Behavioural Modelling of Attackers' Choices

S. Panda, I. Oliver & S. Holtmanns
Cybersecurity Group
Nokia Bell Labs, Finland

ABSTRACT: This paper examines a cyber environment involving attackers and telecommunications operators from attackers' perspective. We incorporate a behavioural approach to understanding attackers' behaviour during the attack process. Traditionally, security games have been analysed assuming the attackers to be of strictly bounded rationality or strategy-less. Furthermore, studies consider attackers do aim to maximise their expected gain which contradicts the assumption of bounded rationality of attackers. We have analysed security interactions considering attackers as rational entities with attack strategies. To understand the thought process and behavioural decision-making choices of attackers, we utilise a decision analysis model capturing the attack process. Based on our analysis, we propose a framework providing a way to enhance attack strategies against cooperating and non-cooperating (competing) operators. This study is intended to capture essential characteristics of an attacker to comprehensively understand and predict their expected behaviour assisting cybersecurity.

1 INTRODUCTION

The major concerns in cybersecurity are to measure the security risks and to determine the effectiveness of one's security investments against perceived threats. As in cybersecurity, security is defined by not only on an individual's security-related investments but also by others' security investments (Anderson & Moore 2006, Laszka, Felegyhazi, & Buttyan 2015). This security interdependence adds additional complexities in quantifying the security risks and crafting appropriate measures against it.

Game theory, being a mathematical modelling tool, has been widely used to study varied aspects of security (Roy, Ellis, Shiva, Dasgupta, Shandilya, & Wu 2010, Merrick, Hardhienata, Shafi, & Hu 2016, Altman, Boulogne, El-Azouzi, Jiménez, & Wynter 2006, Liang & Xiao 2013) and privacy (Manshaei, Zhu, Alpcan, Bacşar, & Hubaux 2013). Most of the work focused on studying defenders' behaviour and have proposed strategic recommendations which include stochastic approaches, frameworks, cognitive and behavioural models strengthening defenders' chances of successfully defending against attempted attacks. Studies have often assumed strategy-less behaviour of adversaries with a prescribed set of actions consistent with the threat models. However, alongside defenders, attackers are also intelligent entities and this assumption is not ideal in real-world situations which consists human adversaries (Camerer 2011).

In cybersecurity, attackers' behaviour has been less explored due to lack of reliable data on their intentions and interactions limiting our understanding of their characteristics and behaviours. (Veksler & Buchler 2016) and (Anderson 2009) have indicated that cognitive approaches can aid in predicting attackers' behaviour addressing real-world security problems.

In addition, over the past years, adversaries have become more financial oriented (Gordon 1994, Gordon 2000, Franklin, Perrig, Paxson, & Savage 2007) making them highly unpredictable. Some intentions behind these malicious activities are instigated by curiosity, or for peer recognition, and are often undecided in terms of ethical legitimacy (Gordon 1994, Gordon 2000). The possibilities of using illegal methods provoke new classes and strategies of attacks creating a need in studying and analysing attackers' behaviour to understand their intentions and decision making criteria.

We performed a game-theoretic investigation on attackers' strategies in the context of cybersecurity. The examined scenarios illustrate security games between attackers (cybercriminals, hackers) and telecommunications operators (defenders). An attacker is an external entity with malicious objectives attempting to break through the security of the targeted entity/system with an intention to hamper the existing state of the target.

A behavioural approach is utilised to anticipate

decision-making behaviour of attackers. We intend to determine attack strategies optimising attackers' effort in performing an attack and improving their perceived utility. A viewpoint this paper aims to highlight is when attack strategies are taken into consideration, what can the choice of not attacking signify?

This study is a step towards understanding the mentality of attackers and their decision-making behaviour from a cybersecurity perspective. Lack of decisive information on adversaries along with the available security information being highly asymmetric - favouring the attackers; results from this study can be used by defenders in assessing their conditions and perceiving the most expected attack strategies.

The rest of the paper is organised as follows. Section 2 covers the relevant literature and highlights the relationship of our work with existing research. Section 3 discusses the behavioural aspects of attackers and presents an attack framework disintegrating the efforts required in an attack process. An analysis of how attack strategies can be optimised using the attack framework is explained in Section 4. Section 5 discusses our findings and concludes this paper.

