Patching security governance: an empirical view of emergent governance mechanisms for cybersecurity

Michel van Eeten

Michel van Eeten is based at Delft University of Technology, Faculty of Technology, Policy and Management in The Netherlands.

Abstract

Purpose – The issue of cybersecurity has been cast as the focal point of a fight between two conflicting governance models: the nation-state model of national security and the global governance model of multi-stakeholder collaboration, as seen in forums like IGF, IETF, ICANN, etc. There is a strange disconnect, however, between this supposed fight and the actual control over cybersecurity "on the ground". This paper aims to reconnect discourse and control via a property rights approach, where control is located first and foremost in ownership.

Design/methodology/approach – This paper first conceptualizes current governance mechanisms through ownership and property rights. These concepts locate control over internet resources. They also help us understand ongoing shifts in control. Such shifts in governance are actually happening, security governance is being patched left and right, but these arrangements bear little resemblance to either the national security model of states or the global model of multi-stakeholder collaboration. With the conceptualization in hand, the paper then presents case studies of governance that have emerged around specific security externalities.

Findings – While not all mechanisms are equally effective, in each of the studied areas, the author found evidence of private actors partially internalizing the externalities, mostly on a voluntary basis and through network governance mechanisms. No one thinks that this is enough, but it is a starting point. Future research is needed to identify how these mechanisms can be extended or supplemented to further improve the governance of cybersecurity.

Originality/value – This paper bridges together the disconnected research communities on governance and (technical) cybersecurity.

Keywords Internet, Governance, Data security

Paper type Research paper

1. Main paper

In recent years, cybersecurity has been framed as one of "the most important areas" of the broader "global war" around internet governance (De Nardis, 2014, p. 88). Mueller (2017) characterized it as a battle between two conflicting governance models: the nation-state model of national security and the global governance model of multi-stakeholder collaboration, as codified in global and transnational institutions like IGF, IETF, ICANN, etc. Other authors also observe processes of "securitization" and "militarization" at work (Hurel, 2017). Cybersecurity threats are interpreted in terms of foreign policy and military conflict, which enable state actors to encroach on cooperative forms of global governance. To Mueller (2017), the danger is clear: "the equation of cybersecurity with state responsibility and national security/military responsibilities means that a large chunk of global Internet policy making is in danger of being pushed out of the open, multistakeholder model".

There is a strange disconnect, however, between this so-called battle and the actual provisioning of cybersecurity "on the ground". While the trend of militarization can be

clearly observed in the policy discourse on governance, the same cannot be said of the actual governance of the networks, systems and services that make up the internet. Remarkably little has changed in the past decade, with one important exception: offensive operations by nation states. Headlines alone are enough to see this shift. The Snowden revelations have exposed a plethora of offensive operations by the USA; the attacks on Sony, the Bangladesh Central Bank and the Wannacry outbreak have been attributed to North Korea; the Chinese intelligence services have been observed hacking NGOs; and the attack on the DNC during the US elections has been attributed to Russia. States are now widely seen as one of the most dangerous threat actors.

Offense is not cybersecurity, however – it is the opposite: cyber-insecurity. On the side of defense, we have not seen changes of even remotely similar magnitude. In the West, at least, the actual policies of states to protect the internet – or perhaps, one should say, "their" internet – are still remarkably hands off. National cybersecurity strategies still rely heavily on voluntary action of the private actors that operate the networks, systems and services that are loosely summarized as "the Internet". Public-private partnerships have been the go-to phrase for at least a decade now. Around this core tenet, a few new developments have emerged, such as intelligence agencies sharing information with selected private companies, mandatory breach notification laws, centralized information sharing mechanisms and sectoral regulation in health and finance requiring the adoption of basic security standards. These do not qualify as a proof of militarization or securitization in the governance landscape, however.

Beyond actual changes, there even seems to be a dearth of concrete policy proposals in the direction of increased state control. To illustrate: US President Trump, who has come to power on a decidedly nationalist platform, has not put forward an alternate, more state-centric approach. His recent "Presidential Executive Order on Strengthening the Cybersecurity of Federal Networks and Critical Infrastructure" has been widely characterized as a continuation this status quo. Also note the rather narrow scope articulated in the title: federal networks and critical infrastructure. That puts out of scope most of what is commonly understood as the internet. Few people would argue that these two areas – the government's own systems and critical infrastructures like energy and transportation – are not within the purview of state regulation.

What about the earlier draft of Trump's executive order, though? Was that not evidence of a push for militarization or more state control? It is not clear. Yes, the language seemed to provide legitimacy for a larger role of the state. But legitimacy for a role is not the same as instrumenting that role, let alone actually performing the role. This is more than an academic distinction. Look at the discrepancy between the alarming description of the security threats faced by the USA and the actual actions proposed in the earlier draft order: a series of reviews.

In short, there is remarkably little evidence of a government takeover of internet governance in the name of security. In this light, the so-called "global war" or "battle" over cybersecurity governance looks rather confusing. What is it really about? To a certain extent, it seems to confuse talking about governance with actual governance.

This paper has a threefold aim. First, it distinguishes between governance as discourse and governance as control. The "war" is taking place at the level of discourse, which is at best loosely coupled to actual control over resources, systems and services. Second, the paper asks: how is cybersecurity governance actually changing? It defines governance as policy-driven control over internet resources, systems and services. As control emanates, first and foremost, from ownership, we can rephrase the question to: how are property rights changing because of security threats, and vice versa? Third, the paper surveys a range of empirical case studies to study how the institutional landscape of security

governance is "patched" to deal with emerging threats. We end with summarizing the main findings and discussing their implications for improving cybersecurity governance.

2. Discourse vs control

We start by observing a simple distinction: talking about governance is not the same as actual governance. Governance – even in its earliest roots – is intrinsically tied to control, steering, influence. Governance without actual control is basically discourse about governance.

This should not be a controversial point, and yet this distinction seems to be lost on many authors in political science and international relations writing about governance and the internet. Their studies focus on the same basic set of institutions, an alphabet soup of acronyms well known to anyone with a passing familiarity of the field: ICANN, IETF, IGF, WSIS, ITU, GCCS, etc. The work of these institutions is what is typically referred to as internet governance.

To understand how internet governance, and the work of these institutions, is related to the governance of cybersecurity, we first need to better understand to what extent they provide governance in the first place. The question, and also the title of an earlier paper (Van Eeten and Mueller, 2013), is: Where is the governance? Take the Internet Governance Forum (IGF). Its mandate focuses explicitly on governance, the conversations taking place within this institution focus on governance, and yet it has basically no impact on the actual functioning of the internet. In other words: all discourse, no control. There is no institutional structure to translate any of the outcomes into something that shapes the behavior the actors operating the resources, systems and services of the internet. In fact, IGF does not even have outcomes in terms of resolutions or policies.

We can go through the same exercise for the other institutions associated with governance. The Global Conference on Cyberspace (GCCS) faces basically the same disconnect as the IGF. Regarding the IETF, some authors have written about its "decision making authority", which sounds like it is doing governance (Mueller, 2010, p. 4). These decisions, however, only refer to its own internal procedures and publication of standards, or rather: RFCs (Request for Comments). What is absent from these analyses is the critical fact that many RFCs are ignored by the actors meant to use them: the operators of the systems that the standard applies to. Standards that are ignored do not constitute governance. Crudely put: anyone can make a standard. Even those RFCs that are not ignored have a hard time getting any traction. DNSSEC is over 20 years old, has been evangelized for at least a decade and yet adoption at the level where it matters, namely, for second-level domains and below, is still at a low rate. In short, yes, the IETF has influence in certain areas, but many of its results are not widely adopted, and it certainly does not have decision-making power over other actors. The degree of influence, and hence governance provided by the IETF, is in the end an empirical question. The sheer process of manufacturing standards should, however, not be equated with governance.

