only through traditional channels of lobbying, policy work and activism, but also try to "steal the fire" (Milan, 2013) by building their own infrastructures and tools which sit on top of the Internet backbone.

Tor – the Dark Web as a privacy infrastructure

In exploring these issues, I focus on the Tor Project, an Internet infrastructure which takes a particularly pure approach to balancing between control and privacy in dealing with online harms. Tor is a network explicitly designed to preserve the privacy of its users above all other concerns, and it drastically restricts its own ability to control what it is used for (being designed to facilitate as wide a range of users and use cases as possible). Built 'on top of' the Internet, Tor uses engineering approaches to restrict the ability of states, Internet Service Providers and others to control, surveil, or censor their users' Internet traffic. As a result, it has also become particularly associated with crime and abuse in public discourse, and hence is a particularly apposite case study for exploring these issues.

The Tor network, often referred to as the "Darknet" or "Dark Web", is an international, volunteer-run infrastructure which 'sits on top of' the regular Internet, providing very strong security and privacy protections to its users. The Onion Routing paradigm which forms the heart of Tor was developed by the US Navy's Naval Research Laboratory as a means of communicating on insecure networks controlled by foreign governments. Onion Routing involves wrapping the routing information of users' Internet traffic in three layers of encryption, then bouncing it around a global network of volunteer-run servers, or 'relays', each of which decrypts a single layer of encryption to find the next relay in the chain. The final, 'exit', relay makes the connection to the site or service with which the user intends to communicate, so no relay (or observer) knows both the origin and destination of the traffic. This generates a crowd of users traversing the Tor network which are hard to distinguish from one another; the larger and more diverse this crowd, the better the anonymity protections for anyone using the network (Dingledine et al., 2004). End users access the Tor network using the free-to-download Tor Browser, which can be obtained at www.torproject.org.

By developing a worldwide, volunteer-run anonymity network used by a relatively wide proportion of the general public, the US naval researchers sought to develop a large crowd of innocuous cover traffic in which to hide their transmissions (Moore and Rid, 2016). In addition to browsing, the Tor network can also be used to host Onion Services: web services which are extremely hard to surveil or to shut down, and which can only be accessed through Tor. Due to the strong privacy and security protections it provides, and its designers' aim to broaden its user base beyond military and high security users, Tor and the Tor community have moved from their initial roots in military and security research to become a key part of the Internet freedom movement. Responsibility for the development of Tor is now in the hands of the Tor Project, a non-profit organisation which develops it as an international security and privacy technology. This organisation includes a variety of roles, including developers who work on the design and source code of Tor and its network, activists and advocates who conduct outreach to Tor's users community, HR and PR professionals who manage the organisation and engage with the media, and others. Tor's vibrant community of volunteer relay operators run the network, separately from the Tor Project organisation. Tor is now used by two million users per day around the world, including activists, journalists, law enforcement, privacy-conscious everyday Internet users, and whistleblowers.

This has not been without controversy. As its user community has grown, Tor has faced many of the same problems as other 'disruptive' Internet platforms, such as Twitter, Google, or Facebook. In particular, it has faced people using its platform for harmful or illegal purposes (Moore and Rid, 2016). Tor's capacity to host Onion Services has led to the rise of so-called cryptomarkets, anonymous online markets for illicit goods hosted on the Tor network which are very difficult for authorities to take down (Aldridge and Decary-Hetu, 2016; Barratt et al., 2016). These markets have caused substantial negative press for Tor and are often collectively referred to as the Dark Web, cementing a public association with drug

dealing, terrorism, and child sexual abuse images which has proven deeply harmful for Tor's reputation and public image.

Internet infrastructures and strategies of control

Most Internet infrastructures have three avenues through which to deal with crime, harm, and abuse. The first of these is the route of design, making changes to the technology in order to shape how it can be used. This is effectively an 'online' Situational Crime Prevention approach (Cornish and Clarke, 2003), making changes to the built environment in order to alter opportunities for criminal offending, or to increase possibilities for guardianship (Reynes, 2010). There are many examples of this, including automated detection systems which scan messages on social media platforms for hate speech or child abuse images and remove them, systems for collecting information on users' real identities, and more subtle changes which can be made to the user experience to 'nudge' people away from abusive or illegal behaviour (see for example, Suzor et al., 2019).

The second approach involves moderation and administration; a range of techniques through which platforms directly police user behaviour. Many platforms make use of moderators and administrators to handle abuse reports, make decisions about suitable sanctions, such as restrictions on use or banning users from the platform, and some adopt a more community-based approach, with moderation of norms and conduct left up to particularly well-established users (Suzor et al. 2019, West, 2017). These processes effectively set up internal policing and governance mechanisms and systems of sanctions for the users of the site through which behaviour is observed by automated systems, paid administrators, or community members, and unwanted behaviour sanctioned.

Finally, platforms can engage with the formal institutions of law enforcement and criminal justice. This involves storing and collecting user data which can be used as evidence in

investigations, the establishment of reporting mechanisms where illegal behaviour is detected, and either replying to subpoenas for user data or, on occasion, developing more active collaboration regimes with secret services (Lyon, 2014). As revealed in the Snowden leaks, and as has become increasingly prominent in discussions about the operation of contemporary criminal justice systems, this kind of collaboration has only been deepening, with some exceptions where companies, such as Apple, have tried to assert the rights of their users against state intrusion (Schulze, 2017).

