Strategic Discovery and Sharing of Vulnerabilities in Competitive Environments

MHR. Khouzani, Viet Pham, Carlos Cid

Information Security Group, Royal Holloway, University of London
arman.khouzani@rhul.ac.uk,viet.pham.2010@live.rhul.ac.uk,carlos.cid@rhul.ac.uk

Abstract. We investigate the incentives behind investments by competing companies in discovery of their security vulnerabilities and sharing of their findings. Specifically, we consider a game between competing firms that utilise a common platform in their systems. The game consists of two stages: firms must decide how much to invest in researching vulnerabilities, and thereafter, how much of their findings to share with their competitors. We fully characterise the Perfect Bayesian Equilibria (PBE) of this game, and translate them into realistic insights about firms' strategies. Further, we develop a monetary-free sharing mechanism that encourages both investment and sharing, a missing feature when sharing is arbitrary or opportunistic. This is achieved via a light-handed mediator: it receives a set of discovered bugs from each firm and moderate the sharing in a way that eliminates firms' concerns on losing competitive advantages. This research provides an understanding of the origins of inefficiency and paves the path towards more efficient sharing of cyber-intelligence among competing entities.

1 Introduction

Businesses across different sectors of the economy, from telecommunication and finance to energy, healthcare and transportation, increasingly rely on cyberspace and IT services. Past incidents of cyber-attacks and consequent damages have left little doubt in the minds of business managers and policy makers about the importance of investment in cybersecurity. Gathering and exchange of security intelligence are identified as key factors in enhancing the effectiveness of cybersecurity measures. Steps have been taken by governments to provide the environments to galvanise and coordinate the exchange of cybersecurity information: UK launched the "Cyber Security Information Sharing Partnership" [1] after a pilot program in 2011/12 as a "joint, collaborative initiative between industry and government to share cyber threat and vulnerability information in order to increase overall situational awareness of the cyber threat". In the US, the "National Coordinating Center for Communications (NCC)" acts as the "Information Sharing and Analysis Center (ISAC)" for telecommunication [2].

While "Information Sharing and Analysis Centers (ISACs)" – such as Information Technology (IT)-ISAC and Financial Services (FS)-ISAC – can provide the platform for exchange of cyber-intelligence, the role of incentives must not be

ignored. Providing the means of communication in the presence of strategic and competing profit-maximizing entities does not necessarily lead to exchange of their cybersecurity information. In order to understand the incentives of firms in creating and sharing information security knowledge, it is important to identify the distinct nature of the security information being shared. Some example categories of the type of cyber-intelligence to be shared are: (a) steps, protocols and measures a firm has taken to improve its security; (b) past incidents of successful or unsuccessful attacks and the resulting privacy, intellectual property and financial losses; and (c) discovered security vulnerabilities. Sharing each of these types of information have specific incentive implications. For instance, "public disclosure" of security breach incidents can harm the consumers and investors' confidence and lead to a statistically significant decreases in the market value of firms [3–5]. In this paper, we particularly focus on the third type of information: sharing discovered security vulnerabilities, or *bugs* for short.

From the societal point of view, sharing knowledge of security vulnerabilities among firms is a positive move: it improves the overall efficiency of bug discovery efforts. It moreover enhances the cyber protection of an entire industry against future attacks by reducing the common exploitable threats. It is often the case that different organizations of an economic sector bear similar vulnerabilities in their information systems [6]. This is partly due to the adoption of common implementations, libraries or operating systems. For instance, the `Heartbleed` bug (formally, CVE-2014-0160), a buffer-over-read vulnerability in the `OpenSSL` cryptographic library exposed in April 2014, affected around half a million certificates issued by trusted certificate authorities [7]. Another reason why different technological companies face common threats is the incorporation of discovered vulnerabilities into hacking toolkits which enables even less sophisticated users to configure the same malware to attack across different organizations [6].

Recognizing the need for cyber-protection, companies may invest in finding their security vulnerabilities. These can be "bugs" for example in their application level software, operating system or implementation of a network protocol, which we will hence generically refer to as the common *platform*. No company knows exactly how many bugs there are in a software they are using. More investment and effort in security research increases the chances of discovering them, but there is always a factor of luck involved. Each company patches and rectifies the vulnerabilities it finds, which is usually the much easier part than finding them in the first place. Each bug that is not discovered by a company, and hence not rectified, is potentially exploitable by cyber-attackers.

