

See discussions, stats, and author profiles for this publication at: <https://www.researchgate.net/publication/313876687>

Stockpiling Zero-Day Exploits: The Next International Weapons Taboo

Conference Paper · March 2017

CITATIONS

0

READS

59

1 author:



[Paul Maxwell](#)

United States Military Academy West Point

16 PUBLICATIONS 25 CITATIONS

SEE PROFILE

Stockpiling Zero-Day Exploits: The Next International Weapons Taboo

Paul Maxwell¹

Army Cyber Institute, USA

paul.maxwell@usma.edu

Abstract: In the current state of global affairs, a market exists for zero-day exploits where researchers, nation states, industry, academia, and criminal elements develop, buy, and sell these commodities. Whether they develop zero-days or purchase them, nation states commonly stockpile them for the future. They may then use them for purposes such as: espionage, offensive cyber operations, or deterrent effect. The immediate effect of this stockpiling though is that the exploit is not divulged to the public and is therefore not remediated. In our increasingly networked and code dependent world, this creates the potential for a cyber disaster with yet unimaginable impacts on global stability. It is therefore imperative that nation states responsibly divulge zero-day exploits through an international framework for the global good. Moving from the current state of affairs to one where responsible release of zero-day exploits is the norm will not be easy. There are many stake holders who argue that keeping stockpiles is beneficial or that this is an area that is not feasible to regulate. However, as we have seen with weapons such as nuclear, chemical, and biological weapons, it is possible to develop international regimes that prohibit the use of such weapons due to their extraordinary capabilities and impact. Alternatively, should these exploits be seen as equally pernicious as contagious diseases, nations may join together to form organizations similar to the WHO that can address international cyber issues. If a taboo against the use of zero-day exploits can be established, i.e., we make their use morally illegitimate, the security of all users will be improved.

Keywords: cyber security policy, zero-day exploit, international norms

1. Introduction

A zero-day exploit is a vulnerability that is previously unknown, i.e., it has not been publicly reported so that remediation can be applied. These exploits are particularly dangerous because of their un-disclosed nature and the lack of protections against them. Nation states acquire and stockpile these exploits through development or purchase to use in actions such as, espionage, deterrence, and offensive cyber operations. The side effect of stockpiling a zero-day exploit is that the nation state contributes to the insecurity of our complex, networked world instead of improving public security.

Over time, the manner in which zero-day exploits are handled has evolved (Miller 2007). The initial trend was that security researchers (and hackers) would divulge their discoveries for the reputation building effects that accompanied the release. This stage was followed by the creation of bug bounty programs that offered both recognition and a small monetary award. Currently, opaque markets have developed where exploits are bought and sold for significant amounts. These trading venues allow “white market” sales to software vendors, “gray market” sales to governments and other offensive minded entities, and “black market” sales to criminals (Stockton and Golabek-Goldman 2013). These markets create a disincentive to voluntary disclosure for the public good by researchers. The markets can potentially incentivize software creators to make insecure code in the hopes of illicitly selling an exploit in the future (Schneier 2012). Within these markets, governments are often the highest bidders for these exploits (Greenberg 2012 and Menn 2013). When nations purchase and stockpile these exploits, a conflict of interest is created between its cyber goals (e.g., intelligence gathering, national defense) and its duty to protect its citizens.

There are numerous arguments and use cases supporting the case for stockpiling exploits of which a few are discussed here. Zero-day exploits can support national espionage requirements. Espionage is widely practiced and within bounds tolerated. Another application is law enforcement where knowledge of exploits can assist with fighting cyber-crime and physical crime. Penetration testing is an area where knowledge of exploits is used to help with security in a positive way. Some argue that the release of exploits provides knowledge to malevolent users about vulnerabilities of which they would otherwise be unaware (Rescorla 2004). Keeping exploits secret thereby limits the potential damage from them. Deterrence is often cited when discussing stockpiling of exploits. If nations are known to maintain these exploits then attackers may be deterred from action. Some would argue that certain nations have hegemonic cyber capability and therefore have no incentive to surrender their

¹ The views of the author are his own and do not reflect the views of the United States Military Academy, the United States Army, or the United States Government.

advantage. Simply put, their losses outweigh their gains if they reveal the exploits. Finally, pessimists would point out that the black markets for these items is uncontrollable and thus any effort to establish a policy of divulgence is wasted.