What about ICANN? It has been a key recipient of the stakeholder views articulated at IGF. Well, tied to ICANNs' corporate governance are its own versions of these multi-stakeholder structures and processes, populated, in fact, by many of the same people. These conversations do have consequences for ICANN. Does ICANN control internet resources, systems and services? Yes it does, only in a limited sense. Some people seem to think that ICANN controls the DNS infrastructure, but this is, in fact, not the case. DNS is a highly distributed system, and each owner of a machine in that system is basically self-governing. To illustrate: this distributed control is why the internet has been plagued for years by millions of open DNS resolvers (a recent count came to a total of 28 million), which are used in so-called amplification distributed denial-of-service (DDoS) attacks[1]. There is nothing ICANN can do about this. The essence of the governance function of ICANN lies in
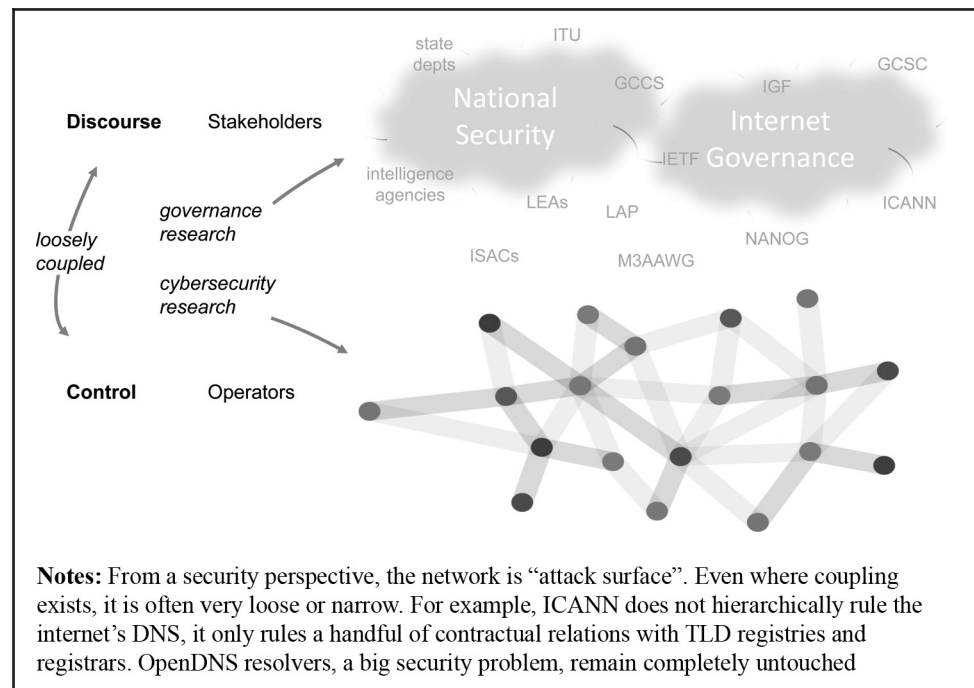
maintaining and adapting a number of contractual relationships with registries and registrars. To be sure, there are important issues at stake in those relationships, as evidenced in the recent IANA transition. It would be wrong, however, to see ICANN as governing DNS. In other words, ICANN manages a very small part of what constitutes the global DNS infrastructure, which, in turn, is a very small part of the actual governance structures of the internet.

It seems fair to conclude that there are only very limited forms of governance going on in the institutions where the "war" over internet governance is raging. The battle between the two governance models – nation states vs global, multi-stakeholder institutions – is a battle over who gets to shape the discourse of internet governance. Observing that discourse, one might argue that securitization is taking place, although a counterargument would be that some of these institutions avoid focusing on security or focus on it only in narrow ways. The position one takes on this question, however, does not negate the fact that the discourse is only very loosely coupled to actual governance, which is located in providers and other actors operating the resources, systems and services of the internet. Figure 1 tries to visualize this point.

At the level of control, very different actors are at play. Depending on what systems and networks one looks at, these actors can number in the thousands or tens of thousands. For example, the market of hosting providers has been estimated to consist of around 45,000 providers (Tajalizadehkhoob *et al.*, 2016). The overwhelming majority of these actors are absent from the processes higher up in the figure. At IFG, you might encounter representatives of larger firms that dominate certain markets, such as Google or Comcast or Microsoft. These representatives tend to be policy or public affairs staff, not the people who actually shape the companies' business policies or run their operations. The language reflects this disconnect: discourse typically refers to stakeholders, whereas at the level of control, people talk about operators and providers.

As far as security is concerned, academia is similarly decoupled. Governance research is heavily focused on the top of the figure, cybersecurity research on the bottom. When you

Figure 1 Governance as discourse vs control

Notes: From a security perspective, the network is "attack surface". Even where coupling exists, it is often very loose or narrow. For example, ICANN does not hierarchically rule the internet's DNS, it only rules a handful of contractual relations with TLD registries and registrars. OpenDNS resolvers, a big security problem, remain completely untouched

search for the term "governance" in the top four academic cybersecurity venues (USENIX Security, CCS, NDSS, IEEE Security and Privacy), you come up empty handed. If policy is mentioned, it is often a rather naïve extrapolation of technical findings, arguing, for example, that a certain technical solution should be adopted by providers or mandated by governments, without an understanding or analysis of the institutional framework within which this is supposed to happen.

Does this mean that discourse-as-governance is irrelevant? Or, to borrow Zittrain's (2008) characterization of the IGF, is there more to it than an array of talkshops? Well, talkshops are not *per se* without influence. Discourse might give rise to norms that shape the behavior of a wider range of actors who do have some kind of control. States could theoretically legislate themselves into more control. (Whether the kind of control that can be legislated is the kind of control that can bring about the desired effects is rather doubtful, but this is another discussion.)

This legitimizing effect of the securitization discourse might be the motivation for the rather dramatic tone in academic publications and white papers – which ironically expresses concerns over securitization by using metaphors of war and battle. Historically, one might argue, legitimacy for state involvement is correlated with actual state involvement.

To put it differently, the fact that discourse and control are only loosely coupled reflects a particular political economy, where many states have not imposed wide-ranging hierarchical control, which leaves a space for distributed, private forms of governance. This particular arrangement is contingent and might change over time, as authoritarian regimes have shown us. Looking at regimes in the West, we cannot discard the scenario that increased state legitimacy at the level of discourse might result in institutional changes to increase state control. As said above, discourse might shape ideas and norms which, in term, influence might lead to an instrumentation of a more dominant role of state. At its core, this is an empirical question that we can only answer over time. At this moment, the claim that control is really shifting in the direction of states does not seem supported by the evidence that is being presented.

3. Changes in security governance

How is cybersecurity governance actually changing? Earlier, we defined governance as policy-driven control over internet resources, systems and services. Where is this control located? Anyone studying how security threats are handled in practice will soon observe that decisions about the protection of a resource are closely tied to the ownership of that resource. Notwithstanding concepts that frame the internet as a "global digital commons" or "public good", the fact is that nearly every resource, system or service is someone's private property. There are some complicated legal "edge cases" around ownership of, for example, IP addresses, but these are minor exceptions. What risks are mitigated, accepted or externalized is driven first and foremost by the incentives of the owner. Institutional mechanisms might constrain or otherwise shape these incentives.

One could basically call this a property rights approach to security governance. There is nothing particularly special about the approach. It has been applied to a wide array of societal challenges and governance issues. Ownership is a conceptually straightforward starting point for thinking about governance. Nevertheless, in the area cybersecurity governance, it immediately brings into focus just why governance is so complicated, or in some ways, absent: ownership is extremely distributed across an interdependent global ecosystem of resources, systems and services.

