Elliott K, Massacci F, Williams J. Action, Inaction, Trust and Cybersecurity's Common Property Problem. *IEEE - Security and Privacy Economics* 2016, 14(1), 82-86.

DOI link to article:

http://dx.doi.org/10.1109/MSP.2016.2

Date deposited:

30/03/2016

Security & Privacy Economics

Action, Inaction, Trust, and Cybersecurity's Common Property Problem

Karen Elliott | Newcastle University Business School

Fabio Massacci | University of Trento

Julian Williams | Durham University Business School

Inaction is sometimes the optimal path—a point well taken by economists. In her substantial volume on optimal control, economist Nancy Stokey begins: "In situations where action entails a fixed adjustment cost, optimal policies involve doing nothing most of the time and exercising control only occasionally."[1] When fixed adjustment costs exist, the investment profile over time tends to be characterized by sudden shifts followed by periods of complete inaction, economists often refer to this as a "lumpy" investment profile.

Much literature on information security focuses on how cybersecurity threats occur and how to best resolve them. However, additional factors, such as the risk environment, interdependent actors, attackers' reuse of exploits, and patching vulnerabilities contribute to Chief Information Security Officers' (CISOs') defense strategies. Indeed, in many cases, the optimal decision is to wait until the degree of uncertainty changes and the benefits of action outweigh the costs.

The defensive investment problem

CISOs typically ask two questions: i) How do we measure our return on security investment?, and ii) How, on an empirical cost–benefit basis, do we know when to patch, fix, or shut down systems as and when new vulnerabilities arise? Both questions address optimal control problems in the presence of fixed adjustment costs. These costs might be known, or they may contain uncertain forward-looking components, and there is a tradeoff between these costs versus uncertain future gains.

It's difficult to measure return on security investment. The constantly evolving state of the "market for attacks" increases the difficulty in determining a security investment's true value. A firm's senior corporate officer rarely knows empirically (for instance, by audit) whether the firm has had no known security incidents because the firm:

- is spending exactly the right amount on security;
- is spending many times more than it needs;
- is spending too little, but attackers haven't stumbled across its vulnerabilities or found it worthwhile to exploit them; or
- is under attack but doesn't know it.

To determine a decisions optimal timing, much information is needed about the nature of uncertain future outcomes, for example, a security manager might wait to see if an exploit will be available for a vulnerability. However, managing cybersecurity investment receives less quantitative support than other typical risk-management activities undertaken by a firm. For instance, most firms actively manage interest rate and foreign exchange risk through their treasury management functions. These activities are carefully accounted for in corporate reports alongside their normal operational activities. Booked losses on hedging can be very large, but senior corporate officers and investors generally understand that it's important to hedge currency and interest rate risk, even if the specifics are hazy.

On the other hand, many CISOs struggle to procure adequate budget until after a significant event has occurred. The decision to invest in a security fix or control appears to be increasingly well understood—for large technology companies at least. The decision process for the deployment of a patch is roughly as follows:

- Determine the security flaw's severity and level of impact on the organization (possibly using the US National Institute of Standards and Technology Common Vulnerability Scoring System calculator);
- Determine the danger of implementing a patch and how much testing is required to ensure that the patch is less destructive than the threat; and
- After weighing the first two steps, triage the update to either an immediate implementation or a regular update cycle.

In a case study on optimal patching, Christos Ioannidis and his colleagues, amongst others, postulated a quantitative tradeoff between the increasing risk of doing nothing and the deterministic cost associated with potentially incomplete mitigation.[2] Indeed, the patching problem is an archetypal fixed adjustment cost in a security setting; part of the objective of this paper is to provide a consistent treatment of this problem.

So, instead of management's failure to provide resources to underfunded information security departments being a catastrophic misstep, a delay in the implementation of security investment controls might be a sensible tradeoff between risk and investment. Many economic models suggest that the tactic of postponing updates might be gaining popularity—not because C-level employees are taking unreasonable risks, but because of an older, much more formidable foe: the tiny adjustments that drive us to the Nash equilibrium, wherein agents continually make choices as they strategize actions and respond to those of others.

In security decision making, we can model the Nash equilibrium problem using three groups of agents: attackers, firms, and government. Attackers decide to invest in a malware or hacking effort, firms in defensive security and regulations, and government in enforcement. Their payoff structures will differ. For instance, hackers might value chaos over money, and firms and governments might value coverage in addition to a simple likelihood × impact calculation. Each will have their own subjective discount factors transforming future value of costs and benefits into risk adjusted current values; therefore, the relative present valuation of costs and benefits will be idiosyncratic across the various agents in the economy.

