



THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Privacy worlds

Citation for published version:

Collier, B & Stewart, J 2021, 'Privacy worlds: Exploring values and design in the development of the Tor anonymity network', *Science, Technology, and Human Values*. <https://doi.org/10.1177/01622439211039019>

Digital Object Identifier (DOI):

[10.1177/01622439211039019](https://doi.org/10.1177/01622439211039019)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Publisher's PDF, also known as Version of record

Published In:

Science, Technology, and Human Values

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.



Privacy Worlds: Exploring Values and Design in the Development of the Tor Anonymity Network

Science, Technology, & Human Values

1-27

© The Author(s) 2021



Article reuse guidelines:

sagepub.com/journals-permissions

DOI: 10.1177/01622439211039019

journals.sagepub.com/home/sth



Ben Collier¹  and James Stewart¹

Abstract

This paper explores, through empirical research, how values, engineering practices, and technological design decisions shape one another in the development of privacy technologies. We propose the concept of “privacy worlds” to explore the values and design practices of the engineers of one of the world’s most notable (and contentious) privacy technologies: the Tor network. By following Tor’s design and development we show a privacy world emerging—one centered on a construction of privacy understood through the topology of structural power in the Internet backbone. This central “cipher” discourse renders privacy as a problem that can be “solved” through engineering, allowing the translation and representation of different groups of imagined users, adversaries, and technical aspects of the Internet in the language of the system. It also stabilizes a “flattened,” neutralized conception of privacy, risking stripping it of its political and

¹University of Edinburgh, United Kingdom

Corresponding Author:

Ben Collier, 1F2 6 Bonnington Road, Edinburgh, EH6 5JD, United Kingdom.

Email: ben.collier@cl.cam.ac.uk

cultural depth. We argue for an enriched empirical focus on design practices in privacy technologies, both as sites where values and material power are shaped, and as a place where the various worlds that will go on to cluster around them—of users, maintainers, and others—are imagined and reconciled.

Keywords

design, privacy, infrastructure, values, social worlds, darknet

Introduction

The history of the Internet is one of bitter struggles over values, imagined future societies, and how these become realized in infrastructure. Although the Internet has changed dramatically in the 21st Century, the controversies and conflicts that underpin it remain embedded in its material, informational, and discursive forms. These issues are far from settled, with battles over online power and governance taking place not only at the level of content, but far further down the “stack,” in the fundamental protocols, technologies, and standards on which the Internet infrastructure is based (DeNardis 2012). In particular, privacy has become a key battleground in which these issues of online politics and power are contested (Lyon 2014; Zalnierute and Milan 2019). Engineers have created infrastructures and tools—Privacy Enhancing Technologies (PETs)—that users can deploy in different ways to create privacy for themselves in an online world defined by surveillance. In this paper, we draw on empirical research and STS theory to explore the role played by values and practices in the development of PETs. We contest a static model, in which relatively stable prior constructions of key concepts like privacy are unidirectionally “inscribed” into technologies, proposing instead one in which privacy values emerge, cohere, and are stabilized throughout development.

Privacy scholarship across a range of domains increasingly renders privacy as a contested concept: what privacy is for, how it might be achieved, the different contexts in which it is realized, and the work involved in producing privacy are revealed to be diverse and heterogeneous (Nissenbaum 2015; Lewis 2017; Bancroft and Reid 2017). To cut through this complex field, in which definitions and realizations of privacy seem to spill out and multiply chaotically, we develop the concept of “privacy worlds”: the sets of coherent, relatively stable formations of practices (and related

values, contexts, and constructions of privacy) that form around privacy technologies. In this paper, we turn to the engineers who design these technologies, whose privacy worlds are particularly crucial as they set in place the technological foundations of all the other worlds (of user groups, maintainers, and others). By empirically exploring the design processes through which these infrastructures take shape, we make sense of the role of values in these design processes—and how these processes create space for and shape the subsequent privacy worlds that develop around PETs (Star and Griesemer 1989).

One such privacy technology is Tor, an anonymity network and secure browser often referred to as the “dark web” (although this is a gross oversimplification of its uses and social relevance). Tor constitutes a network of volunteer-run servers that bounce encrypted user traffic among them, hiding their origin and destination from the Internet Service Providers and nation-states who own and administer the Internet’s fundamental backbone. The software used to access this network is free and accessible via the Tor Project’s website. Tor has become the technological “crown jewel” of the “Internet Freedom” movement as a widely used, successful, and effective defense against metadata surveillance—the practice of surveilling not the content of messages, but the administrative information that can reveal the identities of the people sending and receiving them. We study Tor in its earliest design phase, drawing on empirical data, including interviews and archival research, in order to map the emerging “privacy world” of its engineers.

We begin with a discussion of online privacy and power, then set out our theoretical and methodological approaches and our “privacy worlds” framework, drawn from social worlds theory. After discussing our methods, we introduce the historical context of Tor and Onion Routing. Subsequently, we work through a single, particularly important controversy in the history of Tor’s design: the inclusion of “padding traffic” that would have slowed down the network but potentially made it more secure against the most powerful state adversaries. In doing so, we map the emergence and stabilization of a “privacy world” and the design of a technical paradigm for attempting to produce online privacy. We conclude by reflecting on the implications for values in Tor’s design and the broader utility of the “privacy worlds” framework.

Privacy and the Values of the Internet

The extent to which Internet infrastructures relate to power and values has been a core preoccupation in STS literature for several years. The

scholarship of Musiani and DeNardis contributes particularly fruitful approaches to thinking about and mapping these relationships. This body of work uses theoretical frameworks and methodological approaches from Science and Technology Studies to map how the technical infrastructures of the Internet become both objects of governance (which themselves need to be steered, managed and controlled) and *sites* of governance; acting as technologies of control through which populations are governed and as sites at which issues of societal governance are contested and worked through (Musiani 2013; DeNardis and Musiani 2016).

Musiani and DeNardis describe a “turn to infrastructure” (Musiani et al. 2016) in state-led Internet governance after decades of liberalization, through which states are increasingly attempting to reassert power over the lowest levels of the Internet infrastructure stack to achieve economic, social, and other policy goals in a wide range of areas. Musiani further describes the shaping role that the *architectural design* of Internet infrastructures have on those who use them and the social institutions and processes that rely on or are implicated in them. Each architecture embodies distinct conceptions of users and social facts such as privacy, promoting particular modes of action and interaction (Musiani 2010, 2012, 2015). In this paper, we use the term “infrastructure” to mean large-scale formations of interconnected technological artifacts that are reliant on routine maintenance and can be turned to a variety of purposes (Star and Ruhleder 1996).