From the beginning, the concept of a "dumb" network with intelligence located in the nodes around the edge has been an important design principle for the protocol stack of the internet and its predecessors (Fidler, 2017). This also meant that a lot of decision-making power has been located in those nodes at the edge of the network. Precisely, this

consequence is directly tied to the innovative power that we have come to associate with the internet. The general-purpose computing devices at the nodes have provided a platform for rapid innovation in all kinds of directions, virtually unbridled by gatekeepers elsewhere in the system. What Zittrain (2008) called "generative" computing was coupled with what has recent been referred to as "permissionless innovation" (Thierer, 2014, 2016). The nodes and the network were designed to accept any contribution that followed a basic set of rules – i.e. the functional requirements of an operating system or the protocols of the internet.

A lot of the security threats that have emerged over the past two decades are intimately tied to this design. The fact that nodes can run arbitrary code also allows them to be used or subverted by malicious actors. No one will preemptively stop them; only after malicious activity is detected, are countermeasures sometimes taken. For the attacker, it is not difficult to simply deploy other nodes and keep going. This is why, many security practices on the internet are characterized as whack-a-mole.

The network had not been designed to handle malicious traffic different from benign traffic, or to even tell the difference. As Fidler's (2017) revisionist history points out, this was not a naive choice by a small, trust-based community, but the conscious outcome of national security policies. The defense and intelligence agencies involved in funding and developing the early internet operated under a different threat model: interception and surveillance of traffic in the network (e.g. by other nation states), not subversion of nodes. Over time, some capabilities to mitigate the threats of misbehaving nodes have been tacked onto the basic design of the network itself. Think of blocked ports, anti-spoofing filtering, sinkholing of botnet command-and-control traffic, blackholing of DDoS traffic, etc. These measures did not fundamentally change the status of the network as being a "dumb" network. Authority, also in terms of security, still resided first and foremost with the owners of the nodes.

A key challenge that cybersecurity researchers run into all the time is that they cannot identify a contact point for a device or resources, let alone identify the owner. To illustrate: in a recent study, we tried to notify the owners of approximately 4,500 authoritative nameservers that their server and the associated domain names were vulnerable for a so-called zone poisoning attack (Cetin *et al.*, 2017; Korczyński *et al.*, 2016). We followed IETF's RFC 2142, which specifies how to identify a contact point for DNS servers. The result was depressing: 85 per cent of our notification messages bounced. Simply put, the recipient that the RFC told us to contact did not exist. Needless to say, establishing ownership is even harder, a lot harder actually, than finding points of contact. Even if one could identify the owners, a governance mechanism would then have to scale over, in this case, 117 countries. Of the contact points we did manage to reach, only a fraction remediated the vulnerability. In an earlier project (Cetin *et al.*, 2015), we contacted network operators who were harboring not just a server with a serious vulnerability, but one running malicious code engaged in malware attacks on end users. The cleanup rate here was better: around half of the contacted operators managed to remediate the malicious servers. The other half either did not get the message, ignored it or was otherwise unable to engage with the owner of the server to get it cleaned up.

If control emanates, first and foremost, from ownership, then we can rephrase the question of how governance is changing to: how are property rights changing because of security threats, and vice versa? This is a high-level question that can either be answered in a high-level fashion or by highly granular empirical analysis in a very specific context. The latter is more precise, but it has to go patch by patch of the space called cybersecurity governance. In the context of this paper, we follow the former approach. We explore four trends in terms of how property rights are changing: shifting property rights of resources, systems and services from users to intermediaries; limitations on the property rights of

owners; limitations on property rights of intermediaries; and limitations on property rights of manufacturers.

3.1 Shifting property rights from users to internet intermediaries

Schneier (2012, 2013) has observed that two recent developments are impacting the authority of owners of nodes: the rise of cloud computing and vendor-controlled platforms. The first trend means that more of our data and computing takes place on the networks of others, rather than on our own node. Obvious examples are Gmail, SalesForce, Amazon Elastic Cloud Compute, Facebook, Uber, Spotify, Office 365, Dropbox, etc. The second trend means that more and more of our devices are closed down, or at least less open than general-purpose computers, and controlled by vendors. Think of Apple's iOS or Google's Android operating systems for mobile devices. Both vendors limit what users can do with their devices, i.e. what code they can run, by restricting them to a curated app store with apps that run under limited privileges on the device. Google's model is a bit more open than Apple's, as it allows the use of alternate app stores and the side-loading of apps. Apple actively prevents this by legally fighting users who find ways to circumvent these restrictions on the devices that they own. This difference between the two vendors is important, but it is a gradual one.

These two trends are not just technological advances, they are propelled by powerful economic forces. Most notably: cost. Cloud computing offers dramatic efficiency gains compared to consumers and businesses owning and maintaining their own computing infrastructure. In addition to cost, there is also increased reliability, convenience and yes, security. Google is more competent in securing the Gmail platform than most corporations are in securing their mail servers. Closed operating systems are also easier to defend against compromise, as it is more difficult to exploit vulnerabilities or to trick users to install malicious apps, as these are constantly checked and removed from the app stores.

In other words, these developments are not intrinsically wrong. We derive a lot of benefits from them. But they do mean that power has shifted and will continue to shift from the owners of devices to the cloud operators and platform and device vendors. The OECD (2010, 2011) has referred to these companies as internet intermediaries. Their security practices determine to an increasing degree the security of everyone. And, in many cases, we do not really have a good way to evaluate what they are doing. Cloud operators typically do not allow you to audit their systems and service. There is a fundamental information asymmetry that comes into play.

Schneier (2013) has characterized our increasing dependence for security on these intermediaries as a "feudal relationship". We trade away our some of our sovereignty to these companies and get security in return. Another way to say this is: certain the property rights that we used to associate with the device itself are now tied to the services of cloud and platform operators, and thus are being transferred to those operators. The device itself becoming less important for security governance if my data and computing tasks are no longer performed by that device. In a way, the ownership and associated property rights of data and computing are migrating from the nodes that the user owns to those of the intermediaries.

These intermediary markets are often concentrated in one or a small number of providers. They operate in one- or two-sided markets with positive externalities, which give rise to a "winner takes all" dynamic. These large, concentrated firms offer new control points for governments, as they are much more within the grasp of laws and other instruments of state power than the heavily distributed set of autonomous nodes ever were. We have seen how the interests of corporate and government power have been aligning. A significant number of the so-called internet giants rely on surveillance as much as governments do. Snowden made visible the extent to which NSA was using Google, Facebook, Verizon and others to

get access to data it could not otherwise. An example in the other direction is how the entertainment industry is recruiting governments to enforce their business models.

We do not know exactly how our sovereignty is being eroded or what the security and privacy tradeoffs are that we are entering into. It is hard enough to assess what Facebook is doing with user data. What about the many derivative applications that we are not even aware of? Think of the legal use of Facebook data by firms like Cambridge Analytics, which combines Facebook data with other data of users' online and offline behavior to enable micro-targeted communication for the political campaigns behind Brexit and Trump (Cadwalladr, 2017). Beyond the derivate business models around large intermediaries, there are firms tracking our behavior in ways that are unknown to most of us and whose business models are a mystery (Englehardt and Narayanan, 2016).

In sum, the governance of security has shifted from device owners to large intermediaries, toward a model that Schneier has dubbed feudal security. This shift is generally associated with certain security benefits, most notably more centralized control over the security of devices and services for millions or even billions of users, and economies of scale in terms of concentrating security expertise and resources in a small number of companies that serve many users, rather than users having to fend for themselves. There are also security tradeoffs, however, that are very difficult to evaluate at this moment because of information asymmetries. We currently have no good answer for this problem.

3.2 Limitations on property rights of owners

Next to transfer of property rights from device owners to intermediaries, another trend is the growing number of regulatory constraints on the property rights of device owners. This is mainly happening in sectors that were already strongly institutionalized and regulated, such as health, energy, financial services and transportation.