2 RELATED STUDIES

Even though this paper is confined towards understanding attackers, the complementarity of attackers' behaviour on operators' state is such that modelling them without an underlying operators' state is complex and unrealistic. An operator's state is defined by his security-related investments and relationships with other operators - cooperation (Laszka, Felegyhazi, & Buttyan 2015, Kunreuther & Heal 2003, Varian 2004, Hota & Sundaram 2015) and competition (Jiang, Anantharam, & Walrand 2008, Sun, Kong, He, & You 2008, Khouzani, Pham, & Cid 2014, Panaousis, Fielder, Malacaria, Hankin, & Smeraldi 2014), which induces additional security dependencies.

Attackers, alike operators (defenders), have strategic incentives and work towards maximising their expected utility (Laszka, Felegyhazi, & Buttyan 2015, Hausken 2006). The expected utility is a critical influencer in any decision-making process. For example, an operator invests in a particular security technology only after acknowledging that the investment will attain the expected returns (Hausken 2006). Similarly, the expected utility moderates attackers' strategic choices, especially the motivation (Hausken 2006, Herley 2010), behind attacks. The strategic choices of attackers are also influenced by available resources (Hausken 2006), the context of the interaction and the targeted operator's state which shapes the expected utility.

From an economic perspective, Herley (Herley 2010) pointed out that an attack strategy should be defined by the economics of attacks. He proposed attack strategies distinguishing attacks into scalable attacks and targeted attacks. In scalable attacks, the effort is independent of the number of users attacked. While in targeted attacks, the effort depends on per-user attacked suggesting targeted attacks must be on users with higher than average expected value. A profitable attack strategy involves accurately distinguishing viable from non-viable targets and deciding which viable target to attack based on the expected value (Herley 2012).

Grossklags et al. (Grossklags, Christin, & Chuang 2008) analysed the Nash equilibria and social optima for different classes of attacks and defences in weakest-link, best-sort, and sum-of-effort security games. They introduced a weakest-target game "where the attacker will always be able to compromise the entity (or entities) with the lowest protection level but will leave other entities unharmed." Florencio et al. (Florêncio & Herley 2013) refined this criterion by incorporating the concept of free-riding (discussed by (Varian 2004)) to the lowest protected entity (or entities) stating that even though there exist economically profitable targets, many attacks are extremely difficult to turn into profitable ones grounding it to the economics of attacks.

From an extensive literature review, Hausken and Levitin (Hausken & Levitin 2012) categorised attack tactics on plausible types of attacks such as attacking a single element, attacks against multiple elements, consecutive attacks, random attacks, attacks involving a combination of intentional and unintentional impacts, attacks with incomplete information, and attacks with variable resources. However, a critical difficulty in modelling opponents in general, specifically in the security domain, is due to lack of decisive information regarding potential adversaries and attackers-defenders interactions being highly complex and extensive (Pita, John, Maheswaran, Tambe, & Kraus 2012).

To understand the behavioural aspects of participants in cybersecurity, (Kusumastuti, Cui, Tambe, & John 2015) and (Ryutov, Orosz, Blythe, & von Winterfeldt 2015) have studied not only technical aspects but also psychosocial aspects through a three-player cybersecurity game. Kusumastuti et al. (Kusumastuti, Cui, Tambe, & John 2015) used mini-max solution to identify game parameters and their influence on a player's behaviour. Ryutov et al. (Ryutov, Orosz, Blythe, & von Winterfeldt 2015) aimed at understanding and modelling roles, motivations and conflicting objectives of players. In addition, (Anderson 2009), (Tambe, Jiang, An, & Jain 2014) and (Veksler & Buchler 2016) have demonstrated improvement in the predictability of attackers' behaviour by using behavi-

oural/cognitive modelling in a repeated security game environment.