DeNardis describes the crucial role played by “control points” within Internet technical infrastructures; sites at which control can be exerted over societies by influencing technical standards or the administration of networks (DeNardis 2009, 2012; Bradshaw and DeNardis 2019). The underlying design of the Internet itself embeds qualities that create control points: the protocols that route Internet traffic are inherently traceable, allowing states to exert control by *surveillance* and *ensorship* through the Internet’s physical infrastructure of routers, servers, and wires (Lyon 2014). However, the capacity of the Internet to be extended and adapted through the creation of novel, “higher level” infrastructures and technologies allows the development of further control points that may be out of reach of the state, empowering engineers to build new networks and shape structures of power (Musiani 2012; Milan 2013). This often involves contestation and resistance: political, commercial, and occasionally illicit. Concretely, this body of research depicts power in Internet societies as increasingly fought around *topologies* of design and control in network architectures. The infrastructure of the Internet is still constantly being developed and contested—by the major players such as Google and Facebook, by nation states, by lawyers

and activists, and by hackers attempting to build their own extensions to the Internet or subvert existing infrastructure (Dencik et al. 2016; Milan 2013).

The social issues and contestations of power that emerge at these Internet control points have been framed through a range of different lenses, but the lens of privacy has dominated much of this discussion (Lyon 2014; Nissenbaum 2015; Lewis 2017). Privacy Enhancing Technologies—of which Tor is one of the most well established—seek to transform the properties of the Internet in order to create privacy for their users; more broadly, all Internet technologies embody (for better or worse) a set of privacy properties relating to how they store and manage their users' data. In addition to technical and engineering research on privacy technologies, a growing body of research explores the relationships between understandings of privacy¹ and the design of these technological systems (Danezis and Guerses 2010; Musiani 2012; Crabtree, Tolmie, and Knight 2017), positing that the privacy properties of technologies are material realizations of their designers' understandings of privacy (Musiani 2010; Danezis and Guerses 2010). These constructions of privacy go on to shape the experiences of the users of the subsequent technical systems/infrastructures (Bancroft and Reid 2017; Pfitzmann and Hansen 2005), although the users themselves have their own constructions of privacy and privacy practices that may employ these technologies in ways unexpected by the designers (Lewis 2017).

Danezis and Guerses suggest three distinct modes of framing privacy within discussions of privacy technologies: privacy as confidentiality, privacy as control, and privacy as practice (Danezis and Guerses 2010). One of the strongest strands has been privacy enabled by *data protection*; the translation of values to law, regulation, and into obligations to technically implement data protection by “design and default” (Cavoukian 2013). This is a model for privacy protection in legal situations, based on controlling data flow, and sets of data rights, based on legitimate use. However, for activities that cannot be accounted for solely in legal terms, other approaches are needed. An alternative framing that focuses on social norms is proposed by Nissenbaum (2015): *contextual integrity*, which captures the sensitivity and handling norms attached to information as it flows through different contexts and is accessed or administered by people performing different roles. This problematizes the discovery and evaluation of relevant context and norms in the places where data will flow. This is a useful way of understanding how engineers construct and realize privacy in technical systems, capturing as it does not only information structures, but also the struggles faced by engineers in rendering and reasoning about social aspects

of privacy-protecting systems. However, the precise role of engineers and their values in these processes remains underexplored empirically.

Social Worlds, Values, and Design

Previous sociological research on Tor includes studies of how privacy technology projects cultivate legitimacy in their user communities and broader publics (Gehl 2018), Marechal's work on Tor as an activist organization (Marechal 2018) and a range of works tackling the criminalized uses of the Tor network (most notably Bancroft and Reid 2017). We are instead interested in the relationships between the values of the Tor developers and the technical design of Tor's infrastructure itself. In this section, we briefly discuss these, and our use of a social worlds approach.

The social worlds framework is an interactionist approach to the study of technology, focused not on individual technologies but on technological systems and infrastructures, particularly those involved in the production of scientific or engineering knowledge. It draws from Strauss' (1982) conceptualization of social worlds, which aimed to make sense of naturally emerging forms of organization in productive social life, in which shared sets of practices intersecting around a particular form of collaborative endeavor gave rise to broader "world views" or "universes of discourse" that could be themselves be drawn on by those outside these communities (Unruh 1980). Further work, notably by Star, Bowker, Newmann, and Clarke operationalizes this within STS, inductively studying the microscale interactions, discourses, practices, and interpretations that spill out from technological and scientific work (Clarke and Star 2008). These build up into maps of broad social worlds—the distinct clusters of self-consistent discourse and practice that emerge around complex common endeavors. This focuses analytical endeavor on practices, the discourses they produce and stabilize, and how they become embedded in the category systems and processes that surround extensive infrastructural arrangements.

The "social worlds" perspective is often used to theorize infrastructures, which involve bringing together a heterogeneous mix of types of work and people (including designers, different user groups, maintainers, and others), who often develop radically different "visions" of the technology that manage to coexist rather than compete (Star and Griesemer 1989). Social worlds theory provides a framework for how these often very different ideas, values, and practices come together in communal action to make something in the real world—a social fact that they produce together (such as a functioning museum, an art exhibition, a successful surgical operation, or online

privacy) (Becker 1976; Unruh 1980). This is particularly powerful for conducting empirical research studies of these infrastructures and how these distinct and internally self-consistent “worlds” intersect around them (Clarke and Star 2008).

The “worlds” at issue in social worlds are described as “universes of discourse” (Unruh 1980); the “discourses” here refer to a wide array of disparate forms, including meanings, ideas, aesthetics, associations, constructions of the core concept (in Tor’s case, “privacy,” in Becker’s (1976) case, “art”), links to other concepts, such as anonymity, liberty, and technology. Drawing on Clarke’s (2007) later development of social worlds as “situational analysis,” relevant forms of discourse can be mapped within three broad categories. Within the first of these categories are *shared values, ideas, and interpretations*—ways of thinking about, framing and understanding an infrastructure and the social facts it produces (Star, Bowker, and Neumann 1998). These discourses refer both to category systems and value constructions, and also to broader concepts that structure attempts to reason about the work around which the social world is organized, including overarching goals and justificatory narratives for the value of the project (Star, Bowker, and Newman 1998). The second element is the discourses *embedded in particular practices, tools, and kinds of work* (of the people building and running the infrastructure and its imagined users) that when brought together allow the project to work and the result to be produced. Finally, social worlds include the established technologies and infrastructures and their *materialized discourses*—the technologies, shared protocols, standards, and material elements around which different worlds cluster and the new infrastructural elements that are being created (Clarke 2007).