Navigating crime and power as a rebel infrastructure

Tor, however, has deliberately limited its ability to engage in any of these strategies. Its design intentionally removes any of the control points through which user behaviour might be surveilled, and its foundational design decisions, based around maximising usability, all seek to design out *control* rather than designing out crime. This is both as a matter of principle, and to prevent the people who run its infrastructure and design its code becoming targets themselves. By extension, it has also 'designed out' its ability to administer or moderate user behaviour, and this, along with the anti-authoritarian sensibilities of its community, makes collaboration with law enforcement both a technical impossibility (as the infrastructure collects no useful data on its users) and opposed as a matter of principle. This makes it a particularly interesting case to study in terms of abuse regulation, bringing to the fore the underlying strategies, tensions, and rationalities which sit behind these three more conventional approaches.

Theory and methodology: A social worlds approach to studying Internet infrastructure

There is a wide literature within criminology which deals with the management and governance of crime and harm, particularly within *governmentality studies*, a branch of criminological scholarship which applies the social theory of Michel Foucault to studying the

business of governing societies (Garland, 1997, Valverde, 2009). Foucault argued that government in contemporary societies does not operate from above, but rather, that governing power is spread throughout society amid a range of different institutions, infrastructures, and bodies of multiple kinds (Foucault, 1991). A key contention of this body of research is that in order to understand *technologies of control* – the different ways in which societies govern conduct – one must also understand *rationalities of power* – the ideas and theories about how societies should work, expressed as particular visions of the world – which underpin them (Foucault, 1991; Garland, 1997). The technologies, infrastructures and institutions which result are embedded with these discourses in the form of category systems, such as types of people, types of behaviour, or types of relationship.

Understanding power as operating through *discourses* which are stabilised in a range of material practices and forms is particularly useful for making sense of Internet infrastructures, their role in governing and sanctioning behaviour, how they interact with the broader landscape of power. However, unpicking the ideas and visions which underpin Internet technologies is particularly challenging, as the dense technical details of how these technologies work is hard for social scientists to make sense of. Communications infrastructures are built and maintained by a range of different people, who may understand their job and the role of the technology in rather different ways, so the vision of the world expressed in the end result is often multiple and changing.

Underneath the wires, servers, and infrastructures of Tor lies a substantial cultural life. Although it is tempting to try to unearth a singular perspective which encapsulates Tor's values, its community doesn't speak with a single voice. As an infrastructure, the Tor network relies on a wide variety of different working practices, which involve different relationships with Tor as a technological project and different ways of making sense of how it affects society. This lability to different meanings and uses is a core characteristic of

infrastructure (Star and Griesemer, 1999): the multiple groups they rely on and wide array of different uses they support mean that they accrete dense, heterogeneous, and changing thickets of meaning which are hard to pull apart and separate from the detail of technical practices. Individuals can themselves draw on a range of different understandings of a single infrastructure in carrying out different kinds of work (Star 1999).

In unpicking the *rationalities* and *discourses* which underpin Tor's attempts to deal with these issues of crime and abuse, I use social worlds theory, a theoretical approach which draws on an interactionist conception of technological and scientific work, allowing researchers to inductively pull out the complex landscape of discourses which form around technologies into internally-consistent and coherent 'social worlds' of discourse (Clarke and Star, 2008). These are stabilised in different ways: through group interaction, through the development of working practices, through the creation of documents or policy statements, or through design elements of infrastructures themselves. Rather than focusing on individuals, artefacts, or groups of people within these communities, social worlds theory focuses on discourse and practices, allowing individuals to draw on several social worlds at the same time and focusing on how these discourses and worlds go on to shape and become embedded in material aspects of these technologies, such as design or maintenance (Star and Griesemer, 1989). Understanding how Tor reacts to and makes sense of abuse therefore requires unpicking the social worlds within its community through empirical research.

Research Methods

I conducted online and in-person semi-structured interviews with twenty-six members of the Tor community, including developers, activists, relay operators and others. Given Tor's practices of 'radical transparency', the majority of its core developers and contributors are listed, along with contact details, on the Tor Project website. I approached sixty-two Tor

community members via email (thirty-two responded, of which six refused), with some interviews resulting from connections made at international computer security and Internet freedom activist conferences.

My sample of interviewees was broadly reflective of the diversity of the Tor community, based on the information available on the Tor Project people page. This included nine developers (from fairly new members of the Tor team to some who had been involved since its early days), three other contributors to the Tor Project, eight relay operators, three Onion Services developers, and three other members of the Tor community. My participants were based in a range of countries, including Australia, Canada, France, Germany, Greece, Italy, Russia, Spain, the UK, and the USA. This is fairly representative of the core Tor community, though under-represents members from the global South. These interviews were supplemented by extensive archival research in Tor's freely-accessible online archives, where it stores over sixteen years' worth of internal mailing list discussions and design documents.