To address adversary's bounded rationality, researchers have been pursuing alternative approaches. One approach includes robust optimisation techniques *avoiding adversary modelling* (Yang, Kiekintveld, Ordonez, Tambe, & John 2011, Pita, John, Maheswaran, Tambe, & Kraus 2012, Pita, Jain, Tambe, Ordóñez, & Kraus 2010), while the other approach incorporates human decision-making models for computing *defend strategies* (Nguyen, Yang, Azaria, Kraus, & Tambe 2013). Our work utilises the later approach and differs from the existing research by focusing on modelling adversaries rather than defenders. Firstly, instead of strictly bounded rationality of attackers, we consider attackers with strategic incentives working towards maximising their expected utility. More precisely, we pose a decision model with an intention to understand the mentality of attackers and their decision-making behaviour. We also introduce a generalised attack framework distinguishing the effort required during an attack process. The attack framework is used to evaluate and refine attack strategies. In addition, this framework also facilitates a way of addressing the abstract states of the decision model, which we believe applies to a whole class of security scenarios.

3 BEHAVIOURAL ANALYSIS AND FRAMEWORK CHARACTERISATION

We define the attackers-operators interaction as a game-theoretic model that captures essential characteristics of strategic decision making. The essence of game theory is to study factors influencing behaviour by reasoning what players think other players will do. However, in reality, having complete and perfect information regarding your opponents is never feasible. This applies particularly in the context of security, where the threat is almost always unknown and effectiveness of security investments are very hard to quantify (Laszka, Felegyhazi, & Buttyan 2015). So, every interaction is considered to involve certain degrees of uncertainty in committing to decisions. For example, attackers might have knowledge regarding an operator's investment in security but have no decisive information related to the extent an operator has invested or on what kinds of secure technologies has the operator invested in.

Attackers, alike defenders, being a deficit in resources (Florêncio & Herley 2013) have to act strategically maximising their expected gain and optimising their investment of resources. Aiming this, we analyse the attackers assuming their end goal is to successfully attain the expected results while minimising their effort in the attack process. First, this assumption alleviates the strictly bounded rationality of attackers and facilitate them with diverse attack strategies in contrast to the only choices of attacking or not attacking in traditional game-theoretic modelling approaches. In addition, it supports analysing interaction environments under conditions when attackers do not react - ignore or watch the target; diverging from the traditional approach where attackers follow a prescribed path of invariably attacking.

Introduction of strategic attackers expands the possibilities where an action can bear latent objectives and motives raising concerns regarding the admissibility of proposed defence strategies. To take a concrete example, consider a case of distributed denial-of-service (DDoS) attack, where an attacker attempts to prevent an operator from delivering information or services. With a strategy-less attacker, the resultant action for the operator would be to invest resources in countering the attack with full capacity to minimise the damage. The attacker being strategy-less an attack would precisely be an attempt to harm the existing state of the operator. However, for a strategic attacker, the DDoS attack might merely be a probing attack to assess the strength of the operator or it can be a diversion ahead a powerful targeted attack.

Figure 1 presents a glimpse of an extended action set for strategic attackers. However, further characterisation of attacks based on the severity of attacks and dependencies between attacks are beyond the scope of this paper. Additionally, we consider investment in security as discrete (Grossklags, Christin, & Chuang 2008, Kunreuther & Heal 2003, Lelarge & Bolot 2008), providing insulation towards all forms and degrees of attack, with no further distinctions in their capabilities to defend specific attacks.

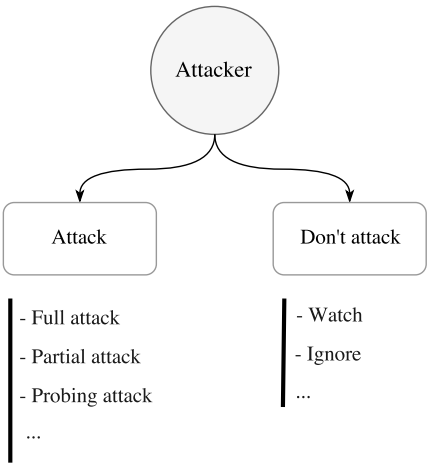

Figure 1: Attacker's decision space

Extracting the intent and decisions made to achieve the intended results from an instance of a highly restricted interaction scenario achieved through classical game theory is challenging. Indeed, it is the very problem in decisively predicting the behaviour of human players administring strictly bounded rationality (Camerer, Ho, & Chong 2004), especially while

addressing human adversaries (Camerer 2011) where there is no evidence on generated forms of motivation and intention behind attacks. One approach to tackle this problem is by understanding the context of interaction, as a decision must be made within a specific context and can be best represented through a hierarchy of decision states (Lewis 2013). Figure 2 is a hierarchical decision analysis tree capturing the mentality of attackers. The lowest level of the hierarchy represents concrete actions or choices of an attacker. As we ascend the hierarchy, states become increasing abstract and can be further fragmented into transitional stages precisely representing and supplementing an interaction scenario.