Information and communication technology has been permeating more and more processes in these regulated sectors, initially without much guidance in terms of standards or regulatory guidelines. Some of the security incidents over the past decade and a half reveal the presence, and failure, of commercial of-the-shelf products in these environments. In 2003, for example, the Slammer worm infected Microsoft Windows systems in the Davis-Besse nuclear power plant in Ohio, disabling a safety monitoring system for nearly five hours. In healthcare, many medical devices have been found to be running outdated or even deprecated operating systems. A ransomware attack in May 2017 crippled operations in several UK hospitals. There has been a range of malware families that target ATMs running some version of Windows.

It has taken a long time for regulatory mechanisms in these sectors to catch up, mostly because of the complexity and speed of the changes. Slowly but surely, though, security standards are being recommended or mandated in these sectors. Many of these standards are process-based ("adopt adequate safeguards"), rather than mandating specific technical security measures. A baseline standard that is being adopted across many sectors is ISO 27001. In the European Union, the directive on security of network and information systems (the NIS Directive) identifies "essential service operators" within the energy, transport, banking, financial market infrastructure, health, drinking water and digital infrastructure sectors and requires them to adopt security precautions and mandatory breach notifications.

Beyond these generic frameworks, there are many sector-specific standards. The US Health Insurance Portability and Accountability Act (HIPAA) of 1996 has a security rule that requires certain controls to be put in place. The North American Electricity Corporation (NERC) has adopted mandatory patching policies for operators. Financial services have had the PCI DSS standard in place for a long time already, but its norms are updated more

quickly in light of new threats, and they are also being extended to other players in the value chain, such as third-party service providers.

While these regulations, in principle, constrain the property rights of owners in terms of mandating certain precautions, it is unclear just how much impact they have. These standards are typically articulated quite abstract (e.g. deploy access controls to ensure confidentiality of data) or process-oriented (e.g. allocate responsibility for information security in a Chief Information Security Officer role). Flagrant breaches of payment systems at large US retailers have occurred even though they were PCI DSS compliant, to name one example. On the other hand, even non-binding guidance, such as the US Food and Drug Administration's (FDA) "postmarket management of cybersecurity for medical devices" might impact security practices, under the threat of law suits by patients whose care was impeded because of a medical device that did not meet best practice standards. Beyond security precautions, there are other regulations in place, such as mandatory breach reporting to various institutions, requirements for sharing data with the government or rules on fraud reimbursement. These regulations, however, do not directly affect the property rights associated with the IT infrastructure that the organizations own.

The importance of sector-specific frameworks is likely to increase over the next years, as the proliferation of the so-called Internet-of-Things (IoT) takes off. IoT, pervasive computing, ubiquitous computing, these terms signal the disappearance of the distinction between devices with and without connectivity and computing capabilities. Without that distinction, it also becomes less meaningful to think about cybersecurity governance as a space with a certain structural coherence. A recent outlook on the governance of IoT suggested that, in many sectors, it will merge with the well-institutionalized framework of safety – e.g. car safety, medical device safety and toy safety. Some of the obligations fall onto the owner, some onto the vendor of the device. (We discuss the latter point in a moment).

3.3 Limitations on property rights of intermediaries

In many democratic countries, legal frameworks for online commerce and services have shielded intermediaries from liability for what happens on their networks and platforms. In the EU, for example, the Electronic Commerce Directive, adopted in 2000, declared that transmission ("mere conduit"), caching and hosting services were not liable for the information they handled, as long as the provider does not have actual knowledge of illegal activity or information and as long as, upon obtaining such knowledge, acts to remove or to disable access to the information ("notice and take down"). The USA has a similar framework with the so-called "safe harbor" provisions in the Communications Decency Act and the Digital Millennium Copyright Act. While these liability exemptions are still in place, intermediaries are regularly under pressure to be more proactive in dealing with hate speech, copyright infringement, anti-terrorism, protection of minors.

In principle, security incidents, such as compromised websites or spam sources, can also be viewed within this framework. While their legal liability is limited, intermediaries have acted more frequently against security threats over time, when they perceived this as being in their interest or part of their responsibility – i.e. to protect their services or users.

As intermediaries have become more dominant, the calls to impose more obligations on them have also intensified, including in the area of security. Some of these calls have helped shape new legislation. The EU NIS directive has articulated security requirements for "digital service providers", which are defined as legal persons that provide "service normally provided for remuneration, at a distance, by electronic means and at the individual request of a recipient of services" (European Commission, 2017). The directive then specifies this as applying to three types of intermediaries: "online marketplace", "online search engine" and "cloud computing services". They are subject to the same requirements as the "essential service operators" discussed above: taking adequate security precautions and mandatory breach notifications. These requirements overlap with

the more narrow set of obligations flowing from the EU General Data Protection Regulation, which focuses on the protection of personal data.

Another area of contention and regulation has been around the issue of "net neutrality", which can be crudely summarized as treating all traffic data as equal. While this issue is not about security, neutrality rules do touch on it, as many security practices of network operators by necessity discriminate good from bad traffic. These practices required exemptions to strict neutrality, to remain legal.

In terms of security, intermediaries' property rights have also been shaped by softer mechanisms than formal regulation. For example, in many countries, public-private partnerships have emerged in which ISPs agree to block spam from leaving their network or to mitigate botnets by contacting and quarantining customers whose devices are infected with malware. The additional responsibility is sometimes tied to some form of "duty of care", though the legal basis to subsume security under this duty and what the duty exactly demands of the intermediary are not always clear.

In short: the increasing dominance of intermediaries appears to go hand in hand with more restrictions on their property rights in terms of what they should or should not do. New regulations have required security precautions to be taken. Other security practices, such as botnet mitigation by ISPs, seem the result of softer forms of government pressure and appeals to "duty of care". That being said, security requirements are limited and liability shields still seem firmly in place. Beyond imposing obligations, where security is aligned with their business interests, intermediaries undertake security measures without external pressure or limitations on their property rights. Facebook, for example, exerts quite a lot of effort to protect its users from malicious content on the platform – e.g. links to phishing or drive-by-download sites.

3.4 Limitations on property rights of vendors

The last shift in security governance to highlight here is changes in the property rights of vendors. Conventionally, software and hardware vendors put their products into the market without requirements in terms of how they were secured. Vendors could offer security features and even compete on them, of course, as part of the freedom of contract. In practice, this freedom has meant that the software industry disclaimed liability for the harm when its products fail. Users, whether corporate or consumer, have to accept End User License Agreements (EULAs) to be able to use the product. This is not without benefits in terms of innovation ("go fast and break things"), but the downside is that time-to-market and other economic incentives have often trumped security. The vendor bears few, if any, of the damages that result from security issues.

In a recent court case, the Dutch consumer union (Consumentenbond) took Samsung to court for failing to release security patches for even recent phones. The unwillingness to release such patches is somewhat startling. Most of Samsung's phones run on Android. Google diligently patches discovered vulnerabilities and then releases these patches to all manufacturers using Android. So, the cost to Samsung of passing these patches onto their customers is quite limited. It would need to do some quality assurance testing to prevent adverse effects of the patch on Samsung's own software, which it bundles on top of Android. Still, even this limited effort is apparently too much to ask. The consumer union demanded that patches should be released for at least two years after the phone is purchased, which is in line with the normal period for product warranty. In the first case, the court ruled against the consumer union on procedural grounds. In November 2016, the union initiated a new case that will receive a more in-depth treatment from the court.

An interesting counterpoint to this case is another case that also involves Samsung and a software update – and it points to important changes in security governance. In August 2016, the company released a new flagship phone: the Note 7. Over the next few months,

serious problems emerged as some of the batteries overheated, caught fire or exploded. In the end, the company had to issue a total recall. Not all owners of the Note 7 were willing to return their device, however. In response, Samsung rolled out a software update that prevented the phone from charging, rendering it completely unusable.