Within a social world, development is not necessarily riddled with struggle and antagonism—the picture is often of a coevolution of technologies, practices, and values (Star, Bowker, and Neumann 1998; Clarke and Star 2008). While nascent social worlds may fail to stabilize, or only coalesce temporarily, a social world truly emerges when these formations remain stable for a period of time, even if only through common goals and ideas, rather than necessarily “running code,” satisfied users, or broader social legitimacy. Social worlds are not hermetic or exclusive: membership is often multiple and shifting, with people inhabiting several of these worlds simultaneously and able to “borrow” or refer to discourses of other worlds when necessary to assist in translation between different perspectives. Worlds can and do interact and shape one another, and the shared infrastructures play a key role in stabilizing or crystalizing arrangements of these worlds.

Through empirical research into the discourses of particular social worlds, we can identify their core interpretive frames. By understanding these frames, we can better understand how these worlds work internally, how they manage conflict and consensus, how they interact with wider society, and how they succeed or fail. In a study of engineering work, this allows us to map how the core values, practices, and technologies emerge, and how engineers attempt to bring in, interpret, and make sense of crucial factors involved in design, such as users, adversaries, and constraints. By doing this work, we can therefore better understand the links between the design of technologies and infrastructures and the constructions of social facts with which they are implicated, such as online privacy.

Empirical explorations of engineering work suggest that values themselves are often actively contested and worked out *within* groups during the design process; the practices and values of engineers emerge in reference to particular projects, approaches, constraints, and contexts and cannot be taken for granted (So-rensens and Levold 1992). Nissenbaum frames the different modes of thought and discourse that emerge in design and development as a set of “balls in play,” with engineers needing to juggle between technical knowledge, abstract values, and the empirical data about systems that emerge during these processes (Nissenbaum 2001). Frequently, the engineers’ values and motivations run up against the informational or material constraints of technology, with some proving mutually exclusive in practice and needing to be discarded.

Our use of the social worlds framework seeks to deepen and extend Nissenbaum’s (2001) framing of the processes by which ideas become materialized in design. Where social worlds theory conceptualizes how value systems shape the material properties of infrastructure, it does so through the idea of convergence—a process by which human actors who take up infrastructures and form practices and discourses around them find their category systems and understandings beginning to “converge” with the discourses embedded in the technology itself as they mutually shape one another (Star, Bowker, and Newmann 1998). This paper develops this approach to explore how these processes work when an infrastructure is created for the first time—how the discourses of engineers and the properties of the technology converge as the social world is being created. We argue that at the heart of a social world is often a key framing discourse (or set of discourses) drawn from the three categories (noted above), which provides a central epistemic and interpretive frame that structures how all these elements interrelate and are incorporated into a self-consistent “world.” This framing discourse is the cipher through which all the other

discourses are interpreted and given meaning, and through which they are deployed to do work and produce, in Tor's case, the core concept of privacy. A focus on these central "cipher" discourses, as we discuss in more depth in our discussion section, develops the social worlds framework to better account for the nascent, world-forming processes that we study in this article.

Privacy Worlds

In order to cut through the complex and heterogeneous landscape of online privacy and make sense of exactly how the developers of privacy technologies and their values feature, we draw on the social worlds framework to theorize PETs as a site around which "privacy worlds" coalesce. Much as Gray's (2018) attempts to theorize *data worlds*, our privacy worlds map the different imaginaries, practices, and technologies that constitute distinct ways of making sense of and producing "privacy." We aim, however, not only to map broad topologies of privacy discourse, as Gray does with data worlds, but to draw more deeply on the social worlds framework to make sense of how these worlds emerge and the crucial role played by engineers.

The landscape of online privacy is characterized by a set of fairly coherent privacy worlds—core, shared constructions of privacy that unite disparate groups of actors. As debates over privacy and security have been fought and technologies and practices have emerged and stabilized, these "universes of discourse" have also coalesced as privacy worlds, be they of law enforcement and security services, of the Internet giants, of the Internet Freedom movement, of the engineers who build the Internet and privacy technologies, or of particular user groups. The developers involved in designing a successful privacy infrastructure—and their privacy world—are particularly important, as they directly shape the key design elements of the technologies around which all the other worlds cluster. The privacy worlds that characterize the contemporary Internet are constrained and shaped by the embedded models, representations, and values stabilized in the technology by the developers—even if they are able to exert a degree of their own appropriation, subversion, and reconfiguration (Fleck 1994; Dourish 2003).

In Tor, as we have set out in previous research (Collier 2020), one of the foundational worlds is that of the "engineers"—counterposed with the worlds of the people who run and maintain the infrastructure and the activists for whom it is a center of struggle, these are the people who design and develop the technologies that underpin the Tor network. Mapping this world, we find a rich and diverse history and set of ideas and

aesthetics—user models, political views, aesthetics, and identities, which we discuss in the empirical section of this article. The discourse at the heart of the engineer world, which provides the “key” to the relationships between its other component discourses, can be summarized as “privacy as a structure”—the idea that privacy is a matter of the structures of power and information created by the design of technical systems and the “control points” they create. This idea has been very influential; it wasn’t invented by Tor, but it characterizes much of the political tenor of contemporary privacy engineering (Musiani 2013). This contrasts with traditional cryptographer ideas of privacy, which are based in mathematical proofs and models. It also conflicts with the core discourses of Tor’s activist world, which sees privacy through the lens of the struggle and the social movement, and its infrastructuralist world, which understands privacy as a neutral service (Collier 2020). Here we are interested in how this central structuring discourse emerges and becomes the interpretive frame that connects all these other discourses, and how it is stabilized in technology and practices. Thus, we show the world of the developers as it comes together across the design process; as developers wrestle with political and social issues and try to realize privacy through engineering practices.

Method

We draw on extensive archival analysis and qualitative interviews conducted with Tor developers to explore empirically the privacy values of Tor and how they were translated into material properties of Tor’s design. This involved twenty-six semi-structured anonymous interviews with core developers, activists, relay operators, and other members of the Tor community (though in this paper we focus on the twelve interviews with developers), and extensive study of the Tor Project’s public archives. Interviews were conducted in 2016 during a particularly turbulent point in Tor’s history, as it came to terms with allegations of abuse concerning a prominent member of the community and reworked its organization around more professionalized, democratic, and value-centered internal structures and processes (Collier 2020; Marechal 2018).

The Tor Project makes the full history of its mailing lists openly available, which provides an invaluable record of some of Tor’s formative design discussions. Archives like these are an underexplored resource for sociologists, though they have contributed to some notable studies of technological development (Gehl 2018; Gueddana 2013). Having made the Tor Project aware that we would be conducting this research, we approached

this archive of material by drawing key sensitizing concepts from our interviews. From a handful of key design discussions, one core controversy was identified as fundamental to Tor's material construction of privacy. The discussions we use in this paper largely date from between 2002 and 2003—the first two years of Tor's development. Where we present quotes from the tor-dev mailing lists, we do so without identifying their authors as these are often loosely sketched thoughts and ideas rather than settled opinions or positions. Numerous papers detailing the Onion Routing design are available online, as are these archived design discussions (Syverson, Dingleline, and Mathewson 2004). Our discussion of the history of Onion Routing is informed by interviews and by the public mailing list archives available at www.onion-router.net.