The arguments against keeping exploits secret are equally varied. Our society is increasingly networked and reliant upon software. No longer is the risk mostly from information and financial theft. Critical infrastructure is connected to the internet and at risk from cyber exploitation. Failure of these systems would have a substantial impact upon civilians. National and trade secrets are frequently exposed to network attacks. These secrets are almost universally stored and accessed digitally and are at risk of theft via cyber exploits. More data about users is accumulated and stored daily for various uses such as marketing, tracking, and medical care. Much of this data is collected unbeknownst to the user and the revelations the data may provide are only beginning to be discovered as the big data sciences advance. The average user of cyber systems is not technically capable of digitally protecting themselves and thus they rely upon companies and their government to secure their data and systems. The individual citizen has no means to ensure the local power plant is safe from cyber-attack. Many have no ability to secure their own devices never mind the ability ensure company 'X' is taking appropriate measures to secure their data. Lastly, stockpiling an exploit is no guarantee of its uniqueness. An unscrupulous vendor may sell an exploit to multiple buyers or researchers may independently discover the same exploit (Ozm 2005).

These dangers to civilians form the basis of a proposal to treat cyber security as a public good and to make the stockpiling and use of zero-day exploits an international taboo. It is proposed that nation states form international agencies and create norms that require the divulgence of zero-day exploits with the goal of improved cyber security for all. To make this point, the rest of this paper is organized as follows. In section 2, foundational work for this proposal discussing public goods, international taboos, and cyber deterrence will be explored. Section 3 discusses why cyber security should be treated as a public good. Developing a taboo against zero-day exploits is contained in section 4 and section 5 suggests how to create international agreement on cyber norms. The paper concludes in section 6 with recommendations for future research.

2. Foundational work

The author is unaware of work to date that calls for nation state divulgence of zero-day exploits. Geer (2014) comes closest with his concept of "vulnerability finding hegemony". He suggests that a single power could corner the vulnerability market, buying all vulnerabilities, and then divulge them. This idea assumes that the U.S. or some other power could accomplish this feat and then relies on that power to then act responsibly. The idea of a hegemonic power reigning in the international cyberspace seems unlikely to occur given the varied goals of nation states.

There is a body of work addressing elements of this divulgence proposal. One subject area is the concept of cyber security as a public good. On this topic, White et al. (2013) argue for cyber security as a public good and suggest that a legal framework should be established obligating governments to protect its citizens. Mulligan and Schneider (2011) offer that similar to public health, cybersecurity should be treated as a public good by actions such as developing education programs, funding studies, and mandating reporting of certain events. In their paper from 2015, Sedenberg and Mulligan extend the comparison of cybersecurity to public health. They recommend a "Doctrine of Public Cybersecurity" to ensure better security for the populace. Rosenzweig (2011) believes that cybersecurity is not currently a public good but instead that information sharing (e.g., threats, vulnerabilities) should be a public good. In this work, he suggests that governments can provide a baseline security similar to how law enforcement secures society.

In the realm of international taboos, Price (1995) has written about how chemical weapons became morally illegitimate and are now taboo. In that work, he asserts that a stigma against chemical weapon use was a necessary condition for non-use and that other conditions such as uncertain military value, moral and legal constraints, resistance of tradition-bound military culture, and increased logistical burdens contributed to the international ban. These elements led to the international agreements (Geneva Protocol of 1925, Hague 1899 conference) which forbade first use of these weapons.

Similar to chemical weapons, nuclear weapons have become taboo and international norms have developed to deal with their use. Gizewski (1996) describes how a nuclear weapons taboo developed due to the extreme

power of the weapons, their unmitigated effect on civilians, and the loss of control accompanied with uncertainty that results from using them. Adding to this is the idea that the mutual vulnerability of the nuclear powers helped to solidify the taboo against their use.

Taboos developed around the use of cluster munitions and land mines as well. Rappert and Moyes (2009) illustrate how a cluster munitions treaty developed. Concerns about “unacceptable harm to civilians” led the international effort to ban these devices. The authors point out that the Geneva Conventions role is to limit weapons in warfare and that the cluster munitions treaty addressed that issue by detailing accuracy, duration, and area of weapon effects. Similar concerns about cyber weapons use exist.