Note the stark contrast between the lackluster attitude to updating software for security vulnerabilities and the aggressive approach to updating once Samsung was working within the product liability framework. Why is this pointing to a shift in governance? Because the pervasive presence of software in all kinds of products – also known as IoT – will mean that cybersecurity will increasingly become subsumed into product liability. Another clear illustration is Fiat Chrysler's two recent recalls for software vulnerabilities in its cars. No accident had occurred, nor had the vulnerability been exploited by criminals, as far as we know. Its mere presence was enough to trigger a recall. The most recent one took place in May 2017 and affected 1.3 million cars. Almost at the same time, the FDA sent a scathing letter to St Jude Medical, a manufacturer of pacemakers, faulting it for knowing about grave security issues in its implantable medical devices as early as 2014, but failing to address them with software updates or by replacing those devices.

The product liability regime is much better institutionalized and provides very strong incentives for vendors. It places constraints, sometimes severe ones, on their property rights. One-sided EULAs disclaiming liability for damages, which users have been habituated to click though because they have no meaningful other option, are overruled by the laws and jurisprudence of product liability.

Of course, not all computing devices or software vulnerabilities will fit within the product liability framework. In fact, this framework has been in place all along, but it did not fit the type of damages associated with security flaws: typically small economic damages, rather than fatalities or injuries, which are distributed over many parties, sometimes millions of them, across many jurisdictions. For those types of security issues, other mechanisms are emerging. In 2015, the US Federal Trade Commission filed a complaint against ASUSTeK Computer that alleges that the company did not take reasonable steps to secure the software on its routers. Researchers had uncovered a variety of vulnerabilities in these devices. As per normal industry practice, ASUS did not address these security flaws in a timely manner, nor did it notify consumers about the risks posed by the vulnerable routers. In the end, the case was settled out of court. As part of the settlement, ASUS is required to establish and maintain a comprehensive security program subject to independent audits for the next 20 years.

ASUS is just one example in a sheer endless list of vendors with similar practices. Most of these firms we have never heard of, yet their products enter markets and networks across the world. In September 2016, the Mirai botnet captured over one million of such poorly secured IoT devices; routers, webcams, digital video recorders, IP cameras and many unknown types of "things". Mirai was used to conduct the largest DDoS attack ever recorded, reaching a peak of over 1 Tb per second during an attack on Dyn, a DNS service provider for Netflix, Twitter, CNN and many other companies, disrupting the services of all these organizations.

Researchers and the security journalist Brian Krebs started calling out vendors whose devices were observed to be compromised and participating in the attacks. Normally, such naming and shaming campaigns have little effect. But the sheer size and headline-grabbing character of the attack made it different this time. In response, a Chinese vendor of network cameras named XiongMai vowed to initiate a product recall. (At the same time, it was threatening legal action against the accusers.) It also said it had changed the default settings of the camera to turn off telnet – a key attack vector – and ensure that users changed the factory-default passwords.

Another development that Mirai triggered is that the European Commission is preparing new legislation to protect IoT devices from security threats. It is looking at the certification and labeling systems for approved secure devices (Krebs, 2016). This could mean that customs will stop devices from entering the EU market that do not have the required certification.

In sum: a number of forces are converging that are reshaping the governance of security and constraining the property rights of vendors. It will require many vendors then to put in place at least basic security practices, such as patching known vulnerabilities, and expose many others to product liability regimes and certification requirements tied to import barriers.

4. Case studies: emerging governance mechanisms for security externalities

In the previous section, we surveyed high-level shifts in cybersecurity governance. Complementing this analysis, we now turn to several small case studies of governance mechanisms that have emerged around negative externalities. Externalities are the effects of economic transactions on third parties. Typical examples are network effects (a positive externality, where a service becomes more valuable to a user as other users join the service) and environmental pollution (a negative externality, where the cost of pollution is borne by society). Many important security threats are examples of negative externalities: compromised nodes (hacked end-user machines, websites, IoT devices) that are used to support attacks on third parties, rather than the owner of the machine.

From a governance perspective, security externalities are a crucial, if not *the* crucial, challenge. As long as owners of resources, systems or networks incur most of the cost and benefits of their own security decisions, these decisions will be aligned with social cost and benefits. In other words, the current situation, where highly distributed ownership also means highly distributed security governance, could work well enough in terms of producing socially optimal outcomes as long as the security decisions of the owners affect primarily themselves. When they invest too little in security, they suffer damages which will incentivize them to invest more. Through a variety of signals – e.g. trial and error, herding behavior or insurance pricing – they will converge on security practices that reflect a rational tradeoff of cost and benefits.

When the cost of security failures by owners end up with third parties, however, then the security incentives of owners are eroded. They invest too little in security or otherwise adopt practices that are too risk-seeking. These are the situations where governance mechanisms are necessary to realign the incentives with the social cost and benefit of security. In the remainder of this section, we explore three areas (botnets, compromised websites and spoofed traffic) where a variety of mechanisms have emerged that try to mitigate the externality.

The case studies are based on a variety of research projects we have been involved in, supplemented with studies from other researchers in these areas. We briefly introduce each threat, then go through the different mechanisms and end with the evidence that is available as to the effectiveness of the mechanism. In line with our discussion at the start of this paper: does the mechanism actually influence the operational process where control is located?

4.1 Mitigating botnets

The global malware outbreaks of the early 2000s, like the ILOVEYOU and CODE RED computer worms, were disruptive and highly visible. Their authors seemed motivated mainly by the quest for notoriety. Then, as more economic transactions moved online and the cost of abusing vulnerabilities decreased, profit-driven criminals entered the scene and rapidly expanded their activities. Their incentives changed malware from

visible and disruptive to stealthy code that kept the machine of the victim up and running as part of a criminal infrastructure with a command and control mechanism: a network of robots, a.k.a. botnet. Criminals discovered an expanding array of business models to monetize these infected machines: sending spam, extorting ransom, performing DDoS attacks, harvesting user credentials, selling fake anti-virus software, committing financial fraud, hosting phishing sites, performing click fraud against advertising networks and more.

Most of the attacks are targeting third parties. This undermines the incentives of end users to protect their machines from malware, as they do not bear the full cost of the infections. As end users were not keeping up, the ISPs have faced increasing pressure to become more active in mitigation. A variety of national anti-botnet initiatives have emerged, as well as other voluntary mechanisms, in which ISPs commit to contacting and remediating infected customers.

To evaluate the effectiveness of different mitigation strategies, we developed an approach to generate comparative infection rates per subscriber for ISPs in over 60 countries, drawing on a variety of data sources on infected machines (Asghari *et al.*, 2015a). We found out that the 262 ISPs in our analysis harbor over 70 per cent of all infected machines worldwide. We also observed widely different infection rates among ISPs, as well as among countries, signaling differences in the effectiveness of mitigation. With those metrics in hand, we looked at different governance mechanisms.

4.1.1 National anti-botnet initiatives. Around five years ago, national anti-botnet initiatives have emerged in seven countries. We compared infection peaks and cleanup rates for those countries using sinkhole data of the Conficker botnet, a botnet that was sinkholed (i.e. the command and control infrastructure was taken over by defenders) and abandoned by the attackers, many of whom were actually apprehended by law enforcement agencies. This offered a unique view into cleanup efforts, as there were few interfering factors. We found no discernable impact of the national anti-botnet initiatives on cleanup rates (Asghari *et al.*, 2015b). Some did well, others performed average or poorly. Perhaps, the ISPs in those countries did not act on Conficker infections, because the botnet was already inactive. If that was the case, then it was a mistake. The infected machines were vulnerable to basically every malware family that showed up after Conficker, which dates back to 2008. We did indeed find that significant portion of the infected population also was present in data from more recent, and still active, botnets.

4.1.2 Active regulators. Another factor that we studied was the relationship between regulatory involvement and infection rates. No regulator had a legal mandate to compel ISPs to do cleanup, except in Finland. In other words, we were interested in the effects of softer mechanisms, such as the fact that the regulator is active in this area and asks ISPs to step up, sometimes pressuring the market players into some form of voluntary self-commitment, such as the Dutch anti-botnet agreement among the ISPs or the US anti-botnet code for ISPs, instigated by the FCC and adopted by the leading ISPs. To measure these effects, we used a proxy for regulatory involvement, namely, whether the regulator was a member of a transnational initiative called the London Action Plan (LAP). LAP came out of anti-spam efforts, but had since broadened its scope to deal with other forms of abuse. Our analysis found a robust relation between LAP membership and lower infection rates in that country (Asghari *et al.*, 2015b).