Cypherpunks and the Navy: The Strange Foundations of the Tor Network

The early development of Tor and its precursors in the 1990s and early 2000s brings together many of the foundational issues of Internet privacy, governance, and ownership. Following the collapse of the Soviet Union, the Internet represented a key front in the cultural and economic battle for soft power, representing to many the core liberal tenets of the unrestricted international flow of speech, ideas, capital, and communication (Curran 2012). The US was grappling with the newly commercialized Internet and the conflicting interests of control, capital, and liberty that it posed (Chenou 2014). The tensions within the neoliberal model, in which extremely strong state power is used to protect and guarantee unrestricted laissez-faire markets, were equally at play within the Internet, as it became clear that the Internet might undermine state control and sovereign power. Commonly known as the “Cryptowars,” the most overt aspects of this conflict centered on the use of encryption and its ability to prevent US law enforcement and security services from responding to crime and terror threats. Arrayed on one side were US politicians, law enforcement, and the security services, who argued that strong encryption should be restricted to government and military uses (DeNardis 2007). On the opposing side were various civil society groups, a reflexively libertarian tech community, and, their apotheosis, the cypherpunks: a loose community of encryption experts and libertarian technologists who saw privacy as crucial to the liberatory potential of the Internet (Coleman and Golub 2008).

However, the contestations of power manifested in the Cryptowars had a second, paradoxically opposed aspect. While it was in the domestic interests

of the US state to break or regulate encryption, they found that the Internet control points held by other governments in their own countries caused difficulties for US agents operating over the Internet in other nations, who could be spotted fairly easily through analysis of the patterns of traffic that they generate. At a level above the *content* of messages, a separate issue was thus rearing its head: this metadata, which is used to route signals around the Internet, is particularly sensitive, allowing the establishment of control points through the observation of patterns of behavior over time that can be extremely revealing (Dingledine and Matthewson 2006).

Onion Routing, the anonymity paradigm that Tor embodies (and of which Tor is the prime working example), was born in the service of achieving these goals in the US Naval Research Lab in the mid-1990s. Onion Routing is an approach to achieving the separation of the *identity of the user* from the routing information used to guide signals around the Internet. This effectively nullifies one of the major “control points” built into the Internet, making the origin and destination of traffic no longer trivially identifiable by the Internet Service Provider (and hence the state). In the Onion Routing design, users’ Internet traffic is bounced around a network of volunteer-operated servers (known as “relays”) located around the world in order to disguise its origin and destination. First, the administrative information that routes users’ traffic around the Internet is wrapped in three layers of encryption. This traffic is then sent as a series of packets to the relay network. The traffic is first sent to an entry relay—a server that decrypts the first layer of encryption and reveals the address of the next relay in the chain. This next, “middle” relay decrypts the next layer of encryption, revealing the “exit” relay’s address. The exit relay then decrypts the last layer of encryption, finds the final destination of the traffic, and sends it on. Thus, no part of the network knows both the origin and the destination of the traffic, and anyone observing a particular user only sees them connecting to the network, not which websites they are accessing. This allows users’ identities to be concealed, becoming indistinguishable in a crowd of millions of other users (Syverson, Dingledine, and Mathewson 2004; Dingledine and Matthewson 2006).

In this design, anonymity is produced through the size of this crowd. The bigger this crowd of users, the harder it is to deanonymize anyone, and the more *diverse* users, the less information an adversary can deduce simply from the fact that someone is using the network. Thus, an Onion Routing system cannot be used only by the military (as an antagonistic nation state would know immediately that any connections to this network were from

military users); it needs mass public use to provide “cover” for its military users. Speed and usability are core security properties, as the network needs to be fast and simple enough to use to encourage the general public (or at least the more privacy conscious among them) to use this system for everyday, innocuous web browsing.

Well, you know, the technology’s cool, and it’s nice to make something that’s actually going to be useful and help people, but one of the really nice things about it is that you build something which by its very nature takes people who think they ought not to trust each other and work together at all, and forces them to collaborate in order to get the results that you want. And I just like the idea that you are forcing people who thought that they should never work with these other people to do so. (*Tor core developer*)

The designers in the Navy needed a wide user base for their system to provide effective cover traffic, and a chance meeting with a group of “cypherpunk” privacy technologists at an academic workshop provided a key opportunity to cultivate legitimacy among the broader community of potential users. They formed a tentative collaboration in which the cypherpunks would act as early adopters of the network, reviewing its design and development and acting as trust brokers with the wider security and privacy community. This collaboration brought the privacy worlds of these two discrete (and often opposed) groups together. The design of the Onion Routing technical paradigm itself stabilizes the conjunction of these two worlds’ distinct understandings of privacy. It depends on a symbiosis between the cypherpunks’ libertarian “everyday” conception of privacy for large numbers of relatively innocuous day-to-day Internet users, and the US navy’s “high risk” form of privacy for a small number of military and intelligence users. These distinct constructions of privacy become *mutually productive* in Onion Routing.

Crucially, this meant that the system’s code and development would need to be completely open, with the system’s ultimate design goal to be usable even by those who didn’t trust its (largely US and European) developers or administrators. A “trust-neutral,” open-source design was needed in order to cultivate the greatest possible diversity of users. The servers themselves would be run by a separate network, initially of cypherpunk privacy enthusiasts, but eventually by volunteers all over the world. However, beneath the top-level design of the Onion Routing paradigm lay a number of contentious decisions about how a real network would work in practice, one of which we now explore.

The Controversy: Padding Traffic and a Global Passive Adversary

After a few test networks, largely developed in the late 1990s, a group of developers, including some of the original designers of Onion Routing (OR), took the code base for an implementation of Onion Routing developed for an undergraduate project at the University of Cambridge and used it to develop a system for use by the general public. In this paper we do not study the initial development of the OR idea, but instead explore how Tor's engineers developed and implemented it in a design for a practical system. These developers included academics, military security researchers, and others, drawn from the initial OR project and other early anonymity systems. Almost immediately, they faced a crucial design decision: whether or not to implement "padding" defenses within Tor against a particularly powerful class of attackers. These defenses add additional fake traffic to the network, making it harder for adversaries to deanonymize users at the cost of making the network slower and more difficult to use for everyday browsing.