The Attackers' Economy

Agents working as 'attackers' are economic actors with preferences—who are the attackers, and what do we know about them? Prior security investment literature typically views attackers as essentially random-number generators.[3] Generators consider a set of vulnerabilities in commonly used software, firmware, and hardware and then throw malicious agents at this set. Eventually, technical proficiency and vulnerability combine to create a tool that can threaten the economic and physical well-being of the selected targets.

Looking at the relative scale of a threat versus the scale of investment to mitigate the threat, the 2015 UN estimated that global annual gross domestic product (GDP) was estimated between $60 and $80 trillion in 2014.[i] Estimates for the size of the cyber security industry are somewhat difficult to ascertain, in 2014 Gartner estimated that the cybersecurity industry accounts for approximately $77 billion—less than one-tenth of 1 percent of global GDP compared to conventional security expenditure on defense equipment and physical security, which is approximately 4 percent of the GDP at just under $400 billion.[ii]

On the other side of the attack-defense equation, in a study examining transactions in a Russian online hacker market (which Google and the US Federal Bureau of Investigation indicate accounts for a majority of online deployed malware tools), we found that transaction sizes are quite low, often in the hundreds of dollars, and only rarely in the tens of thousands.[4] Although the underground hacker market appears to be a well-functioning economy, it is potentially, significantly smaller than the opposing security industry.

Of course, the unit of account for losses might differ dramatically from the unit of account for rewards. If we look at insurance claims against cyber attacks from industry surveys, the claims from US firms are similarly very small; between 2011 and 2013, the median claim was $750,000 and the high was $13.5 million. This individual claim represented approximately 10 percent of total claims made.[5]

What do we take from this? The data on the insurance market and our understanding of the level of available coverage is incomplete. However, if the level of actual damage is so small, then the balance of investment and coverage would indicate an economic puzzle that deserves more research.

Attackers' Motivation

An additional puzzle comes from a study exploring the menus of vulnerabilities in attackers' malware kits, concluding that attackers are in fact, lazy.[6] Owing to the costly effort in developing new tools, attackers persist with malware based on existing vulnerabilities, long after effective patches have been introduced to the market as opposed to exploiting new vulnerabilities in the system(s). Fixed costs appear to make attackers investment decisions as similarly "lumpy" (i.e., uncertain) as those surrounding the defense dilemma decisions of their corresponding targets.

Another interesting facet of cyber attackers is their psychological profile and self-perception in terms of criminality, which affects software engineers' decisions to deploy labor for legal productive efforts or those deemed illegal, such as taking control of vulnerable machines and then selling these 'shells' for exploitation of financial records or deploying them for large computational tasks to other criminals.. Attackers appear to be able to switch liberally between standard software engineering projects and those that would normally be deemed illegal. The criminology literature indicates that the profiles associated with a cyber attacker reveals a far lower persistence in offending type; hackers choose to do work they feel is optimal for their own welfare rather than identify themselves by the offending activity.[7,8]

Although this is somewhat unhelpful for quantitative work, we can reasonably conjecture that the pool of threats that security industry faces is uncertain. If attackers fixed costs change, we could see a sudden and dramatic increase or decrease in attacking intensity, with little way to predict such shifts. As discussed, because of a lack of robust historical data surrounding attackers and predicting their behavior, each observation might be the result of an equilibrium formed from a very different experiment. Identifying causal relationships directly from data is an inherently fraught process, and the lack of detailed understanding of the attacker production function compounds this problem (for more information, see the "Econometrics" sidebar).

sidebar goes here

Externalities and Dependencies

How do firms operational level (micro) security decisions aggregate to the macro and hence the public policy level? Aggregation brings certain benefits as idiosyncratic impacts from events on single firms even out. However, public policy mandates on security policy must be implemented at the micro level, and inappropriately onerous requirements could generate costs for the productive side of the economy that are potentially unwarranted and most certainly unfair. The cybersecurity literature is starting to demonstrate an emerging awareness that the ability to control both the risk-generating mechanism and the source of contingent compensation in the event of a breach can lead to onerous rents.[9] Ranjan Pal and his colleagues illustrated this problem by devising a model in which a security vendor provides both a monopoly service and a monopoly provision of insurance, and the combined monopolies generate a substantial profit. However, the realization that the "substantial" profit from insurance and security activity is not globally desirable.