I teased out the different discourses from my interviews to arrive at three distinct social worlds of Tor (Collier, 2020) which favour contrasting strategies for navigating the challenges Tor faces. These social worlds are ideal types, not necessarily completely encapsulating the viewpoint of any individual actor. Although some of the participants appeared to be anchored in a single social world, others subscribed to views spanning two or more such worlds. I draw from these interviews and documents to characterise the different worlds of discourse in the Tor community, the way they frame Tor's implication in crime and harm, and the material strategies and practices through which they attempt to navigate the pressing problems of crime and criminal justice which Tor today faces. Although these different worlds are associated with different kinds of work, attain prominence at different periods, and change and influence one another over time (Collier, 2020), they also coexist throughout Tor's history and I have aimed to reflect this by situating them in their historical

context. In the following findings sections, I discuss in turn each of these three social worlds and how they encounter and conceptualise abuse – the engineer world, the infrastructuralist world, and the activist world - before discussing some broader changes to Tor's cultures in recent years.

Findings

Privacy as a structure: the *engineer* world and standardisation

Beginning with a small team of Naval Research Laboratory researchers in the late '90s, and expanding to include privacy-conscious researchers and academics in the wider information security community across the early '00s, Tor's early development centred around implementing and designing a network which could mitigate some of the privacy and security issues posed by the way the Internet works, using the 'Onion Routing' framework I describe above. A particular way of understanding Tor and privacy technology emerged from this early development work which underpinned much of Tor's initial mission – I term this the *engineer* social world of Tor.

This engineer world views Tor as a direct actor in power relations, with power and privacy arising from structural forms coded into the topology of the technical networks of the Internet which Tor aims to remake and redesign. They see the design of the Internet as concentrating power in 'choke points' in these systems (in particular, the traceability of administrative information such as IP addresses which are visible to internet service providers) which can be used by nation states as a mechanism of surveillance and control. Tor's design aims to flatten this terrain of structural power (whether in the service of everyday privacy, or to help the US military communicate in countries where they don't control these choke points) by separating this administrative information seen by the ISP and other platforms from the actual identity of Internet users. When its attempts to reshape this landscape of power run into challenges in

practice, Tor's engineer world has its own distinct understanding of these issues of crime, harm, and control, and its own strategies for dealing with them.

> It's kind of a bit like MP3, where you say, OK, society might not be ready yet and we will kill a lot of stuff and, and… video killed the radio star! And it's like, technology comes first and then there's a struggle in society on how to restructure itself to be able to cope with that change… All these structures are becoming more and more stale and static and the only way to change them would be to break them. And I like fluid systems. I like this structurelessness and chaos, and I think that's a value by itself.
>
> *Participant L – Tor core contributor*

At the heart of many of Tor's design decisions is the desire to remove as much as possible any ability to control or censor traffic from the Tor network itself. This, therefore, dramatically reduces the ability of the Tor Project to prevent its abuse for crime. From the engineer perspective, conversations about the crime, deviance, and harm with which Tor is associated are a red herring. Their understanding of crime mirrors that of critical criminologists such as Box (2002), arguing that 'crime' is in fact constructed and enforced by and in the interests of the powerful, designed to distract the public from real questions of power in society. They see crime and harm as an unfortunate but unavoidable consequence of disrupting these vested power interests: rather than promoting positive or negative use cases, Tor works in the interests of those without power over those with power. Equally, they argue that some uses which may be illegal in particular countries may, in fact, be key Tor use cases: for example, LGBT rights activism.

Engineer discourse is not as anti-policing as might be expected, and in fact many expressing this perspective were in favour of the use of targeted police powers to tackle crime. What it opposes is the adoption of engineering and architectural solutions for social control. They

argue that policing through automated mass online surveillance is a dangerous and authoritarian centralisation of power to the state and the unelected software developers who build these platforms, and that social issues should be tackled through democratically accountable institutions.

This does not mean that the engineers do not recognise that crime poses a problem for Tor. They see Tor overcoming these problems through *standardisation*, growing Tor beyond a single technology to become instead a fundamental part of how the Internet works (in much the same way that encryption is). This involves trying to get Tor 'built in' to other technologies, a toolkit for developers rather than only a tool for users. This has the benefit of reframing Tor's crime problems as consequences of a broader shift in how the Internet works, rather than the result of an upstart activist technology. Tor was designed with this inter-operability in mind from the beginning, much like the Internet itself. Many of Tor's core design decisions (such as allowing it to browse the regular internet) are aimed at enabling these interfaces with other technologies, and there is a substantial degree of work undertaken by the Tor Project in convincing other developers to make use of Tor in their own platforms.

> I see Tor and Onion space right now roughly where… web encryption was around, like 2001 or so. Back then if you set up encryption for a webpage, people said… what are you trying to hide, what kind of criminal thing do you have going on? And now it's recognised as the fundamental enabler of e-commerce... ideally [in the long term], I'm out of a job, or doing something else, because this is [now] just the way the Internet works. *Participant I – Tor core developer*

Particularly successful examples of this are the Onion Toolkit, developed by Alec Muffet, which allows anyone to easily set up an Onion Service. Tor has been built into chat messaging apps such as Ricochet and Cwytch (https://cwtch.im). The whistleblowing

platform SecureDrop is another example of an Onion Service technology, which has been widely used by news organisations to take anonymous submissions. Tor has become a go-to tool for security researchers who research adversarial websites as it allows them to collect information without being blocked or revealing their location.