The developers' initial privacy discourses were drawn from Onion Routing's two precursor worlds—the cypherpunks and the military researchers. These render privacy at a rather abstract level: that of democratic values, nation state power, and whole societies. In implementation, however, issues arose that brought these discourses more concretely into conversation with much lower-level constructions of privacy—actual models of real users. The developers needed to turn these abstract privacy discourses into a "threat model"—an approach that refines an initially-abstract set of discourses about risk and mitigation into a design for a practical system through exploring particular use cases, attacks, and defenses:

The problem with anonymity is that we can build such threat models, stronger than any adversary, but then we don't know how to build a system that actually works, or, at least, is usable in that case . . . So, we would like to be in the situation where we can come up with a threat model that covers everyone, but I think, in anonymity, there's a trade-off between threat models, and then other design requirements. So, if we started off with a strong threat model, then that will naturally lead to design choices that will bring us to high-latency, and then we get something that drops usability. So, I think, what probably more happens is that there is some estimate of what attackers can do, the design consequences are worked out, and then there's iteration. In order to work out what is actually useful to people, that feeds into that process, you're right, it's hard. (*Core Tor developer*)

When this process began, the developers had two initial *category systems* drawn from the Onion Routing paradigm: a set of imagined users and the types of adversaries against which they were trying to defend. These initial category systems were very abstract as they had very little information about the capabilities of those they were trying to defend against (foreign nation state security services) and were designing for potentially all conceivable users of the Internet. Because they were trying to design as mass-use a system as possible, these category systems defined users and adversaries in terms of structural forms (rather than any particular use case or opponent). Onion Routing constructs two main categories of user: a large group of “everyday privacy” users who provide cover traffic in exchange for privacy and a smaller group of “high risk” users (Syverson, Dingedine, and Matthewson 2004; Collier, 2020). Early discussions about users do drill down into more detail, but the user discourses cluster within these two broader abstract categories:

“Somebody is watching cnn.com, say some guy in China.”

“A group of CIA agents are deployed around the world, and check back with the cia.gov site periodically.”

“Amnesty International allows anonymous story submission. Reporters risk their lives going to rural Asian countries, and surface every so often to submit a story, to pass back lists of contacts, etc.”

“Anne logs in every day and checks these 4 news sites; it would make Anne unhappy to not be able to use our system for that. (*Selected quotes from developers, tor-dev mailing list, 2002*)

Similarly, the developers began their design work on Tor with very little information about adversary capabilities. As such, Tor’s category system of adversaries is not based around specific real-world actors, such as the Chinese or Russian governments, but is rather conceptualized as a set of abstract categories based on how much of the network adversaries can observe, and the power they can exert in different places. These categories included the *global passive adversary*—an adversary that can passively monitor traffic between all users and servers around the world; the *global active adversary*—an adversary that can actively interfere with or otherwise modulate traffic between all users and servers; and the *roving adversary*—an adversary that is able to control and observe smaller parts of the Internet and the Tor network and whose coverage changes over time.

Global adversaries are particularly problematic for the Onion Routing design, which relies on adversaries having only a partial view of the international Onion Routing network (Murdoch and Danezis 2005). These adversaries who are able to observe the whole Internet, much like the picture of the NSA's global surveillance revealed by Edward Snowden in 2013, are able to observe the timings of the "cells" of information sent around the Tor network, then use them to trace traffic around the network and deanonymize users through "timing attacks." Where people have particularly recognizable patterns of activity—visiting the same websites regularly, or speaking to the same people—this can aid deanonymization.

In the original Onion Routing design discussions, this attack is thwarted by using additional fake cells of data to complicate this traffic analysis (Dingledine and Matthewson 2006). This "padding traffic" can be added in a variety of different ways—for example, by forcing all the servers in the network to constantly send data to one another at random—that make these timing attacks much harder. However, these defenses often slow down the system, and hence would reduce Tor's usability and suitability for everyday Internet browsing. Whether or not to include this padding traffic was one of the foundational controversies of Tor. The subsequent section follows the developers through this controversy and its consequences in detail.

Implementing Onion Routing

The engineers began with the assumption that some form of padding would be necessary to defend against the global passive adversary. In tension with this were the other core discourses underpinning the Onion Routing paradigm: its technical design explicitly links privacy to usability and speed, which would be impacted by padding traffic. The developers needed to work out the tension between the core discourses driving the project through exploring the practical consequences of different potential designs.

The way these decisions were made looked rather different to the traditional picture of inscription, in which different groups of actors compete to inscribe their own more-or-less coherent sets of values and understandings into an artifact. Rather, in Tor's case, there was remarkable consensus around goals, and their initial privacy discourses, though strongly held, were fairly amorphous in practical terms, only stabilizing across the course of a substantial amount of work:

When you were saying, did you do this, or did you do that . . . I was going to say, yes! Because I do think that it evolved over time. I know, sometimes,

security research goes where somebody has a very well thought-out, theoretically analysable, mathematical argument for something, and then they try to design a system that meets that. But for us I think the idea of what security we wanted, and how to reason about it, and the system design, all kind of grew up together. And I think that actually makes sense, because if you're doing something that's really new you don't know what makes sense, and you could start with the abstract model, but you have not so much reason to think that that model is the right one. I mean, we went back and forth, and people do analyses, and then argue about the nature of the properties that are useful. And I mean, sometimes we were aware of things on an abstract level, but the data changed things. (*Tor developer*)

The abstract models of users and adversaries provided by the Onion Routing paradigm quickly proved insufficient for the work of making decisions about how to build a usable system. The developers attempted to reason about the levels of anonymity protections provided by different designs, and cast about for different ways of doing this. They experimented with “entropic” models of anonymity that render anonymity through a mathematical lens—that is, anonymity can be calculated as the size of the group of indistinguishable (based on the information available to attackers) users of which one is a member. Thus, if attackers can only find out the operating system used by users of Tor, and twice as many users use Linux as do Windows, the Linux users gain twice as much privacy from the system as those using Windows. This has some advantages, allowing for easy reckoning of the anonymity conferred to users in different scenarios.

As described by Syverson elsewhere (Syverson 2009), this framework quickly became inadequate as more complex scenarios were brought in, especially where ideas about trust, motivation, and behavior were represented. For example, particular users or parts of the network might be so valuable to an attacker that they are willing to pay a substantial amount of money to compromise them, making reasoning that casts all users as equal participants in an entropy calculation flawed in practice (a similar point is made in Guerses, Kundnani, and Van Hoboken (2016)). The developers needed to be able to reason about threats and risk in a way that brought social, technical, and mathematical factors into conversation with one another. They needed to render these distinct discourses in a common language that could translate between these three domains. This involved transforming social factors, such as properties of users and adversaries, into technical representations, by mapping them as topological patterns of information, power, and risk in the system. This takes much the same form as

DeNardis and Musiani's depictions of "control points" in Internet infrastructures, but worked out in fine technical detail.