In a network of firms, the provision of security has several dependencies, both indirect and direct (that is, through direct technical interconnections, such as shared data facilities and electronic communications networks for financial institutions). Direct connections have been studied extensively in the recent literature, whereas indirect connections are a more recent research interest. For a classic description of the interdependency problem in security, see "Interdependent Security," and for a full network game with contagion, see *Network Security and Contagion*.[10,11]

Indirect connections address the risk environment. This is the change in a firm's risk profile due to the choices of other firms in the network—not through direct linkages but as a result of changes in the overall number and intensity of attackers as a result of their perceived returns on investment. The provision of public goods in networks has been the subject of significant interest in recent research.[12,13] From a security perspective, it's important that investment has a public good component in addition to the private benefits to the firm.

A good is considered public if it is non-rival and non-excludable—that is, the good is enjoyed

simultaneously by an unlimited number of consumers, and it is not possible to prevent others from gaining free access to the good. Note; that only aspects of security have a common property through the aggregate effect on attackers' expected payoffs. Moreover, as our title suggests, it may be more appropriate to consider aspects of security to be closer to a common property good, that is the cost of exclusion in consumption of security investment is very high as opposed to impossible in the pure public good case.

Taking the above forward, it seems that if I increase my effort in an activity and it has a positive spillover effect to you (e.g., I invest in more security and discourage a small amount of the aggregate number of attackers, thereby reducing my own risk), then all agents in my network engage in this virtuous cycle until a Nash equilibrium is reached. (Note: that this might not be as desirable as a coordinated action mediated by policymakers). However, consider the patching problem for network or client software. Many firms' information platforms modular components are specialized and interconnected. Applying a patch in one system might have unintended consequences for other systems (e.g., if a vendor drops legacy support). As such, patches commonly need to be tested, particularly for critical systems. This means that applications of patches have fixed costs, and as we ramp up these fixed costs, the degree of patching coverage drops, and the firms in an economic network suffer through the interdependency in security as we forestall or neglect investments at critical points. This vulnerability does not stem from direct interconnections but via the attractiveness for attackers to invest in attacks that often have very little specific targeting other than a certain platform or a vulnerable library still in use with a legacy system. The attacked firm may not have been very profitable for the attacker, but their hit is the result of other more attractive prospects still being out there. Even though the attack may not have been particularly profitable to the attacker, the damage to the firm may still be quite severe.

The "lumpy" investment profile is also reflected in the security interdependencies with other firms; more importantly, the lumpy profile of a large firm can be felt across the network either directly or indirectly. Indeed, this observation formed the basis of early research on the importance of liability sharing in security patch management.[14] Hence, fixed costs appear to exaggerate already problematic issues of externalities, transmitting costs between firms, and form the basis of our conjecture that unpredictable investment generates excess aggregate security threats. If attackers can expect to make a good profit because somebody out there is unpatched, they will continue to invest time and effort in their current technology before switching to a new one. The important point is that the opportunity set and expected reward for attackers will be formed from the aggregation of all of the unpatched vulnerabilities, many of which will be the result of small delays in investment. Hence, there is potentially a feedback mechanism that is sustaining risks to firms beyond the process of new vulnerability discovery.

Waiting to invest in cybersecurity may be deemed a poor risk management strategy, despite many standard models indicating that delaying investment until the nature of the uncertainty is clear is often the most appropriate course of action. A manager investing in cybersecurity infrastructure must determine the costs of inaction (which might be a function of accumulating risks) versus the cost of action (which might be fixed or have a random forward-looking component). These optimal control problems will likely depend on the joint decision making of all actors in a security context. This is in contrast to models of optimal decision making that treat a threat as a random external event in their environment (i.e., an emergent risk). Furthermore, if attackers have the same type investment decision making problem (upfront fixed investments and continuous variable costs), we might find that the adjustment path for the intensity of attacks on firms increases the unpredictability of associated risks. Ergo, small changes in regulatory policy could lead to substantial unexpected changes in the threat environment.

Acknowledgments

The authors gratefully acknowledge the support of the UK Technology Strategy Board grants Cloud Stewardship Economics and Trust Economics as well as the European Commission Seventh Framework Programme for Research and Technological Development (FP7) project "SECONOMICS" grant agreement 285223.

References

1. N.L. Stokey, The Economics of Inaction: Stochastic Control Models with Fixed Costs, Princeton Univ. Press,

2008, p. 1:1.