There are numerous works on the concept of cyber deterrence. Many would argue that stockpiling zero-day exploits aids in accomplishing cyber deterrence. Included here are some descriptions of work done on cyber deterrence that argues against the cyber deterrence.

As discussed in (Trujillo 2014), deterrence is limited by attribution, signalling, and credibility. It is well known that attribution is currently very difficult and is unlikely to be solved in the near future. Signalling requires actors to have well known intentions and reactions to the events that they wish to deter. Current secreting of exploits and the uses of them does not support this deterrent factor. Finally, credibility is lacking when the true capabilities of an actor’s cyber arsenal is unknown. Cyber attacks to date have been met with a high tolerance (at least publicly) and therefore the credibility of cyber deterrence is weakened. Add to this the fact that detection of cyber attacks is often delayed months or years. This weakens deterrence even further. The author concludes that instead of pursuing cyber deterrence in a traditional sense, nations should strengthen their defenses which is a form of latent deterrence. One way to do this would be to reveal zero-day exploits so that patches can be created and distributed.

The work of Stevens (2012) suggests that deterrence by denial (denying the source of an attack) is one of the current policies being implemented by nation states. He argues that for deterrence to be effective, norms must first be developed for cyberspace that promote constraint and overall good versus maximum utility. It is his opinion that because of the many non-state actors in cyberspace the development of deterrence norms may not be possible.

Lewis (2010) agrees with the idea of publishing constraints on cyber use and creating cyber norms that increase defensive strength thereby helping cyber security by lowering the perceived gains of cyber attacks. He states that deterrence is less effective when asymmetric losses are an outcome. States like the U.S. have more to lose in the cyber realm than some potential attackers. Additionally, a military response to cybercrime and espionage, though approved by the U.S., would be a departure from existing international norms on the use of force.

In the paper by Moore et al. (2010), a game theoretic model is used to argue that nations will not all choose to divulge exploits. Instead, depending on the social cost of exploit use and the technological capabilities gap between them, nations will either divulge or stockpile exploits with at least one nation stockpiling. The basis for their argument is a simple two state model using two pure strategies of stockpile or defend (divulge). They use the model to show that a Nash equilibrium exists for different social costs and technological gaps except when both nations defend. Their model however is limited in its usefulness for several reasons. Besides the fact that the model only deals with two nations states, their assumption that nations patching systems against an exploit protects against that exploit and precludes its future use is invalid. It is well known that patches don’t achieve full implementation for a variety of reasons. Systems can remain vulnerable for many years after patch release. Next, their technological sophistication variable is overly simplistic. Certainly some nations may be more sophisticated than others as a whole but given the market for exploits, a nation only needs money to obtain exploits. The use of two pure strategies reduces the value of the model as well. Given that the NSA claims to release 91% of the vulnerabilities it discovers (Menn 2015), pure strategies are not realistic. Finally, Nash equilibriums assume that the players have perfect information about the other players’ strategies. Given that attribution is difficult and that nation state policies on cyber-attacks are not well known, this assumption is invalid.

3. Cyber security: A public good

A public good is generally something that consumption of does not reduce its availability and is difficult to exclude individuals from using it (White et al. 2013). Examples of things that are public goods are health, safety, and defense. An individual's use of the protections provided by a nation's armed forces does not limit another individual's benefits from that same protection. Additionally, it is not possible to restrict an individual's benefits from the national defense protections. The benefits apply to all residents.