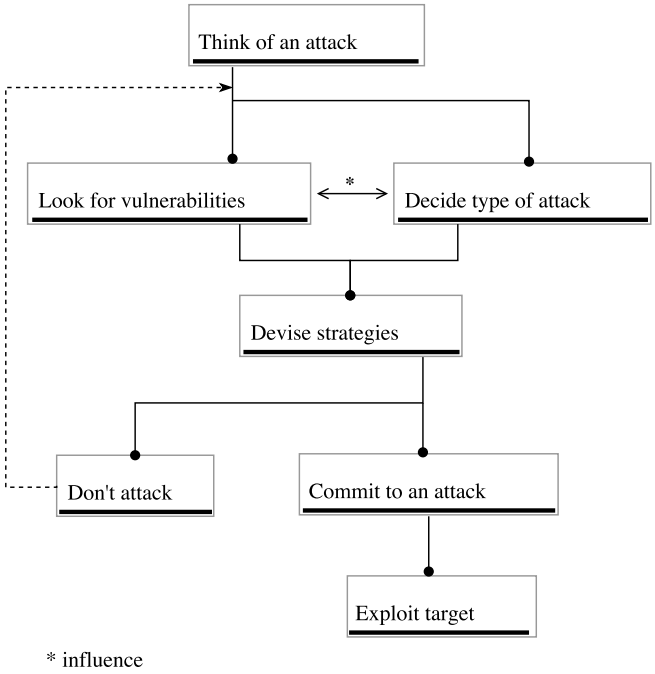

Figure 2: Attacker's Decision Analysis

The hierarchical decision analysis tree is a highly simplified decision model representing cognitive workflow of attackers, initiating from the thought of an attack and terminating on a definitive decision on attacking or not attack, replicating an attack process. The thought of an attack is supported by deciding on whether to search for vulnerabilities or the type of attack to perform within the attacker's capabilities. Based on the context of interaction there could be numerous other decision paths to choose from for attackers. These intermediate choice of paths are latently, or innately, or precisely influenced by factors backing the intended goal. The subsequent steps down the hierarchy include designing attack strategies and then deciding whether to commit to an attack. The low-level decisions which demonstrate certain behaviour are being modelled using game theory. Game theory being a mathematical modelling tool supplements in determining and quantifying elements influencing decisions and assist in predicting behaviour (Burke 1999). The low-level behaviour can be used to infer higher-order objectives that are likely driving such behaviour augmenting our understanding of the intention behind attacks. An improved understanding of intention and motivation will support rigid estimations of attackers' behaviour. However, understanding the abstract states - top levels of the hierarchy demands a multi-disciplinary approach with effective application of concepts from behavioural psychology and cognitive science.

We have characterised the attack process into different efforts required in performing an attack, acknowledging attackers to be rational entities. Figure 3, presents the attack framework demonstrating the efforts required in the attack process. The overall effort required can be broadly divided into *searching effort* and *breaking-in effort*. The searching effort includes efforts required in searching victims (targets), gathering information and searching vulnerabilities to exploit. Breaking-in effort represents the efforts required to compromise a system after choosing a target and a vulnerability to exploit. Based on the total effort required to compromise the target, an expected value can be derived. The expected value is one of the crucial factors moderating an attacker's decision (Herley 2010, Laszka, Felegyhazi, & Buttyan 2015).

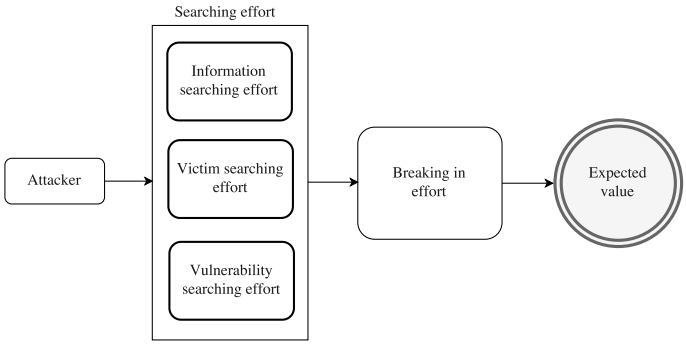

Figure 3: Attack Framework

The conversion of attackers' decision model into efforts required in the attack process disintegrates the top abstract levels of the decision model into modellable units. These modellable units can be used to quantify the expected utilities revealing the incentives behind attacks offering a better understanding of attackers' decision-making behaviour. Additionally, the attack framework assists in evaluating and enhancing attack strategies by effectively regulating the efforts strengthening the efficacy of attacks and ensuring better profits.