4.1.3 Benchmarking. We also performed an in-depth study on the Dutch ISP market and how it performed in terms of botnet mitigation. We first solicited feedback on our approach for measuring comparative infection rates across providers, which would function as benchmarks. Once the approach was supported, we calculated the actual infection rates. We presented them during a closed meeting with only the ISPs. Two interesting things happened as a result of this effort. First, the staff of the worst performing ISPs contacted us

and asked for the report, so that they could go to their management board and ask for more resources. In the months after the meeting, they quickly improved and their infection rate dropped to lower than the market average. The second effect was tied to one of our other findings: all ISPs, including those working within the voluntary "anti-botnet contract", were only cleaning up a fraction of the bots. We had more data than they had, which simply reflected that they had never gone out of their way to acquire this type of data, which is available for free from volunteer organizations like Shadowserver and others. One of our recommendations for them was to also acquire these data, perhaps in a centralized manner, as that would reduce the cost for all of them. This led to another mechanism: a public-private clearinghouse for abuse data.

4.1.4 Public-private clearinghouse for abuse data. In the two years after our study, the government and ISPs agreed to jointly fund the development of a clearinghouse, called AbuseHub, which would give the ISPs much better data on which of their customers was infected. After the initial investment, the ISPs would pay a fee to cover all the cost of operating the clearinghouse. In parallel to the startup of AbuseHub, we ran a longitudinal measurement of infection rates of the members of AbuseHub and those who we not. We did not find very compelling evidence of a direct impact of the clearinghouse itself, but we did find that the Dutch ISPs had improved substantially. In our data, the country as a whole was now among the cleanest in the world, in terms of botnet infections (Van Eeten *et al.*, 2016).

To sum up: several governance mechanisms have emerged to deal with the threat of botnets, supplementing the efforts of the owners of the machines, and of Microsoft, who has been trying to protect owners by making its platform more resistant to infections. ISPs have been engaged through network forms of governance, with national governments in facilitating roles, sometimes combined with "the shadow of hierarchy", i.e. the threat that if the ISPs do not act themselves, a regulatory response will follow. Some of the efforts are transnational, as with the regulators collaborating via LAP, many are national. The ISPs that have successfully brought down infection levels have constrained the property rights of their customers, one could argue. Filtering and even quarantining infected end-user machines – that is, restricting their access to the internet – seems to work, if done at sufficient scale. The incentives of ISPs to internalize these externalities, for which they are not the root cause, seem to be a mix of soft pressure, reputation effects via benchmarking and cost sharing and reduction around remediation. Contrary to the alarmist reports about the growth of malware, end-user infection levels have held steady for more than five years now, at a global average of around 1 per cent of all Windows PCs (Asghari *et al.*, 2015a).

4.2 Mitigating Web compromise

Websites and other hosting services are abused for criminal purposes. A wealth of research has identified how hosting infrastructure shows up in various criminal business models. Think of phishing sites, command-and-control servers for botnets, child pornography, malware distribution and spam servers. Nobody contests that hosting providers play a key role in fighting cybercrime. Much of the criminal activity runs on compromised servers of legitimate customers, some on servers rented by the criminals themselves. In either case, hosting providers typically become aware of the problem only after being notified of the abuse.

Somewhat analogous to our work on ISPS, we have mapped the global markets for hosting providers, which consists of around 45,000 firms (Tajalizadehkhoob *et al.*, 2016). We leveraged a variety of data sources on compromised servers in that market to build comparative metrics for compromise rates at providers (Noroozian *et al.*, 2015, 2017). A key problem we needed to solve is how to normalize these rates and correctly account for size and other provider properties. Larger providers, with more customers and servers, will incur more incidents than smaller providers, even with the same security effort. We did indeed find that over 84 per cent of the variance in abuse incidents can be explained by

a number of simple size proxies (Tajalizadehkhoob et al., 2017a). Once we control for this difference in exposure to threats, we were able to compute comparative compromise rates.

4.2.1 Rule of law. In general, in the hosting market, we see fewer initiatives to mitigate abuse compared to the ISP market. This might be related to the much large number of providers: around 45,000 firms versus fewer than 300 in the ISP market. In most countries, engaging with fewer than 10 ISPs is enough to cover over 90 per cent of the whole market. In hosting, there is no such concentration. There are a few big players, but a very long tail of increasingly small ones. In the Netherlands alone, there are over 1,000 providers. Most of them are outside the reach of the type of network governance initiatives we have seen against botnets. They depend on relatively large companies with enough staff to send people to meetings, on established relationships and peer pressure within a known group and on publicity in which companies are mentioned, providing some level of public accountability. In the absence of identifiable provider initiatives against Web compromise, we first looked at institutional factors: are compromise rates lower in countries with better law enforcement and more mature infrastructure? We found no significant relationship (Tajalizadehkhoob et al., 2017a).

4.2.2 Pressure by law enforcement. In reports by security vendors, the Netherlands had long featured in the top when it comes to hosting malicious content. This is partially because the country's hosting infrastructure is disproportionately large, but it also seemed to reflect the presence of a few rotten apples in the market. The National Police, Public Prosecutor's Office and the Authority for Consumers and Markets asked us to use our metrics to identify the top 10 worst providers in the Netherlands – worst being defined as having the highest concentration of criminal activity in their network. After we produced the list, the authorities sent these firms a letter and visited them for a conversation about how they planned to improve the situation. One provider refused to meet. The meetings with the others took place and after these rounds of talks, we looked at changes in the abuse concentrations in these provider networks. We found no discernable impact. The worst provider remained the worst, one got worse than before, others moved a little down, a few other ones entered the top 10. There was no legal course of action to put more pressure on the top 10. What did happen was that a number of pro-active providers, not present in the top 10, took the initiative to try to mobilize the whole market in dealing better with abuse. Similar to the idea of AbuseHub, they wanted to provide better abuse data to providers and argue that, in the end, they would be better off by being more proactive. This process has only just started, and it is too early to tell if it will manage to change the security performance of Dutch providers.

4.2.3 Promoting best practices for hosting providers. Different industry organizations have suggested best security practices for providers (M3AAWG, 2015). One key best practice is to keep server and customer software patched, as servers get compromised at scale by exploiting known vulnerabilities. It is unclear, however, if providers have a strong enough incentive to pursue better patch levels and, if they do, whether they have enough control over the software stack to ensure better patching. Customers also install software, and providers rarely want to interfere with those installations, as they might break something on the side of the customer. In short, it is unclear whether these best practices have any influence.

We undertook a large-scale measurement study of so-called shared hosting (Tajalizadehkhoob et al., 2017b). This is a part of the market that offers inexpensive services by letting many customers share the same server. For technical reasons, this also means that customers have restricted privileges, so they cannot change the lower levels of the software stack. This is in the hands of the provider. We first mapped all shared hosting providers in the world. This amounted to a set of 1,259 providers. Then, we scanned a range of software packages on the server, from the operating systems and other infrastructure software, up to content management systems, which can be installed by

customers even when working under restricted system privileges. We found that there was a significant provider effect on patch levels. To put it differently: providers can make a difference, and a significant number of them do. We also found that this effort pays off: looking at the number of phishing sites, we found that the best 10 per cent of the market, in terms of patching, has four times fewer phishing sites than the worst 10 per cent of the market. In other words, even though the economic incentives to adopt best practices seem weak, a portion of the market is actually doing it and improving the security landscape in the process.