> They're not all these, these drugs undergrounds. Like, the majority of them are these ephemeral things that are just in the background. And I think we're going to start seeing a lot more of them as Tor is sort of built into things in ways where you don't even know it's there…where Tor is more of a security toolbox, where you can pick and choose which features you want… this is what's needed to get Tor into everything as the underlying technology for communication. *Participant D - Tor core developer*

Most importantly, Tor has also begun to try to get incorporated into other browsers, with Brave Browser recently integrating Tor into its private browsing mode, so that its users can access Tor in their browser with the click of a button. The much more widely used browser Firefox (which has 250 million users) has long been considering a similar move, in the meantime incorporating a range of Tor's security improvements and anti-tracking technologies. As Google increasingly becomes known for its surveillance operations, competitors to its Chrome browser are increasingly using privacy and anonymity protections as a distinguishing feature.

Much as strong encryption became the norm for online technologies (despite much resistance from the US government) not only due to the tireless campaigning work of activists, but also to the enormous security benefits which this offered to online banking and commerce, the increasing preponderance of high-profile cyberattacks and breaches could well lead to the protections which Onion Services offer being increasingly in demand, leading to the

increasing *standardisation* of Tor. Despite Tor's emphasis on privacy over control, issues of abuse are important considerations in its design and development processes. In pursuing standardisation, the Tor Project seeks to frame its action in relation to broader dynamics of online power and control rather than become caught up in the complexities of what constitutes crime in different jurisdictions. However, as will become apparent, design is not the only site at which these issues raise their head for the Tor Project.

Privacy as a service: the *infrastructuralist* world and neutralisation

As the Tor network began to expand from a prototype to a fully-functioning infrastructure, the maintenance and administration of its network took on an increasingly central role. The growing community of relay operators, the volunteers from around the world who run Tor's relay network, began to experience the consequences of Tor's attempts to restructure online power and privacy first-hand. As the first abuse complaints began to trickle in, largely in the form of copyright enforcement notices and then from services such as Wikipedia, the relay operators began to take on a substantial amount of work: responding to these notifications, dealing with ISPs who began to become reluctant to house Tor relays, and with blocklists and law enforcement. From these practices of maintenance and administration arose a distinct perspective from that of the engineers – the *infrastructuralist perspective*.

This can be summarised as an ethic of 'privacy as a service', with Tor functioning as a neutral service provider, protecting its users' privacy online without claiming to act as a political actor in its own right or take a view on what people use it for. The *infrastructuralist* world aims to denude Tor of explicit values, withdrawing it from public conversations about politics and social meaning in order to permit as wide a range of people and perspectives to contribute, regardless of conflicting understandings of 'privacy' and its importance. In accordance with this perspective, for much of Tor's life, the Tor Project has avoided making

strong public commitments to a particular set of values other than a dedication to privacy, generally preferring to frame the technology itself as 'neutral' in order to avoid contentious political debates and permit the widest possible community of contributors and users.

The infrastructuralist perspective baulks at the assertion that Tor as a community or organisation should take any view at all on the particular ways in which it is used, even if these are abusive or criminal. The majority of the relay operators whom I interviewed felt this way, often comparing Tor with a knife or similar tool, with no intrinsic politics or values. This amounts to an assertion of 'technological neutrality': the argument that technologies themselves possess no agency and are mere conduits of human action.

> It's like, *sighs* it's like having a knife – with a knife you can cut an apple and with a knife you can kill a man… so the Tor network is just a knife which is laying on the table without anyone touching it. That's my opinion. *Participant Q - Tor relay operator*

> Because the tool is something that does, something that helps you to do something. But what you will do with this tool is up to you. Crime happens not on the hard drive of the Bond movie producer, crime happens not on the Silk Road drug store, no. Crime happens inside people's mind. The criminal mind is a way of thinking… Neither Tor or other software authors, nor people who are running even exit nodes, no they're not responsible… for another people's thoughts and actions… Tor is just a tool. *Participant N - Tor relay operator and open source contributor*

In terms of managing the abuse itself, Tor substantially limits the ability of its operators to censor how it is used. While individual operators have the ability to blocklist particular websites and *kinds* of web traffic (for example, email, or Internet Relay Chat) from passing

through their relay, they have no effect on the network as a whole and cannot censor based on the content of communications, so the network takes a 'neutral' approach to how it is used.

In practice, any censorship of Tor traffic is taboo for operators. The fact that Tor is designed specifically to preclude any mechanisms for censorship on the basis of content allows Tor relay operators to take advantage of laws which offer 'mere conduit' protections, absolving them as service providers from responsibility or liability for the actions of their users. This legal and moral neutrality is very important for the people who run the Tor network, given the content which flows through their servers. The reality of Tor is that, as well as providing substantial social benefits, it also facilitates (as does any infrastructure) a range of activities which are unambiguously harmful. Although the relay community justifiably defends their decision to help Tor, they do need ways to reconcile this tension and the stigma it brings. This makes taking a view on user traffic of any kind dangerous, with infrastructuralist discourse preferring to recuse itself from moral judgement and avoid articulating Tor as 'for' any particular use case other than a broadly-conceived privacy.

> Under European law, I am not allowed to alter the packetflow. As long as I am pushing packets from A to B I am protected as a ISP. Would I like to kick the botnets out? Yes! Am I allowed to do this? I don't think so. *Participant P - Relay operator*

Adopting this way of understanding abuse provides them with a way of coming to terms with the reality of how some people use Tor in practice. The more diverse groups which use Tor, even including the police, or those who use it for harmful purposes, the more its relay operators feel they can abrogate responsibility for the traffic which flows through and maintain their 'neutral' status: as soon as they begin to take a moral view on this traffic or try to shape how Tor is used, they risk becoming culpable.