4 OPTIMISING ATTACK STRATEGIES

Lack of complete and perfect information against target induces uncertainty in attackers' decisions. The attack process eventually converges to a point of choice where an attacker has to decide on whether to attack or not to attack. The fate of an attempted attack depends on the target's capabilities to defend against

attacks which further depends on the extent of security investments. Figure 4, represents the expected payoffs of an attack against a targeted operator. We consider the investment in security as discrete - successfully preventing all forms and degrees of attacks. In the Figure 4, secure represents a system capable to successfully defend an attack and insecure represents the alternate.

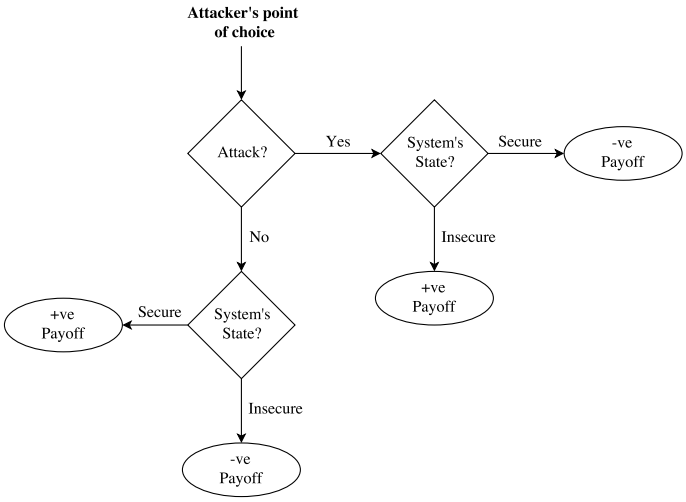

Figure 4: Attacker's expected payoff

In reality, defenders outnumber attackers. A critical problem attackers face is identifying targets such that a committed attack would yield something. For example, let's say the telecommunications domain has a total of N operators with N_c cooperating operators sharing security dependencies and N_n non-cooperating operators competing against security. It is an extremely expensive task for attackers to choose a viable target from such an intertwined mesh of operators. Here, the number of systems under each operator is ignored as considering it magnifies the complexity by many folds. Possible tactics attackers might adopt addressing this situation are

1. To randomly choose an operator and try breaching through the operator's defences. This approach would involve a heavy searching effort and a heavy breaking-in effort. This approach adds additional uncertainty as the attacker is unsure regarding his capabilities in successfully compromising the target.

2. To search for a certain type of vulnerability and then trying breaching it. This approach would involve heavy searching effort but a small breaking-in effort. Even though this approach involves high searching effort, chances of successfully compromising the chosen operator are very high.

The expected Utility (U) represents the probable payoff an attacker will receive on attacking the chosen operator. Based on the attack framework in Figure 3, the expected Utility for an attack can be determined as

$$U = cost(Information\_searching + Target\_searching + Vulnerability\_searching + Breaking\_in) - expected\,Value$$

where from (Herley 2010),

$$cost(Information\_searching) < cost(Target\_searching)$$

and any other forms of relationships cannot be defined from the existing literature.

Gathering and sharing of security-related information is a key factor heightening cybersecurity in both cooperating (Hausken 2017) and non-cooperating (Khouzani, Pham, & Cid 2014) environments. However, it is a known fact that the proposed information by defenders supports attackers in strategic decision-making. The following analysis illustrates how commonly available knowledge on operators can be used to reduce the cost of an attack. The use of available information reduces the information-searching effort to a static cost, represented as C_i, rather than a variable cost. In addition, attackers must bear the vulnerability-searching costs, represented as C_v, as a common cost irrelevant to any choice of target. t represents the choice of a target from the set of operators.