As discussed previously, cyber security could be considered a public good. Using cyber security does not prevent someone else from having it and if networks and systems are secured, then it is difficult to exclude individuals from benefitting from that security. As stated in the UN General Assembly Group of Governmental Experts on Developments in the Field of Information and Telecommunications in the Context of International Security (2015), cyber threats are a significant threat to public safety, nation security, and the stability of the global community. If nations were to treat cyber security as a public good, then they could create regulatory baseline security standards, develop education programs and messages, mandate reporting of critical events, and perhaps create an organization dedicated to providing network security for the populace (Mulligan and Schneider 2011). Similar to how law enforcement operates, nations could focus public cyber security efforts on overall security and not individual security. This could take the form of ensuring critical infrastructure is secured, that strict data security regulations are emplaced, or that liability for insecure software is placed on the creator. As mentioned by Sedenberg and Mulligan (2015), nations can assuage concerns over things such as exposing industry secrets and privacy issues by limiting access to certain information like data on vulnerabilities and sensitive research topics. In terms of nations stockpiling zero-day exploits, why are governments not held liable for damages resulting from known zero days in a manner that corporations are held liable for loses caused by bad faith or wilful and intentional departures from duty (Frye 2002)? If an exploit that a nation is stockpiling is used against its citizens, then that nation is in part complicit in the resulting losses. The obligation that governments have to protect their citizens was ignored in the interest of other dubious needs.

4. Developing a cyber weapons taboo

Even without declaring cyber security a public good, it is feasible to acknowledge cyber-attacks and the exploits that enable them as morally illegitimate and to create a taboo against their use. Several classes of weapons arguably have a taboo against their use such as chemical, nuclear, and cluster bombs/landmines. Understanding how those taboos developed can provide a framework for establishing a cyber taboo.

In discussing the chemical weapons taboo, Price (1995) proposes conditions for the development of weapons taboos. The first condition is that a stigma against the use of the weapon must exist in the public realm. Contributing conditions are: uncertain military value, resistance by military culture for usage, existence of moral and legal constraints, and extra logistical burden. Sometimes these taboos defy logic. In the case of taboos against nuclear weapons use, some modern conventional explosives are more powerful than some nuclear weapons but are not considered taboo (Gizewski 1996). However, nuclear weapons effect more than soldiers thus creating a negative public perception. Rappert and Moyes (2009) point out that limiting weapons use (making them taboo) is an inherent part of the Geneva Conventions. Their discussion on the 2008 Convention on Cluster Munitions illustrates how a weapon that poses "unacceptable harm to civilians" was banned due to concerns of the weapon's accuracy, duration, and the area of the weapon's effects.

Using Price's conditions, cyber weapons and the exploits that facilitate them could lead to a taboo against their use. Already, a public stigma is developing against cyber-attacks through the onslaught of personal data theft, identity theft, and cybercrime that occurs daily effecting civilians, industries, and nations alike. Attacks that disable critical infrastructure such as the electrical grid, the water supply, or the financial system would inarguably have enormous impact on ordinary citizens similar to nuclear weapon usage. This alone could form the basis of making their use morally repugnant.

However, other elements contribute to the case of a cyber weapons taboo. Legal constraints currently exist prohibiting cybercrime and the illegal use of networks for criminal activity. Next, despite a few documented military successes (e.g., Stuxnet, Syria 2007) there is uncertainty over the value of cyber attacks for military purposes. This in part exists because of the accuracy, duration, and area of effect of cyber weapons is generally unknown. Can cyber weapons be used in a precise manner similar to conventional bombs to effect only the target or at least satisfy collateral damage minimization? Given that exploits broadly effect machine

architectures and their installed software, limiting the effect an exploit enables is difficult and limiting the exploit to only that target is impractical. How long will the weapons persist after their usage? Once a cyber exploit is utilized, it is safe to assume that others will analyze it and potentially attempt to weaponize it. Exposed systems can remain vulnerable for many years despite the release of patches and signatures (which may be delayed for months or years until the exploit is discovered) as the experience with Conficker shows. As with chemical weapons, it may be argued that military leaders may be hesitant to use cyber weapons because they are not well understood and a perception that they are less than honorable and inglorious. Furthermore, using cyber exploits and payloads will often require transiting the networks of other nations enroute to the target potentially violating the sovereignty of those nations. A physical parallel of aircraft transiting the airspace of a nation to attack a foe would create a difficult diplomatic problem. The cyber version could as well putting our relationships at risk. Finally, even Moore et al. (2010) recognized that if the social cost of not divulging exploits was extremely high (at its maximum), then a strategy of universal divulgence would be acceptable. A taboo against cyber weapons and the exploits that enable them would indicate just such a situation.