4.2.4 Abuse and vulnerability notification mechanisms. A lot of mitigation against hosting abuse is based on a voluntary mechanism called abuse reporting. Some entity – this could be another provider, a security company, an intermediary like Google or Facebook – sends a message to the abuse contact point and informs the provider that some system on its network is being abused for criminal purposes. The request is to act and stop the abuse. Does this mechanism actually work? We ran an experiment sending reports on newly discovered abuse incidents to providers. We not only tested cleanup rates, but also whether the reputation of the sender of the abuse report mattered. We found that, first of all, around half of the providers, in collaboration with the affected customer, remediated the situation. That is a significant effect for a voluntary mechanism with basically no means of enforcement or even visibility of the community into whether recipients acted on the reports. We also found that the reputation of the sender made no difference in our setup. Other studies had found that including the full technical details of the abuse event, rather than an easily digestible summary, led to more cleanup. Because of this, we included the full event details. Perhaps, this amount of information diminishes the need to trust the sender.

A related area of work that has recently emerged is the practice of vulnerability notifications. It is basically the same mechanism, but this time, the provider is informed of vulnerabilities that might get exploited, rather than actual abuse. With the rise of powerful new scanning software, many groups and organizations are running internet-wide scans for vulnerabilities. The results are then sent to the provider or resource owner via automated emails. The results of these efforts have been somewhat disappointing. As we discussed earlier in this paper: many of the relevant parties cannot be reached (Cetin *et al.*, 2017). Of those that are reached, most do not remediate the vulnerability (Stock *et al.*, 2016; Li *et al.*, 2016). This could potentially be the result of a rational tradeoff, as not all vulnerabilities get exploited, and perhaps the required remediation has negative effects for the resource owner. That being said, in light of the fact that many recipients are not even reachable, it seems overly optimistic to assume that most notifications are carefully evaluated.

4.3 IP source address spoofing

Our final empirical case is on the problem of IP source address spoofing. The basic TCP-IP protocol stack allows a device that is sending a packet to basically lie about it source address, the IP address from which it has been sent. The protocols in use do not contain any mechanisms to validate the source address. Despite source IP address spoofing being a known vulnerability for at least 25 years, and despite many efforts to shed light on the problem, spoofing remains a viable attack method for redirection, anonymity and DDoS amplification, as evidenced in February 2014 during a 400 Gbps DDoS attack against Cloudflare. Defeating amplification attacks, and other threats based on IP spoofing, requires providers to filter incoming packets from customer networks with spoofed source IP addresses – in other words, to implement BCP 38, a best current practice also known as source address validation (SAV).

4.3.1 Promoting source address validation. SAV suffers from misaligned incentives: a network that adopts SAV incurs the cost of deployment, while the security benefits diffuse to all other networks. That being said, SAV is a widely supported norm in the community.

One initiative where network providers commit to implementing BCP38 and promote wider adoption is called MANRS (Mutually Agreed Norms for Routing Security).

Several measurement projects are underway to increasing the visibility of what networks have or have not adopted SAV reduces the incentive problem by leveraging reputation effects and the pressure of other providers and stakeholders (Lone *et al.*, 2017; Beverly and Bauer, 2005). In reality, providers are not really aware which networks are BCP38-compliant. In some cases, they have been surprised by the results regarding their own networks. In such an environment, why would anyone adopt BCP38, incur all the cost of implementing and maintaining the dynamic or static filters and get basically no benefits, not even in the sense of recognition from the community?

Surprisingly, many networks have, in fact, adopted BCP38. The exact numbers are hard to provide, mostly because not all networks are tested. The most reliable data, collected via the Spoofer project[2], have limited coverage, but it still finds that more than half of all tested networks are compliant. Other data sources also find significant adoption rates. The security community tends to focus on the many networks that are not compliant. This is understandable, but should not obscure the important finding that adoption is, in fact, the more unlikely event and it is happening at a non-trivial scale.

4.4 Comparing the case studies

Summing up, we can observe that a variety of governance mechanisms have emerged around the security externalities that we studied. Remarkably, in all three areas, we find evidence of private actors partially internalizing the externalities. We encountered a surprising amount of voluntary action. This is not meant to say that these problems are solved. Not at all. One could easily argue that the current governance mechanisms are not nearly enough. That being said, it is an important finding that forms of collective action exist in this environment where control is extremely distributed and hierarchical structures are, by and large, absent. Given that environment, it will be difficult to replace or supplement them with something better.

In terms of scale, we have seen governance mechanisms that are geographically concentrated and involve a small number of actors within a specific country, to mechanisms that span the global market, involving thousands of actors. One might speculate that the mechanisms that operate on a smaller scale are more effective, as they can directly engage with relevant actors. However, the cases show mixed results in this respect. Think of the lack of impact of the pressure of Dutch law enforcement on poor performing hosting providers. Furthermore, on the other end of the scale, we find evidence that social norms propagated in the industry community turn out to have measurable impacts.

We can also map the cases onto the four canonical types of governance: market, hierarchy, network and community (Tenbensel, 2005; for a more in-depth discussion in the context of cybersecurity, see Kuerbis and Badiei, 2017). We found no examples of market mechanisms. We did encounter several examples of network governance, all at the level of countries. We also found two examples with a weak form of hierarchy: soft pressure of regulators to improve botnet mitigation seemed to work, while the "shadow of hierarchy" in the form of law enforcement visits made seemingly no impression on the worst Dutch hosters. And, we found two examples where community governance, in the form of social norms (more specifically: best practices around patching vulnerable servers and around anti-spoofing measures), had a measurable impact.

Reputation effects as an incentive could, in theory, work in all four governance types. As we did not study a market mechanism, we did not see it at work there. That said, the hosting providers we talked to were skeptical that customers in their market cared at all how well their provider did in fighting abuse. We did find that reputation effects worked in (national)

networks of actors, where even the private sharing of ISP benchmarks had an impact. It is unclear to what extent public reputation – i.e. naming, shaming and praising – can be scale to transnational networks or even global markets. One experimental study found it might (He *et al.*, 2016), but this is a subject of future work.

If nation states were involved, it mostly took the form of civil agencies participating in an industry network – in contrast to the predictions of the "militarization" narrative we discussed in earlier part of the paper. The one example where the state did assert its hierarchy, the law enforcement initiative against the hosters, was also where the state looked weak. The results were disappointing in terms of effectiveness.

5. Conclusion

In this paper, we defined governance as policy-driven control over resources, systems and services. In that light, much of what is now studied as governance is actually discourse about governance, rather than actual governance. At the level of discourse, the maligned trends of securitization and militarization are indeed observable, but there is little evidence that control is actually shifting toward nation states – the notable exception being offensive operations. States have become an important part of the threat landscape. This is something different, however, from providing governance over the defensive challenge of cybersecurity.

Control is located first and foremost in ownership. This brings into focus just how globally distributed authority over security policies are practices currently is. Security governance will have to function within this challenging environment. We explored four high-level shifts in governance in terms of changes in the property rights of owners, intermediaries and vendors as a consequence of security threats. One of the effects of the rise of the IoT is that the notion of "the Internet", or even "ICT", as a coherent space of governance might become obsolete. Security governance seems poised to become subsumed into the well-institutionalized regimes of safety and product liability, broken down into different sectors such as health, energy, transportation and so on.

In the final third of this paper, we explored a number of governance mechanisms that have emerged around three important security threats. While not all mechanisms are equally effective, in each of the three areas, we found evidence of private actors partially internalizing the externalities, mostly on a voluntary basis. No one thinks that this is enough, but it is a starting point. Future research is needed to identify how these mechanisms can be extended or supplemented to further improve the governance of cybersecurity.

Notes

1. For an ongoing count, visit http://openresolverproject.org/

2. www.caida.org/projects/spoofer/

References

Asghari, H., Ciere, M. and Van Eeten, M. (2015b), "Post-mortem of a zombie: conficker cleanup after six years", *24th USENIX Security Symposium (USENIX Security 15)*, Washington, DC.