In a cooperating environment, the state of an operator is not only influenced by his decision but also by other cooperating operators' decisions. An attacker knowing that a set of operators (N_c) are cooperating refines the target-searching scope from N operators to N_c operators, where $N_c < N$, reducing the effort to an extent. The expected Utility (U_c) for attacking cooperating operators can be defined as

$$U_c = C_i + cost(Breaking\_in) + C_v \sum cost(Target\_searching) N_{c(t,-t)} - expected\,Value$$

Whereas, in a non-cooperating environment, an operator's security investment might encourage competing operators to invest in better security measures. On the other hand, it might also increase the likelihood of attacks on competing operators as the attacker will prefer a victim will lower resistance. Knowing operators are competing reduces the victim-searching effort considerably, as it is economically beneficial to attack the losing operator. Reduced victim-searching

effort can facilitate in allocating additional resources for vulnerability-searching and for breaking into the operator's defences. The expected Utility (U_n) for attacking competing operators can be defined as

$$U_n = C_i + cost(Breaking\_in) + C_v \sum cost(Target\_searching) N_{n(t,-t)} - expected\,Value$$

Desired Gain represents the amount of gain the attacking wants from an attack. From an economic perspective, an attacker would prefer the attack that maximises his desired Gain. That is, from the available range of attacks which would successfully compromise the found vulnerability, he chooses the attack which max($U-$ expected Gain). This indicates co-existence of several classes of attacks on a point of attack. The expected payoff and the desired gain from an attack would moderate the decisions of the attacker. As

$$\text{decision} = \begin{cases} Attack, & \text{if } U \geq \text{ expected Gain} \\ Donot\ attack, & \text{if } U < \text{ expected Gain} \end{cases}$$

5 CONCLUSION AND FUTURE WORK

We investigated cybersecurity environment from attackers' perspective. Our results show that taking into consideration and admitting that attackers have strategies, incentives etc, implies that defenders (telecoms operators in our studies) need to change how they perceive, defend and react to attackers. The implications given the rise of targeted/coordinated attacks versus uncoordinated attacks (eg: DDoS) mean that operators must significantly reassess their investment in security technologies towards the former, despite the latter having better 'security theater'[1].

Traditionally, security games have been analysed **assuming the attackers to be of bounded rationality with limited set of prescribed choices**. Furthermore, studies consider attackers do aim to maximise their expected gain and this consideration contradicts the assumption of bounded rationality of the attacker. We study the attackers considering they **share similar characteristics as defenders** with attack strategies maximising their expected gain.

In particular, we model security interactions with an extended set of actions available to the attackers. This expands the possibilities where an action can bear latent objectives and motives raising concerns regarding the admissibility of proposed defence strategies. We present a hierarchical decision tree capturing the mental model of attackers during the attack process. The decision model is supported by a generalised framework representing the attack process in terms of efforts required by the attackers addressing the abstract levels of the decision model. Using this framework attack strategies against cooperating and competing operators are derived optimising attackers' effort resulting in a better gain. Furthermore, it facilitates a way of understanding the strategic decision-making abilities of attackers.

Not all attacks are intended towards achieving economic targets. A novice attacker might not aim to maximise his economic payoff rather aim in gaining experience, or reputation and the interaction might end on an attempted attack. However, it might be a completely different picture for an experienced attacker. When such personality traits of the attackers are considered, specifically the strategic option of not attacking, it unsettles the traditional security modelling approach, particularly the Stackelberg approach (Kar, Nguyen, Fang, Brown, Sinha, Tambe, & Jiang 2017), where the game proceeds with the assumption that the attacker acts (invariably attacks). This raises a number of research questions challenging the traditional approach used in modelling cybersecurity. For example

- Is every interaction between an attacker and defender a repetitive process or is it a single-point interaction which ends on an attempted attack?
- Is using Stackelberg Security Games to model security interactions an appropriate choice?

Considering the economics of scalable and targeted attacks discussed by (Herley 2010), would it be an effective strategy to launch a small scalable attack to determine the strength of the target and then launch a specific attack incapacitating the target?