5. Implementation of cyber norms against stockpiling exploits

If the ideas of a cyber weapons taboo and cyber security as a public good are combined, then it is reasonable to argue that nations should not stockpile exploits but should instead divulge them in the interest of public security and protection against cyber-attacks. Stevens (2012) argues that norms for cyberspace that promote doing the right thing vice seeking maximum utility should be developed. These norms can help channel, constrain, and constitute action in cyberspace. Lewis (2010) states that cyber norms that make states responsible for their citizens' actions may deter cyber-attacks. The bottom line is that norms provide a framework upon which policy can be developed and implemented. The challenge lays in how to obtain this goal. How do nations with frequently conflicting goals agree to a policy that limits their freedom of maneuver in cyber space?

Fidler (2003) describes how international norms evolved for protection against the spread of disease. Because nations cannot control activities outside of their borders, they collaborate to make international laws and organizations such as the World Health Organization (WHO). The WHO created many medical norms such as surveillance protocols where information about diseases is shared. This global governance model enables the elevation of public good and human rights above utilitarian goals like patent protection. An example of this is the TRIPS agreement that makes patients' rights to procure certain medicine more important than patent protections. An international cyber organization of this type could follow this example to develop an intelligence/threat sharing process.

The authors in Shiffman et al. (2002) suggest that a "punctuated equilibrium" model is how international norms and policies are formed. In this model, there needs to be a widely considered problem, perception of a fix, and a powerful coalition of actors desiring to solve the issue. Part of this model also requires elements of social messaging such as using symbols and images to provoke emotions that are supportive of the goals trying to be achieved. Cyber-attacks and crime are without doubt a widely considered problem. Reading professional materials or news articles and attending cyber conferences, it is easy to see the concern in industry, academia, and government. There are also many perceived fixes ranging from better coding to software/hardware solutions to improved laws. Agreeing not to harbor zero-day exploits can be part of a larger cyber norms package. Additionally, actors large and small desire to solve the issue. Large industry actors such as Microsoft have expressed the desire in their 2016 cyber norms document (Charney et al. 2016) to develop cyber norms and solve the problems faced. Even governments such as China have publicly stated the desire to curb hacking and support international efforts against cyber-crime (Strom and Talev 2015).

Florini (1996) believes that international norms are like biological genes that evolve and survive. She holds that the norms need to be prominent, interact well with non-competitive norms, and have favourable conditions to survive. This model suggests that norms can develop and spread in several ways one of which is when a new issue emerges in which prevailing norms are not well established. She suggests that a "norm entrepreneur" can help a norm develop and grow. This model fits cyberspace well. The issues with cyberspace are prominent as seen by the weekly reports of breaches, disclosures, and attacks. Additionally, norms against stockpiling exploits would work well with existing norms that support international cooperation in terms of criminal activity and health disclosure. Certainly, this is a relatively new issue in historical context and norms are clearly not well established. As Florini suggests, if transparency techniques such as the "trust and verify" method can be adapted to zero-day exploits, then cyber norms can be developed.

Finally, Bolton and Nash (2010) discuss the emergence of Non-Governmental Organizations' (NGOs) impact on the development of international norms. He uses the example of the roles played by NGOs in the Ottawa Convention of 1997 that banned cluster bombs and landmine use. He also illustrates how NGOs are helping to set norms on a variety of issues such as conflict diamonds, child soldiers, and the International Criminal Court. In this model, NGOs could be formed that influence the development of cyber norms. Organizations such as IEEE (working from a standards position) or cyber non-profit groups could work with nations to influence their policies and also to provide independent trust mechanisms. These groups can leverage societal opinions and emotions to create improved cyber security policies that are focused on universal benefits.

Overall, movement towards cyber norms has occurred and is not without reason. In the cyber-crime realm, the Council of Europe 2001's Budapest Convention on Cybercrime shows that international agreement on cyber can be achieved. Forty-nine states have ratified that agreement and six more are signatories. The agreement has broad goals that attempt to harmonize laws on cyber-crime and establish cooperation. The UN's Group of Governmental Experts is also making strides towards norm development. If nations, NGOs, and industry work together they can achieve similar agreements that benefit the whole. The efforts however should avoid focusing on the social elements of cyberspace that may not appeal to non-western nations. Instead it should focus on providing maximum security to all with a framework that is verifiable.