Informally I think [the roving adversary] reflects the capability of an attacker to root several machines very quickly but can't hold on to them for very long (sysadmin having a late night and figures out something is going on or some other form of [intrusion detection system] etc). (*tor-dev mailing list 2002*)

But, what is reasonable in [the roving adversary] is the partial compromise of the network. An adversary has a budget, and short of a systemic vulnerability, he must compromise individual network elements or set up his own. (*tor-dev mailing list 2002*)

This allowed the developers to assess the practical consequences of different implementations of padding traffic for usability, resilience, security, and a range of other system factors. In the following quote, the developers use these discursive design practices to translate social factors into engineering considerations, allowing them to reason about how long it would take to deanonymize different use cases by mapping out the information structures and patterns in each case:

- * If there are more users, it may take longer [to deanonymize them].
- * If Alice's behavior isn't very odd (that is, if she behaves similarly to other users), it may take longer.
- * If other users are online more often, or Alice is online more often, or Bob is online more often, it may take longer.
- * If Alice sends requests to a bunch of people besides Bob, it may take longer (or it may not improve anything at all—wouldn't it be neat to be able to show that.)
- * If Alice refrains from talking to Bob as often, then it may take longer. (*tor-dev mailing list 2002*)

Once these representations were formalized, the developers could engage in "attack brainstorming": iteratively attempting to work out the consequences of different kinds of attack, adversary, or use case.

It's like, someone presents a solution to this problem. And then usually what happens is that a bunch of people think through this and then come up with attacks to it. Um, and it's like, hey, what if someone did this, what if someone did this, what if someone did this? And you kind of iterate on it until you come to a point where all of the attacks you can think of in this space fail

against your solution. I mean, unless someone comes up with something that's completely different, or comes up with an attack that completely subverts that, that is your working model of how things are going to be. (*Tor community developer*)

In doing this, the developers take user models and render them into forms of discourse that can be tested, mapped, and incorporated into the emerging internal logic of the system. They interrogated each of their core adversary and user categories in this way, mapping different potential geographies of information and control, and the consequences this bore for Tor's users in each case.

As they worked through these different scenarios, refining their abstract user and adversary categories, a conclusion began to emerge—rendered as discourse, but appearing to reflect a *material* constraint revealed through testing, discussion, and engineering practices. Firstly, the everyday types of online activity that they were trying to protect are inherently patterned; users want to speak to the same people repeatedly, have long-term, linkable relationships, and regularly visit the same websites. The administrative traces left by these activities are extremely distinctive, providing attackers with a wealth of different ways to characterize individuals and deanonymize their Tor traffic. As they mapped these patterns in practice, they realized that protecting against traffic analysis attacks would require a degree of padding so onerous that the network would become unable to support everyday browsing:

Here's my point about padding. Right now I'm not convinced there can be padding/throttling regimen that is both useful and practical, or maybe even either useful or practical. (*tor-dev mailing list, 2002*)

Secondly, as they refined their adversary categories, they realized that the idea of the global passive adversary was both “too weak, and too strong.” In practice, a global view of the Internet is extremely hard for even nation states to attain. Equally, they realized that any adversary who is able to maintain a global view of the Internet passively will have access to a range of other, “active” attacks, such as delaying or modulating signals entering and leaving Tor nodes that padding does nothing to stop.

I have a basic problem with the idea of global passive adversaries. As an academic exercise, it seems fine, but it is hard for me to imagine an adversary that is powerful enough to be global but weak enough to be entirely

passive . . . The global passive adversary is a fairly clean notion so perhaps it should still be pursued for abstract analysis purposes, but I need way more convincing than I've seen to design against it. (*tor-dev mailing list, 2002*)

In the end, what we said was . . . because it's so easy to do the end-to-end timing correlations, we weren't going to bother to add overhead of any-padding, until somebody could come up with a design where we thought that it was reasonably helping to, to raise the bar. You know, so that it was actually worth it. (*Tor core developer*)

The emergence of these two discourses through testing and discussion led the developers to make the decision to remove padding traffic and global adversaries from the design of Tor. This was deeply consequential for Tor, enabling it to provide a relatively fast network that was usable for everyday browsing of the Internet. This has continued to the present day, where Tor is able to be used even for file transfer and video streaming. Tor, as a result, has the widest possible spectrum of use cases for an anonymity network that also offers high levels of security against all but the most well-resourced attackers.

This process stabilized not only a key aspect of Tor's design, but also a set of practices that embodied a core structuring discourse that organizes the "engineer" world to this day: the idea that online privacy is a matter of structural power traced in technical networks.

I see the work that I do as decentralising and distributing power. Because I think that's always a good thing. I see that as a fundamental . . . like, if nothing else is true in the world, distributing power in this world is a good thing. (*Tor community developer*)

This discourse, which constructs "privacy as a structure," has had a significant impact on the way that the Tor community thinks about itself more broadly, shaping the way it justifies its political status, its approach to governance of its network, and its public statements about how it is used, including the harm and crime that it inadvertently facilitates (discussed in more depth in Collier 2020). It emerged (and was enacted and rehearsed) through the practices the Tor developers used to navigate the design discussion. It also structures the interrelations between the distinct discourses that make up the engineer world: the discourses that shape their design practices, the ideas and values that animate the project or that emerge in its wider political context, and the technical discourses about how Tor ought to work.

Discussion: Privacy Worlds

Tor has its roots in the fusion of two precursor privacy worlds—of the cypherpunks and the military developers—around a technological design whose principles aligned at the overlaps between these two distinct visions of the future Internet. The resulting privacy world—of the engineers—that stabilized across Tor’s development embodies a set of privacy discourses and practices. Tor’s developers saw the Internet as a democratizing force and global good—a space of liberation—but also potentially a catastrophic source of authoritarian control. They intrinsically understood that the real power is in the infrastructure; this infrastructural thinking is often at the heart of privacy technology and tends to reflect, as Tor does, a universalist approach to privacy.

This meant that the Tor developers needed to reckon directly with the “control points” and topologies of power that they were trying to change by extending the Internet with new infrastructure. In doing this, they had to render not only these control points in the discourses and design of the system, but also a range of other factors—users, practices, and adversaries. In attempting to reckon with the vast, heterogeneous set of different potential users and their privacy values and practices, Tor’s engineers turned on their head what are now considered core aspects of user-focused design (Stewart and Williams 2005), beginning instead with broad abstractions and refining generic constructions of users based on their patterns of behavior (Pollock and Williams 2008). This was crucial for realizing the users in discourse but also necessarily “flattened” the formations of practice and discourse that would make up the rich privacy worlds of Tor’s future users, stripping them of their wider contexts.