When a bug is indeed successfully exploited, the victim suffers direct losses. These can include outage of their services, recovery costs, losses of important data, user compensation, legal fines, etc. However, a company may also be affected by incidents of cyber-attacks on other companies in that economic sector: On one hand, the whole sector of the economy may suffer a blow: as customers may lose confidence in the whole "service" and seek alternative "safer" means. For instance, if one or a few major online banking companies fall victim to a cyber-attack, then some customers may lose confidence in the whole sector and switch to traditional banking altogether. Moreover, investors and stock holders

may too lose confidence in the whole industry in favour of alternative options for investment. These two effects translate to a net market value loss of the whole sector, which bites all of the companies upon a successful attack on anyone of them. However, on the other hand, if (and once) a bug is exploited in competitor(s) that a company has discovered before (and has hence taken care of), it can have the opposite effect of boosting the confidence of customers as well as the investors: customers may switch to use and investors redirect their capital to the "safer" company. In other words, discovering a bug in a common software may give a company a "competitive edge" compared to others.

The two effects work in the opposite direction of each other in terms of incentives for sharing the found vulnerabilities. The sharing strategies, in turn, affect the investment decisions to discover the bugs in the first place: On the one hand, sharing information translates to a more effective discovery process and hence encourages investment, as the findings of one company is fortified by another's since the process of finding the bugs is probabilistic in nature. But on the other hand, there can be a tendency of free-riding on the discovery investment of other companies and hence get away with less investment. Further complicating the problem is the presence of uncertainty and information asymmetry: companies ought to make their discovery investment decisions in the face of uncertainties about the total number of bugs, and they need to make decision about sharing of their findings not knowing the number of findings of the other company.

Contributions of this paper are as follows: In Section 2, we model the interdependent vulnerability research investment and information sharing decisions of two strategic and competing firms as a two stage Bayesian game. We fully determine the Perfect Bayesian Equilibria of the game in closed-form in Section 3. Specifically, in Subsection 3.1, we derive the Bayesian equilibrium strategies of the firms about sharing of their finding for a given investment pair, and given their findings. In particular, we establish that the sharing strategies are unique and dominant, and are in the simple forms of "full-sharing" or "no sharing", completely determined by the competitive nature of the security findings. In Subsection 3.2, we derive the investment strategies of the firms knowing their subsequent sharing strategies. We show how "full sharing" leads to free-riding and inefficiently low investments. Also how "no sharing" is socially inefficient by preventing mutual benefit of sharing, double-efforts and potential over-investment. Finally, in Section 4, we provide a light-weight mediation mechanism free of monetary-transfers that enable (partial) sharing of the information when the firms fail to achieve any sharing on their own.

Comparison to Literature: Information sharing in the context of cybersecurity is investigated in papers like [8–15]. These works build on microeconomic models of information sharing in a general oligopoly (e.g. [16–18]) where the effect of information sharing is captured as improvement in the efficiency of production, i.e., reducing the marginal cost, or improving demand, or both. A common feature of the models is that there is no specification of the type of security information to be shared. The decision of how much information to share is modelled as a normalized continuous variable between zero and one, zero corresponding to no

sharing and one corresponding to full sharing. In contrast, we specifically model the information as the discovered security vulnerabilities by each player, and hence, the sharing decisions in our model is the "number" of bugs to be shared. In addition, the relation between security investments and information sharing is rather loose in the previous literature. For instance, the effective amount of shared information is heuristically chosen as the product of the investment decision and sharing decision. In contrast, we specifically model the process of investment for "generation" of security information and subsequently, sharing of them. Moreover, we develop a mediation monetary-free mechanism that enables sharing in the face of competition as a novel contribution. More distantly, this work is related to research on R&D rivalries, e.g. [19], with at least one major difference that vulnerability discoveries are inherently not patentable.

2 Model

Our model considers a game between firm i and firm j where each decides how much to invest in security research on a common "platform", and subsequently how many of their found security vulnerabilities to share with the other. The platform has an unknown number of security vulnerabilities, or "bugs", which, if not discovered and rectified, may be exploited with ramifications for both firms. Before the game starts, the nature determines the total number of bugs following some distribution. Let the random variable representing the total number of bugs be B with the sample space of $\mathbb{N}^{+1}$ and known mean value λ. The realisation of B is not observed by any of the firms. The game play consists of two stages: *investment* and *sharing*, as described in the following:

1- Investment: In this stage, the players, while unaware of the total number of bugs in the platform, "simultaneously" decide how much to invest in bug discovery, and make it publicly known. Note that simultaneous move in the context of game theory just implies that neither one of the players can assume pre-commitment to a decision by the other players. A player's investment c determines the probability $p \in [0, 1)$ that each bug is discovered. For simplicity, we assume that the bugs are homogeneous, in that they are equally difficult to discover. Moreover, we assume discovery of each bug is independent across the bugs and across the firms. The research investment c and discovery probability p are related through function π as $p = \pi(c)$, with $\lim_{c\to\infty} \pi(c) = 1$. We naturally assume that $d\pi(c)/\,dc > 0$, as well as $d^2\pi(c)/\,dc^2 \leq 0$: The chance of finding bugs should be improved with more investment, and it is increasingly more difficult to improve the success of bug discoveries. In general we assume that the two firms have distinct cost-probability relations, denoted as $\pi_i(c)$ and $\pi_j(c)$. Because we assume both π_i and π_j are strictly incising, there is a one-to-one mapping between investment and discovery probability. Indeed, $c_i = \pi_i^{-1}(p_i)$ and $c_j = \pi_j^{-1}(p_j)$. Hence, we can equivalently represent each player's strategy in this stage by its choice of discovery probability, i.e., p_i and p_j.