6. Conclusions and future work

Zero-day exploits will continue to be discovered and some will be utilized by actors for illicit purposes. Nation states should develop international norms for these exploits that require responsible divulgence of these exploits when discovered. The security of our networks, systems, and the data therein is crucial to the global community and must be buttressed where possible. Cyber security should be treated as a public good and as a result, governments should act to defend their citizens and the global community against cyber threats. One way to do this is to minimize security risks. Divulging zero-day exploits instead of stockpiling them is one way to obtain this goal. The gains from stockpiling these exploits are dubious and are outweighed by the costs to public security. Of course this cannot be accomplished by lone actors. Instead this requires nations to work together to create international bodies that develop and promote norms for cyberspace.

Naturally, to achieve this end state, much work needs to be done. Agreement on what are the cyber norms and how to support them is necessary. Finding common ground amongst actors with different goals is not simple. Similar to the Cluster munitions ban of 2008 (Rappert and Moyes, 2009), perhaps a starting point is the definition of acceptable norms (developed by NGOs or the like) followed by nations arguing why exceptions should be allowed to those norms instead of starting from nothing and building consensus. Techniques for verifying compliance need to be developed. To maintain the trust necessary for bans on zero-day exploits, verification will be critical similar to the case with nuclear weapons. How exploits are divulged will need to be determined so that software vendors can develop remedies and distribute them to minimize harm. It is well known that attacks often increase after divulgence (Bilge and Dumitras, 2012). However, vulnerabilities will eventually be disclosed and so this is only a question of "when" and not "if". Another area to work is how to encourage zero-day discovery and responsible release amongst researchers. It is necessary to prevent the currently quasi-open zero-day markets from migrating to the dark web where it is harder to monitor. In conclusion, there is much to be done but as with chemical weapons, landmines, cluster bombs, and biological weapons, international norms can be developed that make the world safer against these exploits and their weapons they enable.

References

- Asllani, A., White, C., and Etkin, L. (2013). Viewing Cybersecurity as a Public Good: The Role of Governments, Businesses, and Individuals, *Journal of Legal, Ethical and Regulatory Issues*, Vol. 16 (1), pp. 7-13.
- Bilge, L. and Dumitras, T. (2012). Before We Knew It: An Empirical Study of Zero-Day Attacks in the Real World, *Proceedings of the 2012 ACM conference on Computer and Communications Security*, Raleigh, NC, 2012, p. 833-844.
- Bolton, M. and Nash, T. (2010). The Role of Middle Power-NGO Coalitions in Global Policy: The Case of the Cluster Munitions Ban, *Global Policy*, Vol. 1 (2), pp.172-184.
- Charney, S., English, E., Kleiner, A., Malisevic, N., McKay, A., Neutze, J., and Nicholas, P. (2016). From Articulation to Implementation: Enabling Progress on Cybersecurity Norms, Microsoft, [online], Available at: https://mscorpmedia.azureedge.net/mscorpmedia/2016/06/Microsoft-Cybersecurity-Norms_vFinal.pdf, [Accessed 7 Sep 2016].
- Council of Europe (2001). Convention on Cybercrime, *Treaty No. 185*, [online], Available at: <https://www.coe.int/en/web/conventions/full-list/-/conventions/treaty/185>, [Accessed: 7 Sep 2016].