This neutral framing has a further practical purpose. While Tor is not explicitly criminalised in many countries, it does become entangled in criminal justice processes, which brings it into conflict with the technologies of control through which states maintain online order. This is primarily experienced by Tor's relay operators, as police and ISPs investigating illegal online conduct by Tor users follow their digital trails to the door of the relay operator, who appears from these administrative records to have been responsible for significant illegal activities.

Although successful prosecutions are rare, especially as Tor provides a service which allows investigators to establish proof that the traffic originated from the operator's Tor relay, rather than their personal computer, operators understandably try to avoid getting caught in this process to begin with. As a result, Tor's relay operators have attempted to mitigate this threat through the cultivation of fairly sophisticated mechanisms to de-intersect Tor from these administrative processes while still protecting its users from state surveillance. Crucially, Tor's operators have realised that having a relay in their own name operating from a home internet connection appears very different to police than a relay owned by an Internet Service Provider hosted in a private datacentre:

> When you, as an ISP, interact with law enforcement, you're interacting with people who know what they're dealing with… Like, literally, the only contact I have had with them is on that kind of level, where somebody is doing something bad on the Internet, oh, it's a Tor node, oh, OK, we know what that is, we'll… go find evidence some other way! Right? *laughs* Um, and that's kind of the way it should be.
>
> *Participant U - Tor relay operator*

Accordingly, Tor's exit operators often set up small companies or charitable organisations which they register as a service provider and use to host their relay, which they refer to as

"legal entities". This means that when police look up an IP address associated with illegal activities, they find what appears to be a company providing hosting for its customers, rather than an individual's home connection. For the relay operator, this is the difference between a dawn raid for child pornography charges, including the seizure of computers and a potential court case, and a polite letter informing them that one of their customers has broken the law.

Prospective relay operators are advised to avoid running a Tor relay from their home connection, instead setting it up in a datacentre on a rented server. These servers tend to be clustered within countries and service providers who are sympathetic to Tor, don't bother to ban nodes over abuse complaints, or have jurisdictions where investigating foreign cybercrime cases is not a police priority.

> When I run an exit, I want it to be owned by a legal entity that's not me… when someone uses that exit for something bad, and some police investigation happens, which unfortunately might happen, I want it to go to the company that owns it, and then at least it'll mean that they'll ask a question before they bash my door down…. I want it to be obvious when a police investigation is happening that this is a proxy
>
> *Participant R - Tor relay operator*

Overall, therefore, there are a range of practical and cultural drives which arise from the administration and maintenance of the Tor network which push towards the neutralisation and de-politicisation of what are the (clearly still deeply political) values of privacy, anonymity, and anti-censorship at the heart of the network as a response to abuse of Tor. These attempts to paint the technology as a mere neutral carrier chime with the particular experience of the technology associated with infrastructuralist world, pushing them to 'make Tor invisible' rather than to become involved in complex conversations about its values in public. They do not stop the operators' experience of this work being deeply bound up with

politics and values, however they fulfil a practical purpose in facilitating its de-stigmatisation and slipping it between the cracks of the external systems of power against which it clashes.

Privacy as a struggle: the *activist* world and reclamation

This neutralisation of Tor's politics, however, has come into conflict with a third sensibility which has increasingly risen to prominence in the Tor community over recent years. The revelations of mass US surveillance of Internet traffic by Edward Snowden in 2013 both led to a massive influx of new, politically engaged people with activist and policy backgrounds to the Tor community, sparking a wider movement critical of the power and politics behind Internet infrastructures and platforms. In the context of broader attempts to professionalise the Tor Project as a fully-fledged NGO (Collier, 2020; Marechal, 2018), Tor's increasing infamy in the press, including the rise of cryptomarkets to public prominence and a wave of negative stories about terrorism and child sexual abuse images on Tor, contributed to a general sense that a change in tack was needed to reframe public perceptions of Tor. The resulting rise of what I term the *activist* social world has caused a reorientation of this strategy of technological neutrality, and Tor has in recent years become much more engaged in public discussions about the values it represents.

The *activist* world, which stems from the work of policy workers, activists, and managers in the Tor community, views privacy as a struggle. Privacy technologies are framed as part of a political movement for civil rights, wielding political power and embodying a coherent set of values of their own. Asserting these values in public is therefore, for this world, an important way in which privacy technologies exert power and shape societies. Accordingly, Tor has engaged far more in recent years in asserting its values, through more engagement with the news media, a co-ordinated communications and branding strategy, the professionalisation of the organisation, and through the publication of documents such as the Tor Social Contract:

> Tor is not just software, but a labor of love produced by an international community of people devoted to human rights… We advance human rights by creating and deploying usable anonymity and privacy technologies… Our vision of a more free society will not be accomplished simply behind a computer screen, and so in addition to writing good code, we also prioritize community outreach and advocacy. *Excerpts from the Tor Social Contract*

This activist social world frames problems of crime faced by technologies like Tor as stemming at least partly from questions of public image and perceived values. In this framing, technologies like Tor attract crime problems (and the attention of the criminal justice system) when they become associated with crime and deviance and legitimate users become dissuaded. Hence, the activist world contends that Tor's reputation as a 'Dark Web' full of illegal content is the prime factor in shaping its use for crime, and if it becomes known as a tool for free speech and liberal democracy it is likely to attract a wider range of more positive use cases.