In (Kusumastuti, Cui, Tambe, & John 2015), attackers are facilitated with an option to not attack and invest in enhancing their capabilities enabling in launching stronger attacks in the future. This consideration would reduce the breaking-in effort and vulnerability searching effort. From a psychological perspective, Rogers (Rogers 2000) classified hackers depending on their expertise (from novice to experienced), areas of interests (software, hardware, etc.) and behavioural patterns.

Modelling players to be able to predict expected behaviour in a more realistic way requires a profound understanding of their incentives, motives and the context of the interaction. Behavioural modelling of the attacker will not only assist in understanding the expected intentions and behaviour of attackers but will also assist in devising comprehensive defences against such characteristics of attacks.

[1] Bruce Schneier - Beyond Security Theater: `https://www.schneier.com/essays/archives/2009/11/beyond_security_thea.html`

However, each instance of an interaction is unique, with a unique set of parameters characterising and moderating it. Modelling these unique interactions under common grounds is highly ineffective. They demand to be modelled based on the context of the interaction and using only game-theoretic concepts restricts the context and the interaction environment to a larger extent through biases, heuristics, and convenience.

This preliminary exploration will guide our future studies in aptly modelling behavioural aspects of attackers and in refining the attack strategies and characteristics of attackers by incorporating proposed concepts from the existing research work. This would further aid in comprehensively modelling the behavioural aspects of the participants in the context of information-cyber security.

6 ACKNOWLEDGEMENTS

The work was made in conjunction with the EU SCOTT and SECREDAS Projects and has received funding from the Electronic Component Systems for European Leadership Joint Undertaking under grant agreement No 737422.

REFERENCES

Altman, E., T. Boulogne, R. El-Azouzi, T. Jiménez, & L. Wynter (2006). A survey on networking games in telecommunications. *Computers & Operations Research 33*(2), 286–311.

Anderson, J. R. (2009). *How can the human mind occur in the physical universe?* Oxford University Press.

Anderson, R. & T. Moore (2006). The economics of information security. *Science 314*(5799), 610–613.

Burke, D. A. (1999). Towards a game theory model of information warfare. Technical report, AIR FORCE INST OF TECH WRIGHT-PATTERSONAFB OH.

Camerer, C. F. (2011). *Behavioral game theory: Experiments in strategic interaction*. Princeton University Press.

Camerer, C. F., T.-H. Ho, & J.-K. Chong (2004). A cognitive hierarchy model of games. *The Quarterly Journal of Economics 119*(3), 861–898.

Florêncio, D. & C. Herley (2013). Where do all the attacks go? In *Economics of information security and privacy III*, pp. 13–33. Springer.

Franklin, J., A. Perrig, V. Paxson, & S. Savage (2007). An inquiry into the nature and causes of the wealth of internet miscreants. In *ACM conference on Computer and communications security*, pp. 375–388.

Gordon, S. (1994). The generic virus writer. In *Proc. Intl. Virus Bulletin Conf*, pp. 121–138.

Gordon, S. (2000). Virus writers: The end of the innocence? In *10th Annual Virus Bulletin Conference (VB2000), Orlando, FL*.

Grossklags, J., N. Christin, & J. Chuang (2008). Secure or insure?: a game-theoretic analysis of information security games. In *Proceedings of the 17th international conference on World Wide Web*, pp. 209–218. ACM.

Hausken, K. (2006). Income, interdependence, and substitution effects affecting incentives for security investment. *Journal of Accounting and Public Policy 25*(6), 629–665.

Hausken, K. (2017). Security investment, hacking, and information sharing between firms and between hackers. *Games 8*(2), 23.

Hausken, K. & G. Levitin (2012). Review of systems defense and attack models. *International Journal of Performability Engineering 8*(4), 355–366.

Herley, C. (2010). The plight of the targeted attacker in a world of scale. In *WEIS*.

Herley, C. (2012). Why do nigerian scammers say they are from nigeria? In *WEIS*.

Hota, A. R. & S. Sundaram (2015). Interdependent security games under behavioral probability weighting. In *International Conference on Decision and Game Theory for Security*, pp. 150–169. Springer.

Jiang, L., V. Anantharam, & J. Walrand (2008). Efficiency of selfish investments in network security. In *Proceedings of the 3rd international workshop on Economics of networked systems*, pp. 31–36. ACM.