We have considered how design issues in Tor moved and mutated between practices, discourses, and technological aspects, and studied the relationships and stable formations that grew among them. The developers created a set of design practices (in the form of discursive frameworks, which might also be thought of as semantic tools) to represent and reconcile their values, but these practices themselves were crucial in stabilizing the core discourses on privacy that organize the privacy world of the engineers. There are key points in our story where a leap in one domain radically changes the others: where the user discourses evolve, requiring reshaping development practices, where a change in practices (as previous ones prove obsolete) suddenly spills out a range of new discourses, or where a technical breakthrough necessitates new ways of working and thinking about the project. As these different kinds of discourses came together, they stabilized

around a single core discourse—privacy as a structure. This core discourse, embedded in the threat modeling practice of the developers, served as a cipher—an interpretive framework through which users, adversaries, technical features, and wider ideas about privacy could be passed into the common language of design and engineering. It is embodied in values, practices, and the tech itself, serving as a wider interpretive frame (or “cipher”) through which development work on Tor can be done and privacy can be understood—the heart of the engineers’ privacy world. This gives us a view of privacy not as context but as conversation, with privacy values not only inscribed into tech, but used actively to navigate design problems—a process that spills out and crystalizes its own discourses on privacy.

The Tor engineers are not a lonely star—in fact, they are now orbited by a rich collection of other privacy worlds. We leave for other papers a discussion of what happens when the technology enters the real world and the users stop being abstractions—how messy user practices and privacy discourses then form their own stable worlds around the technology (as documented in Bancroft and Scott 2017), or the supportive worlds that form around the work of keeping the network running and funded (Collier 2020). Remarkably, and against all current wisdom, Tor demonstrates that it is possible to build very powerful and successful privacy systems with the users considered (initially) purely as abstractions. But this is an ethic in itself, bound up in values and with its own pitfalls: these abstractions do not necessarily reflect the privacy worlds that emerge when the technology “goes live” and meets its real users, who have their own stakes in Tor’s social meaning. In recent years, a change in the discourses of Tor’s engineer world has reshaped their practices; in a welcome spirit of user-centered design, they are beginning to build not around abstract categories and atomized individuals but real desired user groups, fleshed out with design research, global outreach networks, and user surveys (Collier 2020).

Conceiving of *privacy worlds* opens up fruitful new perspectives on classic Internet Freedom debates. Although this bears some similarity to Nissenbaum’s (2015) *contexts* of privacy, the privacy world allows for a richer and wider conception of discourses (within which Nissenbaum’s contextual social norms would be only a single aspect) to enter into these arrangements, and establishes a firmer link between the practices of both users and the designers and maintainers of privacy technologies. We argue that future research might profitably explore further privacy worlds, mapping their rich experiences, perspectives, and practices, and the wider stable formations of privacy that they represent. Using a privacy worlds approach to study differing constructions of privacy and how they are developed

suggests an approach to privacy research that retains much of the utility of *contextual integrity*, but provides an expanded framework for mapping the discourses, technologies, and practices that produce privacy.

Acknowledgments

We would like to thank the reviewers and editors for their extensive comments throughout the review process, and their contribution to this much transformed and improved version of the article. We would additionally like to thank Jamie Buchan and other colleagues who provided comments on the draft, particularly at STIS in Edinburgh and at the Cambridge Cybercrime Centre. Finally, we would like to thank the Tor community, especially those who contributed interviews or whose extensive mailing list discussions we have featured here.

Declaration of Conflicting Interests

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

Funding

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

ORCID iD

Ben Collier  <https://orcid.org/0000-0002-9207-3068>

Note

1. Although privacy and anonymity are related concepts (and often used interchangeably), we refer to privacy in this paper using Vedder's (2011) broad definition, as socially constructed category systems pertaining to the control of information about individuals, with spatial, relational, decisional, and informational dimensions, and deep links to social order, power, and culture. We distinguish this from anonymity, which (as befits privacy technologies like Tor) is the material affordance that allows an individual to control the knowability of their own identity in a given situation (Danezis and Guerses 2010). For the purposes of this paper, privacy is a predominantly social concept, linked to a range of cultural factors, while anonymity (as described by Tor's engineers) is a narrower material property concerned with the actual distribution of information about people (though both could undeniably be considered to be equally rich social constructions).

References

- Bancroft, A., and P. Scott Reid. 2017. "Challenging the Techno-politics of Anonymity: The Case of Cryptomarket Users." *Information, Communication & Society* 20 (4): 497-512.
- Becker, H. S. 1976. "Art Worlds and Social Types." *American Behavioral Scientist* 19 (6): 703-18.
- Bradshaw, S., and L. DeNardis 2019. "Privacy by Infrastructure: The Unresolved Case of the Domain Name System." *Policy & Internet* 11 (1): 16-36.
- Cavoukian, A. 2013. "Privacy by Design: Leadership, Methods, and Results." In *European Data Protection: Coming of Age*, edited by S. Gutwirth, R. Leenes, P. de Hert, and Y. Pouillet, 175-202. Dordrecht, the Netherlands: Springer.
- Chenou, J. M. 2014. "From Cyber-libertarianism to Neoliberalism: Internet Exceptionalism, Multi-stakeholderism, and the Institutionalisation of Internet Governance in the 1990s." *Globalizations* 11 (2): 205-23.
- Clarke, A. E. 2007. "Situational Analysis." In *The Blackwell Encyclopedia of Sociology*, edited by Ritzer, 1-2. New York: Blackwell.
- Clarke, A. E., and S. L. Star. 2008. "The Social Worlds Framework: A Theory/Methods Package." In *The Handbook of Science and Technology Studies*, vol. 3, edited by E. Hackett, O. Amsterdamska, M. Lynch, and J. Wajcman, 113-37. Cambridge, MA: MIT Press.
- Coleman, E. G., and A. Golub. 2008. "Hacker Practice: Moral Genres and the Cultural Articulation of Liberalism." *Anthropological Theory* 8 (3): 255-77.
- Collier, B. 2020. "The Power to Structure: Exploring Social Worlds of Privacy, Technology and Power in the Tor Project." *Information, Communication & Society* 1-17.
- Crabtree, A., P. Tolmie, and W. Knight. 2017. "Repacking 'Privacy' for a Networked World." *Computer Supported Cooperative Work: An International Journal* 26 (4-6): 453-88.
- Curran, J. 2012. "Rethinking Internet History." In *Misunderstanding the Internet*, edited by J. Curran, N. Fenton, and D. Freedman, 34-65. London, UK: Routledge.
- Danezis, G., and S. Gürses. 2010. "A Critical Review of 10 Years of Privacy Technology." *Proceedings of Surveillance Cultures: A Global Surveillance Society* 1-16.
- DeNardis, L. 2007. "A History of Internet Security." In *The History of Information Security*, edited by K. M. M. de Leeuw and J. Bergstra, 681-704. London, UK: Elsevier Science BV.
- DeNardis, L. 2009. *Protocol Politics. The Globalization of Internet Governance*. Cambridge, MA: MIT Press.