2. C. Ioannidis, D. Pym, and J. Williams, “Fixed Costs, Investment Rigidities, and Risk Aversion in Information Security: A Utility-Theoretic Approach,” *Economics of Information Security and Privacy III*, B. Schneier, ed., Springer, 2013, pp. 171–191.
3. A. Gordon and M.P. Loeb, “The Economics of Information Security Investment,” *Proc. ACM Trans. Information and Systems Security*, 2002, pp. 438–457.
4. L. Allodi, M. Corradin and F. Massacci. “Then and Now: On The Maturity of the Cybercrime Markets. The lesson black-hat marketeers learned,” *Forthcoming IEEE Transactions on Emerging Topics in Computing.*
5. “NetDiligence Cyber Claims Study” NetDiligence, 2014; www.netdiligence.com/NetDiligence_2014CyberClaimsStudy.pdf.
6. L. Allodi and F. Massacci, “The Work-Averse Attacker Model,” *Proc. European Conf. Information Systems* (ECIS 15), 2015; http://aisel.aisnet.org/ecis2015_cr/7.
7. B. McCarthy, “New Economics of Sociological Criminology,” *Ann. Rev. Sociology*, vol. 28, 2002, pp. 417–442.
8. G. Kirwan, The Psychology of Cyber Crime: Concepts and Principles, IGI Global, 2011.
9. R. Pal et al. “On a Way to Improve Cyber-Insurer Profits When a Security Vendor Becomes the Cyber-Insurer,” *Proc. IFIP Networking Conf*., 2013, pp. 1–9.
10. H. Kunreuther and G. Heal, “Interdependent Security,” *J. Risk and Uncertainty*, vol. 26, no. 1, 2003, pp. 231–249.
11. D. Acemoglu, A. Malekian, and A. Ozdaglar, *Network Security and Contagion*, tech. report 19174, Nat’l Bureau Economic Research, June 2013.
12. Y. Bramoullé, R. Kranton, and M. D’Amours, “Strategic Interaction and Networks,” *American Economic Rev*., vol. 104, no. 3, 2014, pp. 898–930.
13. N. Allouch, “On the Private Provision of Public Goods on Networks,” *J. Economic Theory*, vol. 157, 2015, pp. 527–552.
14. H. Cavusoglu, H. Cavusoglu, and J. Zhang, “Security Patch Management: Share the Burden or Share the Damage?,” *Management Science*, vol. 54, no. 4, 2008, pp. 657–670.

*REFERENCE FOR SIDE BAR: B. H. Baltagi, 2011. *Econometrics*. New York: Springer.

Karen Elliott is a Lecturer in Management at Newcastle University Business School. Contact her at karen.elliott@ncl.ac.uk.

Fabio Massacci is a Professor in the Department of Engineering and Computing Science at University of Trento. Contact him at fabio.massacci@unitn.it.

Julian Williams is a Professor of Accounting and Finance at Durham University Business School. Contact him at julian.williams@durham.ac.uk.

Sidebar

Econometrics

In econometrics, we often seek to identify exogenous and endogenous variables in a system of regression equations. By exogenous variables we refer to regressors that are uncorrelated with the noise term inherent in the regression. Endogenous variables exhibit some level of correlation with the noise term. For instance if

an explanatory variable in one regression equation is an independent variable in another regression, it is an endogenous variable within the system of equations. Several empirically driven approaches, such as the use of linear and nonlinear instrumental variable regressions, have been proposed to correct for the identification issues inherent in empirical models in which knowledge of the underlying process is not well understood. Evidence from a broad range of micro-econometric studies have illustrated that the exclusion of appropriate instrumental variables from empirical models can result in highly misleading inferences. See chapters 9 and 10 of Badi Baltagi's book for a good summary of endogeneity, instruments and multiple equation modeling in regression analysis.

Abstract: Cybersecurity tends to be viewed as a highly dynamic, continually evolving technology race between attacker and defender. However, economic theory suggests that in many cases doing "nothing" is the optimal strategy when substantial fixed adjustment costs are present. Indeed, the authors' anecdotal experience as chief information security officers indicates that uncertain costs that might be incurred by rapid adoption of security updates substantially delay the application of recommended security controls, so the industry does appear to understand this economic aspect quite well. From a policy perspective, the inherently discontinuous adjustment path taken by firms can cause difficulties in determining the most effective public policy remit and the effectiveness of any enacted policies ex post. This article summarizes this type of policy issue in relation to the contemporary cybersecurity agenda.

Keywords: Cybersecurity, return on security investment, fixed adjustment costs, real option value of delay, the public good aspect of security

---

[i] See the UN World Economic Situation and Prospects: http://www.un.org/en/development/desa/policy/wesp/.

[ii] See the Gartner report link here http://www.gartner.com/newsroom/id/2828722.