Promoting Tor's socially beneficial use cases and encouraging more journalistic organisations to set up Onion Service versions of their websites is a large part of this effort at changing Tor's image. The activist social world is also the only one of Tor's social worlds which is occasionally (though not always) willing to condemn Onion Services (also known as Hidden Services) outright, arguing that they pose too great a risk of abuse, (as distinct from Tor's use as a browser).

> I'm not really a fan of Onion Services myself. I think it's nice from a technology point of view. It's nice if you can think about systems, and that's the classical thinking that I was used to before all this public visibility. The technical community accepts that it's currently all crap, and all shit happening on the Darknet because it's

> technically so neat… I'm not sure that just because there are potential worlds where Hidden Services would save the planet, it's maybe not the world we live in - *Participant–L - Core Tor contributor*

> I think it's an absolute disaster… Tor's public perception has been really bad… I think the most important thing they could do is, like, rebrand, and have a decent PR person… if you look at it from the outside, it feels like some underground, dodgy, like, drugs trading thing… this whole "Dark Web" bullshit means that Tor gets lumped in with Silk Road. *Participant R – Tor advocate and relay operator*

This conceives of privacy technologies as possessing substantial power to act as moral agents, shaping public debate and influencing policy and legislation. The activist strategy is to engage directly in these public debates, making explicit cases for Tor as possessing intrinsic political values, and being intended for particular uses and political causes. In doing this, they seek to reclaim Tor as not about crime, but about control; itself at the vanguard of a wave of moral reaction against mass surveillance and authoritarian attempts by powerful groups to control the internet. By engaging in these public conversations, they attempt to get governments, institutions, and public opinion on their side. This involves promoting particular positive use cases of Tor, making the case that Tor 'isn't about' the cryptomarkets and illegal pornography (and arguing that this represents a very small percentage of Tor's actual users). Rather, they claim that Tor was designed for a particular set of intended uses – namely, for journalism, human rights work, and the protection of everyday internet browsing from mass surveillance.

> You need to be working out how to present the good use cases along with the bad ones. I think they're still learning as an organisation how to do that, they've not really had to do that for the last decade, because they've had a bunch of government

> funding, and they've been able to tailor it to what they want to do. Now that they're more reliant on people and outside organisations for funding, well it looks like… they have to get better at selling the technology as a whole to society - *Participant–Z - Tor Onion Services developer*

As a result, they articulate a vision for Tor which is rather different from the neutralised status Tor has asserted over the years, or the structural change through engineering imagined by its designers. This world of discourse is more likely to accept publicly that harmful uses of Tor are a problem, and to condemn particular use cases of Tor, especially those which are associated with crime or the far right.

By explicitly allying Tor with other social movements, such as women's rights, LGBT liberation, civil rights, they attempt to ensure that perceptions of Tor are steered by its community and to *reclaim* Tor's values. They do this by promoting particular use cases, allying with particular causes, and partnering with the particular groups which the activists choose to train in how to use Tor. This has the advantage of empowering Tor to use its substantial clout in lobbying for privacy as a political cause and shaping public perceptions of Tor to improve its image. This, however, faces problems in practice, clashing with *infrastructuralist* sensibilities in the Tor community who are both unused to acting in the domain of public discussion and deeply suspicious of associating technologies with an explicit politics.

Democratisation - from disruption to governmentality

These three distinct strategies – neutralisation, reclamation, and standardisation – have facilitated Tor's disruptive ethos as a privacy infrastructure throughout much of its lifetime. From an initially engineer-focused approach, this expanded to more administrative concerns as the network infrastructure was implemented and grew, and in more recent years, as abuse

became more prominent an issue, the activist approach to conversations about Tor's values has become more important (for a more comprehensive exploration of these worlds and how they are changing, see (Collier, 2020)). However, these shifts in how Tor makes sense of and deals with abuse have also been part of a broader shift in how it understands itself as an organisation, and a fourth strategy has emerged which draws from the engineer and activist worlds: democratisation.

Tor has grown substantially over the last ten years, taking on responsibility for an ever-larger and more diverse global population of users. This move beyond its initial adoption by technologically-skilled privacy enthusiasts has seen the growth not only of problems with crime and abuse, but also increasingly with broader problems of accessibility and the Tor Project's responsibility for its user community. The way in which Tor works (as is the case for any technology, platform, or infrastructure) is fundamentally reflective of values and visions of its developers. When it was a much smaller platform, and its users were more similar to the people who develop it, i.e. US and European privacy-concerned software developers and other technically-skilled individuals, this was less of an issue and approaches to abuse based on neutralisation, making the case for Tor's values, and standardisation appeared largely unproblematic. However, the expansion of Tor's userbase globally has revealed that many of the implicit assumptions built into the technology (and the training required to use it safely) may not actually hold for a number of types of user whom it now sees as critical to its expansion. Tor is increasingly aware of these issues, and what it means for their own intervention in power relationships:

> I think that's a valid argument against Tor. That no matter how much you try to educate people to be able to use it, ultimately you are supporting the power structures… And in that sense, then Tor becomes a weapon against those that just don't know how to use it, right? *Participant–L - Tor core contributor*

As Tor's social worlds have grown and changed, (see (Collier, 2020)), they have increasingly tried to engage reflexively with this problem of the power they themselves wield to shape the Internet for their users. One of the ways in which the engineer world has approached this is through a lens very familiar to developers – that of *usability*. Concerns with usability go back very early in Tor's development, from its earliest desires for as wide a user base as possible, and hence making the network as fast and reliable as feasible, to the compilation of the initially complex range of tools which users needed to configure into a single, easy-to-use Tor browser program. But beyond this, it has become increasingly clear that Tor's broader usability design is a real issue – the ways that language, subtle interface design elements and other technical aspects for Tor do not necessarily reflect how different groups and societies around the world may use the Internet differently and have radically different privacy and security needs.