Kar, D., T. H. Nguyen, F. Fang, M. Brown, A. Sinha, M. Tambe, & A. X. Jiang (2017). Trends and applications in stackelberg security games. *Handbook of Dynamic Game Theory*, 1–47.

Khouzani, M., V. Pham, & C. Cid (2014). Strategic discovery and sharing of vulnerabilities in competitive environments. In *International Conference on Decision and Game Theory for Security*, pp. 59–78. Springer.

Kunreuther, H. & G. Heal (2003). Interdependent security. *Journal of risk and uncertainty 26*(2-3), 231–249.

Kusumastuti, S., J. Cui, A. Tambe, & R. S. John (2015). A behavioral game modeling cyber attackers, defenders, and users. Research paper presented at the AAAI Spring Symposium, Stanford University, Palo Alto.

Laszka, A., M. Felegyhazi, & L. Buttyan (2015). A survey of interdependent information security games. *ACM Computing Surveys (CSUR) 47*(2), 23.

Lelarge, M. & J. Bolot (2008). A local mean field analysis of security investments in networks. In *Proceedings of the 3rd international workshop on Economics of networked systems*, pp. 25–30. ACM.

Lewis, M. J. (2013). Hierarchical decision making. In *STIDS*, pp. 162–165.

Liang, X. & Y. Xiao (2013). Game theory for network security. *IEEE Communications Surveys & Tutorials 15*(1), 472–486.

Manshaei, M. H., Q. Zhu, T. Alpcan, T. Bacşar, & J.-P. Hubaux (2013). Game theory meets network security and privacy. *ACM Computing Surveys (CSUR) 45*(3), 25.

Merrick, K., M. Hardhienata, K. Shafi, & J. Hu (2016). A survey of game theoretic approaches to modelling decision-making in information warfare scenarios. *Future Internet 8*(3), 34.

Nguyen, T. H., R. Yang, A. Azaria, S. Kraus, & M. Tambe (2013). Analyzing the effectiveness of adversary modeling in security games. In *AAAI*.

Panaousis, E., A. Fielder, P. Malacaria, C. Hankin, & F. Smeraldi (2014). Cybersecurity games and investments: a decision support approach. In *International Conference on Decision and Game Theory for Security*, pp. 266–286. Springer.

Pita, J., M. Jain, M. Tambe, F. Ordóñez, & S. Kraus (2010). Robust solutions to stackelberg games: Addressing bounded rationality and limited observations in human cognition. *Artificial Intelligence 174*(15), 1142–1171.

Pita, J., R. John, R. Maheswaran, M. Tambe, & S. Kraus (2012). A robust approach to addressing human adversaries in security games. In *Proceedings of the 20th European Conference on Artificial Intelligence*, pp. 660–665. IOS Press.

Rogers, M. (2000). A new hacker taxonomy. *University of Manitoba*.

Roy, S., C. Ellis, S. Shiva, D. Dasgupta, V. Shandilya, & Q. Wu (2010). A survey of game theory as applied to network security. In *System Sciences (HICSS), 2010 43rd Hawaii International Conference on*, pp. 1–10. IEEE.

Ryutov, T., M. Orosz, J. Blythe, & D. von Winterfeldt (2015). A game theoretic framework for modeling adversarial cyber se-

curity game among attackers, defenders, and users. In *International Workshop on Security and Trust Management*, pp. 274–282. Springer.

Sun, W., X. Kong, D. He, & X. You (2008). Information security investment game with penalty parameter. In *Innovative Computing Information and Control, 2008. ICICIC'08. 3rd International Conference on*, pp. 559–559. IEEE.

Tambe, M., A. X. Jiang, B. An, & M. Jain (2014). Computational game theory for security: Progress and challenges. In *AAAI spring symposium on applied computational game theory*.

Varian, H. (2004). System reliability and free riding. In *Economics of information security*, pp. 1–15. Springer.

Veksler, V. D. & N. Buchler (2016). Know your enemy: Applying cognitive modeling in security domain. In *Proceedings of the 38th Annual Conference of the Cognitive Science Society*, pp. 2405–2410.

Yang, R., C. Kiekintveld, F. Ordonez, M. Tambe, & R. John (2011). Improving resource allocation strategy against human adversaries in security games. In *IJCAI Proceedings-International Joint Conference on Artificial Intelligence*, Volume 22, pp. 458.