- Geer, D. (2014). Cybersecurity as Realpolitik. [online], Available at: <http://geer.tinho.net/geer.blackhat.6viii14.txt>, [Accessed: 13 Oct 2016].
- Fidler, D. (2003). Emerging Trends in International Law Concerning Global Infectious Disease Control, *Emerging Infectious Diseases*, Vol. 9 (3), pp 285-290.
- Florini, A. (1996). The Evolution of International Norms, *International Studies Quarterly*, Vol. 40, pp. 363-389.
- Frye, E. (2002). The Tragedy of the Cybercommons: Overcoming Fundamental Vulnerabilities to Critical Infrastructures in a Networked World, *The Business Lawyer*, Vol. 58 (1), pp. 349-382.
- Gizewski, P. (1996), From Winning Weapon to Destroyer of Worlds: The Nuclear Taboo in International Politics, *International Journal*, Vol. 51 (3), pp. 397-419.
- Greenberg, A. (2012). Shopping for Zero-Days: A Price List for Hackers' Secret Software Exploits, *Forbes*, [online], Available at: <http://www.forbes.com/sites/andygreenberg/2012/03/23/shopping-for-zero-days-an-price-list-for-hackers-secret-software-exploits/#52cf31ce6033>, [Accessed 7 Sep 2016].
- Group of Governmental Experts (2015). Developments in the Field of Information and Telecommunications in the Context of International Security, *United Nations General Assembly*, Seventieth Session, Item 93, A/70/174, p. 17.
- Lewis, J. (2010). Cross-domain Deterrence and Credible Threats, *Center for Strategic & International Studies*, p. 5.
- Menn, J. (2015). NSA Says How Often, Not When, It Discloses Software Flaws, *Reuters Global Energy News*, [online], Available at: <http://www.reuters.com/article/us-cybersecurity-nsa-flaws-insight-idUSKCN0SV2XQ20151107>, [Accessed 7 Sep 2016].
- Menn, J. (2013). Special Report: U.S. cyberwar strategy stokes fear of blowback, *Reuters Cybersecurity*, [online], Available at: <http://www.reuters.com/article/us-usa-cyberweapons-specialreport-idUSBRE9490EL20130510>, [Accessed 7 Sep 2016].
- Miller, C. (2007). The Legitimate Vulnerability Market: Inside the Secretive World of 0-day Exploit Sales, *Independent Security Evaluators*, p. 10.
- Moore, T., Friedman, A., and Procaccia, A. (2010). Would a 'Cyber Warrior' Protect Us? Exploring Trade-offs Between Attack and Defense of Information Systems, In: *NSPW'10*, Concord, MA, pp. 85-94.
- Mulligan, D. and Schneider, F. (2011). Doctrine for Cybersecurity, *Dædalus, the Journal of the American Academy of Arts & Sciences*, Vol. 140 (4), pp. 70-92.
- Ozment, A. (2005). The Likelihood of Vulnerability Rediscovery and the Social Utility of Vulnerability Hunting, In: *The Workshop on Economics and Information Security*, Cambridge, MA., p. 21.
- Price, R. (1995). A Genealogy of the Chemical Weapons Taboo, *International Organization*, Vol. 49 (1), pp. 73-103.
- Rappert, B. and Moyes, R. (2009). The Prohibition of Cluster Munitions, *Nonproliferation Review*, Vol. 16 (2), pp. 237-256.
- Rosenzweig, P. (2011). Cybersecurity, the Public/Private "Partnership," and Public Goods, *Hoover National Security and Law Task Force*, Available at: <http://ssrn.com/abstract=1923869>, p. 29.
- Schneier, B. (2012). The Vulnerabilities Market and the Future of Security, *Forbes*, [online]. Available at: <http://www.forbes.com/sites/bruceschneier/2012/05/30/the-vulnerabilities-market-and-the-future-of-security/#5faa1aa47763>, [Accessed 7 Sep 2016].
- Sedenberg, E. and Mulligan, D., (2015). Public Health as a Model for Cybersecurity Information Sharing, *Berkely Technology Law Journal*, Vol. 30 (3), pp. 1687-1737.
- Shiffman, J., Beer, T., and Wu, Y. (2002). The Emergence of Global Disease Control Priorities, *Health and Policy Planning*, Vol. 17 (3), pp. 225-234.
- Stevens, T. (2012). A Cyberwar of Ideas? Deterrence and Norms in Cyberspace. *Contemporary Security Policy*, Vol. 33(1), pp. 148-170.
- Stockton, P. and Golabek-Goldman, M. (2013). Curbing the Market for Cyber Weapons, *Yale Law & Policy Review*, Vol. 32 (PPP), pp. 101-128.
- Strohm, C. and Talev, M. (2015). China Vows to Curb Commercial Hacking in Agreement With U.S., *Bloomberg Politics*, [online], Available at: <http://www.bloomberg.com/politics/articles/2015-09-25/obama-says-u-s-china-agree-to-curb-hacking-for-trade-secrets>, [Accessed 7 Sep 2016].
- Trujillo, C. (2014). The Limits of Cyberspace Deterrence. *Joint Forces Quarterly*, 4th Quarter 2014, pp. 43-52.