Addressing this has involved the beginnings of a wide campaign of user research, involving global outreach, improvements in communication, the creation of a 'community portal', and a shift in design processes to centre them around the needs of important user groups (which they term 'personas') identified by this outreach, which may be less well served by Tor as it currently exists. Focusing design around particular (positive) use cases is in itself an adaptation to harmful or abusive uses, shaping who Tor is 'for' both in the normative, activist sense, and in the engineer sense of its technical properties.

All this is reflective of a more *governmental* character emerging as Tor has developed, through which it is attempting to learn more about its users, alter the design elements and category systems within its technologies to better represent them, and take responsibility for their experience of using the Internet (and hence important parts of their social, economic, and political lives). In doing this, Tor is attempting to actively and reflexively develop a

governmentality (or rationality of power) of its own, trying to more carefully and democratically shape how it acts in this space of infrastructural power.

Discussion and concluding thoughts: infrastructural power and its limits

Where platforms and infrastructures attempt to deal with and conceptualise abuse and crime, they by necessity become involved in the business of government. The four approaches which I identify in this chapter to dealing with abuse and crime without taking an active role in policing conduct – standardisation, neutralisation, reclamation, and democratisation - are not unique to Tor, and underlie many of the more traditionally policing-oriented strategies which platforms employ.

As I outline in this chapter, how Tor makes sense of abuse and crime is inherently bound to the ways in which it makes sense of the values at its heart, especially its constructions of privacy. I present by necessity only a partial perspective, reflective of the thoughts and experiences of the people who agreed to be interviewed. Although the sample is fairly reflective of the broader Tor developer and maintainer community, there are undoubtedly missing perspectives in this discussion (such as the users of Tor), and it is generally weighted towards respondents in Europe and the USA. Additionally, the characterisation of Tor's values presented is particularly reflective of contemporary, rather than historical, issues: though I have attempted to fill in historical perspectives from mailing lists and through interviews with older members of the community, there are some eras of Tor's history which are more lightly sketched.

Bringing Tor's history up to date, at the time of writing, the world has been undergoing unprecedented social change in response to the COVID-19 pandemic. This has had immediate impacts on Tor, whose recent expansion had, as discussed above, been based on a diversified funding model which has collapsed in a context of massively reduced

discretionary spending. Although Tor will persist, they have had to reduce staffing considerably to focus on the core development and maintenance work of Tor. The cultural orientation of Tor's core group will not disappear overnight, however this will undoubtedly reorient their focus.

The utopian engineers who evangelised the possibilities of the Internet age argued that it promised to open up infrastructural power to ever-smaller groups of people looking to realise a vision of the world through technology (Milan, 2013; Yar 2012). Tor is a particularly striking example of this: a small, relatively poorly-resourced organisation which has nevertheless built a global infrastructure that provides unprecedented privacy, security, and anti-censorship protections for its users. As Tor has grown, it has encountered issues of crime in a range of different domains: design, administration, and public image. These constitute distinct sites at which abuse is managed and conceptualised; each of these packages up its own interpretation of the broader values and purpose of the organisation, and its own way of conceptualising abuse. The way in which Tor survives, and how it goes on to shape society, depends a great deal on the tensions between its social worlds and how they are worked out in practice.

It is important to note that the relationship between Tor and law enforcement is a complex one. Although Tor's engineer world (and hence its design) seeks to limit the structural powers which law enforcement can establish over the Internet infrastructure, they are not *necessarily* anti-police (and although some members of the team advocate an activist-based critique of law enforcement from anarchist political theory, many are supportive of targeted police surveillance where there is a reasonable suspicion of wrongdoing). In fact, Tor has become a vital tool for online law enforcement, allowing them to hide their identity from the administrators of illicit services when conducting investigations. Senior members of the Tor Project have often embraced this and have been keen throughout its history to conduct

outreach with law enforcement organisations in order to dispel its reputation as a tool for crime. In some cases, this has provoked controversy: while the infrastructuralist world welcomes the veneer of neutrality which this brings, many of a more activist disposition are uncomfortable with appearing to support targeted police surveillance (and the structural power on which it relies) or the use of Tor in the service of (often US) geopolitical hegemony in exchange for countering mass-scale techno-policing, due to the fact that targeted policing tends to be focused on groups who already experience structural oppression within society (an argument developed in detail in Guerses et al 2016). These tensions within Tor, designed as it was in the service of covert US operations and from a long history of US soft power moves, grant it a complex status as a liberation technology.