Asghari, H., van Eeten, M. and Bauer, J.M. (2015a), "Economics of fighting botnets: lessons from a decade of mitigation", *IEEE Security & Privacy*, Vol. 5 Nos 16/23.

Beverly, R. and Bauer, S. (2005), "The Spoofer project: inferring the extent of source address filtering on the internet", *Proceedings of USENIX SRUTI, July.*

Cadwalladr, C. (2017), "The great British Brexit robbery: how our democracy was hijacked", *The Guardian*, 20 May, available at: www.theguardian.com/technology/2017/may/07/the-great-british-brexit-robbery-hijacked-democracy

Cetin, O., Gañán, C., Korczyński, M. and van Eeten, M. (2017), "Make notifications great again: learning how to notify in the age of large-scale vulnerability scanning", *16th Workshop on the Economics of Information Security (WEIS 2017)*, UC San Diego, Ja Jolla, 26-27 June.

Cetin, O., Jhaveri, M.H., Gañán, C., van Eeten, M. and Moore, T. (2015), "Understanding the role of sender reputation in abuse reporting and cleanup", *14th Workshop on the Economics of Information Security (WEIS 2015)*, TU Delft, 22-23 June.

De Nardis, L. (2014), *The Global War for Internet Governance*, Yale University Press, New Haven and London.

Englehardt, S. and Narayanan, A. (2016), "Online tracking: a 1-million-site measurement and analysis", *23rd ACM Conference on Computer and Communications Security (CCS 2016)*, Hofburg Palace, Vienna, 24-28 October.

European Commission (2017), "The Directive on security of network and information systems (NIS Directive)", available at: https://ec.europa.eu/digital-single-market/en/network-and-information-security-nis-directive

Fidler, B. (2017), "Cybersecurity governance during and after the postwar order: a historical analysis", Paper presented at "Who Governs? States or Stakeholders? Cybersecurity and Internet Governance, Third Annual Workshop Internet Governance Project", GA Tech School of Public Policy, Atlanta, 11-12 May.

He, S., Lee, G.M., Han, S. and Whinston, A.B. (2016), "How would information disclosure influence organizations' outbound spam volume? Evidence from a field experiment", *Journal of Cyber Security*, Vol. 2 No. 1, pp. 99-118.

Hurel, M.L. (2017), "Unpacking cybersecurity governance: institutional framings", paper presented at "Who Governs? States or Stakeholders? Cybersecurity and Internet Governance, Third Annual Workshop Internet Governance Project", GA Tech School of Public Policy, Atlanta, 11-12 May.

Korczyński, M., Król, M. and van Eeten, M. (2016), "Zone poisoning: the how and where of non-secure DNS dynamic updates", *Proceedings of the 2016 ACM on Internet Measurement Conference, ACM*, pp. 271-278.

Krebs, B. (2016), *Europe to Push New Security Rules Amid IoT Mess*, Krebsonsecurity.com, available at: https://krebsonsecurity.com/2016/10/europe-to-push-new-security-rules-amid-iot-mess/

Kuerbis, B. and F. Badiei (2017), "Mapping the cybersecurity institutional landscape", paper presented at "Who Governs? States or Stakeholders? Cybersecurity and Internet Governance, Third Annual Workshop Internet Governance Project", GA Tech School of Public Policy, Atlanta, 11-12 May.

Li, F., Durumeric, Z., Czyz, J., Karami, M., Bailey, M., McCoy, S., Savage, S. and Paxson, V. (2016), "You've got vulnerability: exploring effective vulnerability notifications", *25th USENIX Security Symposium (USENIX Security 16)*, USENIX Association, pp. 1033-1050.

Lone, Q., Luckie, M., Korczyński, M. and van Eeten, M. (2017), "Using loops observed in traceroute to infer the ability to spoof", *International Conference on Passive and Active Network Measurement, Springer, Cham*, pp. 229-241.

M3AAWG (2015), "M3AAWG anti-abuse best common practices for hosting and cloud service providers", available at: www.m3aawg.org/sites/default/files/document/M3AAWG_Hosting_Abuse_BCPs-2015-03.pdf

Mueller, M. (2010), *Networks and States: The Global Politics of Internet Governance*, MIT Press, Cambridge, MA.

Mueller, M. (2017), "Is Cybersecurity eating internet governance? Causes and consequences of alternate framings", paper presented at "Who Governs? States or Stakeholders? Cybersecurity and Internet Governance: Third Annual Workshop Internet Governance Project", GA Tech School of Public Policy, Atlanta, 11-12 May.

Noroozian, A., Korczyński, M., Tajalizadehkhoob, S. and van Eeten, M. (2015), "Developing security reputation metrics for hosting providers", *Proceedings of the 8th USENIX Conference on Cyber Security Experimentation and Test (USENIX CSET), USENIX Association*.

Noroozian, A., Ciere, M., Korczyński, M., Tajalizadehkhoob, S. and van Eeten, M. (2017), "Inferring security performance of providers from noisy and heterogeneous abuse datasets", *16th Workshop on the Economics of Information Security (WEIS 2017)*, UC San Diego, Ja Jolla, 26-27 June.

OECD (2010), *The Economic and Social Role of Internet Intermediaries*, OECD Publishing, Paris.

OECD (2011), *The Role of Internet Intermediaries in Advancing Public Policy Objectives*, OECD Publishing, Paris.

Schneier, B. (2012), "When it comes to security, we're back to feudalism", WIRED 11.26.12, available at: www.wired.com/2012/11/feudal-security/

Schneier, B. (2013), "You have no control over security on the feudal internet", *Harvard Business Review*, available at: https://hbr.org/2013/06/you-have-no-control-over-s

Stock, B., Pellegrino, G., Rossow, C., Johns, M. and Backes, M. (2016), "Hey, you have a problem: on the feasibility of large-scale web vulnerability notification", *25th USENIX Security Symposium (USENIX Security 16)*, USENIX Association, pp. 1015-1032.

Tajalizadehkhoob, S., Korczyński, M., Noroozian, A., Gañán, C. and van Eeten, M. (2016), "Apples, Oranges and hosting providers: heterogeneity and security in the hosting market", *IEEE Network Operations and Management Symposium (IEEE-NOMS 2016)*, Istanbul, 25-29 April.

Tajalizadehkhoob, S., Böhme, R., Gañán, C., Korczyński, M. and van Eeten, M. (2017a), "Rotten Apples or bad harvest? What we are measuring when we are measuring abuse", Paper in review at ACM Transactions on Internet Technology (TOIT).

Tajalizadehkhoob, S., van Goethem, T., Böhme, R., Gañán, C., Korczyński, M., van Joosten, W. and van Eeten, M. (2017b), "Herding vulnerable cats: a statistical approach to disentangle joint responsibility for web security in shared hosting", *Paper in review at 24rd ACM Conference on Computer and Communications Security (CCS 2017).*

Tenbensel, T. (2005), "Multiple modes of governance: disentangling the alternatives to hierarchies and markets", *Public Management Review*, Vol. 7 No. 2, pp. 267-288.

Thierer, A. (2014), *Permissionless Innovation: The Continuing Case for Comprehensive Technological Freedom*, Mercatus Center, George Mason University, Virginia.

Thierer, A. (2016), *Permissionless Innovation and Public Policy: A 10-Point Blueprint*, Mercatus Center, George Mason University, Virginia.

van Eeten, M. and Mueller, M. (2013), "Where is the governance in internet governance?", *New Media and Society*, Vol. 15 No. 5, pp. 720-736.

van Eeten, M., Lone, Q., Moura, G., Asghari, H. and Korczyński, M. (2016), "Evaluating the impact of abuse HUB on Botnet mitigation", arXiv preprint arXiv:1612.03101.

Zittrain, J. (2008), *The Future of the Internet - And How to Stop It*, Yale University Press, New Haven.

Corresponding author

Michel van Eeten can be contacted at: M.J.G.vanEeten@tudelft.nl