- DeNardis, L. 2012. "Hidden Levers of Internet Control." *Information, Communication & Society* 15 (5): 720-38.
- DeNardis, L., and F. Musiani. 2016. "Governance by Infrastructure." In *The Turn to Infrastructure in Internet Governance*, edited by F. Musiani, D. L. Cogburn, L. DeNardis, and N. S. Levinson, 3-21. New York: Palgrave Macmillan.
- Dencik, L., A. Hintz, and J. Cable 2016. "Towards Data Justice? The Ambiguity of Anti-Surveillance Resistance in Political Activism." *Big Data & Society* 3 (2): 2053951716679678.
- Dingledine, R., and N. Mathewson. 2006. "Anonymity Loves Company: Usability and the Network Effect." In *Workshop on the Economics of Information Security*, June.
- Dourish, P. 2003. "The Appropriation of Interactive Technologies: Some Lessons from Placeless Documents." *Computer Supported Cooperative Work* 12 (4), 465-90. doi: 10.1023/A:1026149119426.
- Fleck, J. 1994. "Learning by Trying: The Implementation of Configurational Technology." *Research Policy* 23 (6): 637-52. doi: 10.1016/0048-7333(94)90014-0.
- Gehl, R. W. 2018. "Archives for the Dark Web: A Field Guide for Study." In *Research Methods for the Digital Humanities*, edited by Lewis Levenberg, Tai Neilson, and David Rheams, 31-51. Cham, Switzerland: Palgrave Macmillan.
- Gray, J. 2018. "Three Aspects of Data Worlds." *Krisis: Journal for Contemporary Philosophy* 2018 (1): 5-17.
- Gueddana, W. H. 2013. "The Open Source Biography: A Multi-stage Approach." PhD Thesis.
- Gürses, S., A. Kundnani, and J. Van Hoboken. 2016. "Crypto and Empire: The Contradictions of Counter-surveillance Advocacy." *Media, Culture & Society* 38 (4): 576-90.
- Lewis, S. J. 2017. *Queer Privacy: Essays from the Margins of Society*. Vancouver, British Columbia: Mascherari Press.
- Lyon, D. 2014. "Surveillance, Snowden, and Big Data: Capacities, Consequences, Critique." *Big Data & Society* 1 (2): 2053951714541861.
- Maréchal, N. 2018. "Use Signal, Use Tor? The Political Economy of Digital Rights Technology." PhD Thesis.
- Milan, S. 2013. *Social Movements and their Technologies: Wiring Social Change*. Cham, Switzerland: Springer.
- Murdoch, S. J., and G. Danezis. 2005. "Low-cost Traffic Analysis of Tor." In *2005 IEEE Symposium on Security and Privacy (S&P'05)*, 183-95. Oakland, CA: IEEE.
- Musiani, F. 2010. "Privacy as Invisibility: Pervasive Surveillance and the Privatization of Peer-to-peer Systems." *TripleC* 9 (2): 126-40.

- Musiani, F. 2012. "Caring about the Plumbing: On the Importance of Architectures in Social Studies of (Peer-to-Peer) Technology." *Journal of Peer Production* 1 (online): 8-p.
- Musiani, F. 2013. "Network Architecture as Internet Governance." *Internet Policy Review* 2 (4): 1-9.
- Musiani, F. 2015. "Practice, Plurality, Performativity, and Plumbing: Internet Governance Research Meets Science and Technology Studies." *Science, Technology, & Human Values* 40 (2): 272-86.
- Musiani, F., D. L. Cogburn, L. DeNardis, and N. S. Levinson, eds. 2016. *The Turn to Infrastructure in Internet Governance*. London, UK: Springer.
- Nissenbaum, H. 2001. "How Computer Systems Embody Values." *IEEE Computer* 34 (3): 118-20.
- Nissenbaum, H. 2015. "Respecting Context to Protect Privacy: Why Meaning Matters." *Science and Engineering Ethics* 24 (3): 831-52.
- Pfitzmann, A., and M. Hansen. 2005. "Anonymity, Unlinkability, Unobservability, Pseudonymity, and Identity Management-A Consolidated Proposal for Terminology."
- Pollock, N., and R. Williams 2008. *Software and Organisations: The Biography of the Enterprise-Wide System or How SAP Conquered the World*. Routledge.
- Sorensen, K. H., and N. Levold. 1992. "Tacit Networks, Heterogeneous Engineers, and Embodied Technology." *Science, Technology, & Human Values* 17 (1): 13-35.
- Star, S. L., G. C. Bowker, and L. J. Neumann. 1998. *Transparency at Different Level of Scale: Convergence between Information Artifacts and Social Worlds*. Urbana-Champaign, IL: Library and Information Science.
- Star, S. L., and J. R. Griesemer. 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.
- Star, S. L., and K. Ruhleder. 1996. "Steps Toward an Ecology of Infrastructure: Design and Access for Large Information Spaces." *Information Systems Research* 7 (1): 111-34.
- Stewart, J. K., and R. Williams. 2005. "The Wrong Trousers? Beyond the Design Fallacy: Social Learning and the User (2005)." In *User Involvement in Innovation Processes. Strategies and Limitations from a Socio-technical Perspective*, edited by Harald Rohrer, 9-35. Munich, Germany: Profil-Verlag.
- Strauss, A. 1982. "Social Worlds and Legitimation Processes." *Studies in Symbolic Interaction* 4 (17): 171-90.
- Syverson, P. 2009. "Why I'm Not an Entropist." In *International Workshop on Security Protocols*, April, 213-30. Berlin, Germany: Springer.

- Syverson, P., R. Dingleline, and N. Mathewson. 2004. "Tor: The Second-generation Onion Router." In *Usenix Security*.
- Vedder, A. 2011. *Privacy 3.0. In Innovating Government*, 17-28. TMC Asser Press.
- Unruh, D. R. 1980. "The Nature of Social Worlds." *Pacific Sociological Review* 23 (3): 271-96.
- Zalnierute, M., and S. Milan. 2019. "Internet Architecture and Human Rights: Beyond the Human Rights Gap—Policy & Internet." *Special Issue: Internet Architecture and Human Rights* 11 (1): 6-15.

Author Biographies

Ben Collier is a lecturer in Digital Methods at the University of Edinburgh. He researches digital infrastructure and how it becomes implicated in crime, power, and control, using perspectives and methods from science and technology studies and criminology.

James Stewart is a lecturer and research fellow at the University of Edinburgh in the Institute for Science, Technology, and Innovation. He works in the field of technology studies, concentrating on the appropriation, consumption and use of new ICTs and the co-evolution of large-scale systems of technology and culture.