It is tempting to imagine that societies can be 'designed' through infrastructural solutions, both for those attempting to exert state control and those attempting to resist it. In fact, this struggle cannot be fought solely in the terrain of design discussions and technological solutions. How a platform or infrastructure deals with crime, abuse, and user conduct stems from the values which underpin it, however, these are rarely monolithic. While Tor has eschewed the internal policing of content and conduct (and the collaboration with state law enforcement) which platforms like Google and Facebook, or Internet Service Providers employ, they too have traversed journeys from disruptive innovators to an essentially governmental character, which can be seen most clearly in changes in their conceptions of users and usability. As such, the development of this character is not dependent on engagement with existing state law enforcement, or on the direct policing of content and conduct, arising instead from the exercise of distinct forms of infrastructural power over populations.

Tor's governmental character is fundamentally different from that of Facebook or Google, focused around user privacy and the democratisation of their own infrastructural power. This

suggests that the governmentalities developed by platforms and infrastructures as they grow arise from and reflect their own particular values. These values are mediated by the different kinds of work on which these infrastructures depend, the sites at which they come into contact with governmental issues such as crime and conduct, and the social worlds which grow from them. Tor's approach to these issues points to alternative governmental approaches for online platforms and infrastructures, which are more democratised and less reliant on policing-through-design, internal abuse handling, or collaboration with state surveillance and censorship[1].

Bibliography

Aldridge, J., & Decary-Hétu, D. (2016). Cryptomarkets and the future of illicit drug markets. *The Internet and drug markets*, 23-32.

Barratt, M.J., Lenton, S., Maddox, A., and Allen, M. (2016), "What if you live on top of a bakery and you like cakes?" Drug use and harm trajectories before, during and after the emergence of Silk Road. *International Journal of Drug Policy*, 35, 50-57

Bauman, Z., & Lyon, D. (2013). *Liquid surveillance: A conversation*. John Wiley & Sons.

Box, S., & Muncie, J. (2002). Crime, power and ideological mystification. *Criminology: A Reader*, 81.

Clarke, A. E., & Star, S. L. (2008). The social worlds framework: A theory/methods package. *The handbook of science and technology studies*, 3, 113-137.

Collier, B., (2020), The power to structure: exploring social worlds of privacy, technology, and power in the Tor Project, *Information, Communication, and Society*

[1] Notable examples of this include www.openprivacy.ca and social media platforms such as Mastodon

Cornish, D. B., & Clarke, R. V. (2003). Opportunities, precipitators and criminal decisions: A reply to Wortley's critique of situational crime prevention. *Crime prevention studies,* 16, 41-96.

Dingledine, R., Mathewson, N., & Syverson, P. (2004). Tor: The second-generation onion router. Naval Research Lab Washington DC.

Foucault, M. (1991). *The Foucault effect: Studies in governmentality*. University of Chicago Press.

Garland, D. (1997). Governmentality' and the problem of crime: Foucault, criminology, sociology. *Theoretical criminology*, 1(2), 173-214.

Gillespie, T. (2018). *Custodians of the Internet: Platforms, content moderation, and the hidden decisions that shape social media*. Yale University Press.

Gürses, S., Kundnani, A., & Van Hoboken, J. (2016). Crypto and empire: the contradictions of counter-surveillance advocacy. Media, Culture & Society, 38(4), 576-590.

Kohl, U. (2013). Google: the rise and rise of online intermediaries in the governance of the Internet and beyond (Part 2). *International Journal of Law and Information Technology*, *21*(2), 187-234.

Lyon, D. (2014). Surveillance, Snowden, and big data: Capacities, consequences, critique. *Big Data & Society*, *1*(2)

Milan, S. (2013). *Social movements and their technologies: Wiring social change*. Springer.

Moore, D., & Rid, T. (2016). Cryptopolitik and the Darknet. Survival, 58(1), 7-38

Marechal (2018). PhD Thesis: Use Signal, Use Tor? The Political Economy of Digital Rights Technology

Nissenbaum, H. (2009). *Privacy in context: Technology, policy, and the integrity of social life.* Stanford University Press.

Pinch, T. (2010). The invisible technologies of Goffman's sociology from the merry-go-round to the internet. *Technology and culture*, 51(2), 409-424.

Sandvig, C., Hamilton, K., Karahalios, K., & Langbort, C. (2014). Auditing algorithms: Research methods for detecting discrimination on internet platforms. *Data and discrimination: converting critical concerns into productive inquiry*, 22.

Star, S. L., & Griesemer, J. R. (1989). Institutional ecology,translations' and boundary objects: Amateurs and professionals in Berkeley's Museum of Vertebrate Zoology, 1907-39. *Social studies of science*, *19*(3), 387-420.

Schulze, M. (2017). Clipper meets Apple vs. FBI: a comparison of the cryptography discourses from 1993 and 2016. *Media and Communication*, 5(1), 54-62.

Suzor, Nicolas P., et al. (2019), What do we mean when we talk about transparency? Toward meaningful transparency in commercial content moderation. *International Journal of Communication* 13, 18.

Valverde, M. (2009). Beyond Discipline and Punish: Foucault's challenge to criminology. *Carceral Notebooks*, 4, 201-224.

West, S. (2017). Raging against the machine: Network gatekeeping and collective action on social media platforms. *Media and Communication* 5.3, 28-36.

Yar, M. (2012). Virtual utopias and dystopias: The cultural imaginary of the Internet. *Utopia: Social Theory and the Future*, 179-95.

Zuboff, S. (2015). Big other: surveillance capitalism and the prospects of an information civilization. *Journal of Information Technology*, *30*(1), 75-89.