[1] We adopt the convention that random variables are denoted by capital letters and their realisations by lower case. Also, $\mathbb{N}^+ := \mathbb{N} \cup \{0\}$.

2- Sharing: After investments are made, each player privately and independently "discovers" some bugs in the platform. Subsequently, each decides how many of its findings to share with the other. Note that the discoveries are not part of the strategies of the players and is rather determined probabilistically –by "nature"– once the investments are made. Since the discoveries are private, they cause an "incompleteness of information" of players about each other. We therefore model this sharing decisions as a Bayesian game. Firms i and j respectively discover N_i and N_j bugs in the platform, which are random variables with the common sample space of $\{0, 1, \ldots, B\}$.[2] The set of discovered bugs may have an overlap, i.e., some identical bugs may be discovered by both firms. We denote the number of common bugs by N_{ij}. The sample space of N_{ij} is $\{0, 1, \ldots, \min(N_i, N_j)\}$. Given the total number of bugs B and investment levels c_i and c_j, the nature determines the number of bugs discovered by each firm and the number of commonly discovered bugs N_i, N_j and N_{ij}. The quadruple (B, N_i, N_j, N_{ij}) is the random variable over the set of possible "states of the world" Ω. Note that due to the revelation of investments at the end of the first stage, the probability distribution of (B, N_i, N_j, N_{ij}) over Ω is publicly known. For each nature state $(b, n_i, n_j, n_{ij}) \in \Omega$, firm i (resp. j) observes n_i (resp. n_j), i.e., the number of bugs it has discovered, as its "type". For each realisation of the number of found bugs and announced investments, a firm must decide how many of its found bugs to share with the other. Due to the homogeneity assumption of bugs, the bugs to be shared can be assumed to be picked uniformly randomly. A (pure) strategy of firm i is thus a mapping $s_i(p_j, n_i) : [0, 1] \times \mathbb{N}^+ \to \mathbb{N}^+$ such that $s_i(p_j, n_i) \leq n_i$.[3] Let $\sigma_i = (p_i, s_i)$ denote the pure strategies of player i for *the whole game*. After both σ_i and σ_j are decided, the the overall utilities of each player is determined as the result of its investment together with the expected losses/gains from security incidents.

In what follows, we describe the expected utility of the two players after two stages of actions. We assume risk-neutral players, that is, the players care equally about their utility of expected outcome and their expected utility. Hence, the utilities are linear sums of the (negative of the) expected costs per each bug minus the investment cost for discovery of the bugs. Note that at the time of taking the decision about sharing the discovered bugs, the investments for discovering the bugs are "sunk" costs, i.e., they are already spent and will not affect the cost to go of different actions to take. Each bug, if not discovered by or informed to a player, will be successfully exploited on that player by attackers with a probability, which without loss of generality, we take to be one. We assume that the exploitation probabilities and the severity of bugs are homogeneously distributed. For each bug there are three types of losses/damages:[4]

- **Direct loss** $\boxed{l > 0}$: affecting only the compromised firm (e.g. outage/denial of its services, compromise/corruption of its data, etc.).

[2] By $\{0, 1, \ldots, B\}$, it is meant that given the realisation $B = b$, the set is $\{0, 1, \ldots, b\}$.
[3] Since p_i is part of player i's strategy, it needs not be included as an argument to s_i.
[4] For simplicity of exposition, we assume the losses and damages are symmetrical; it is straightforward to generalise the results to non-symmetric cases.

Table 1: List of main notations

Parameter	Definition
B, b	Random variable for the total number of bugs, and a realisation
N_i, n_i	Random variable for the number of bugs discovered by i, and a realisation
N_{ij}	Random variable for the number of common bugs discovered by both
a_i	Action of player i: how many discovered bugs to share
λ	Expected number of the total number of bugs
p_i, p_j	Probability that each bug is discovered by player i,j
u_i, u_j	Expected utilities of player i, j
c_i, c_j	Discovery investment cost of player i,j
l	Direct loss upon exploitation of an (undiscovered) bug by attackers
δ	Loss (gain) in utility of the player who is the only one attacked (not attacked) – capturing the market competition effect
τ	Loss in utility of both players if a bug is exploited in either one of them – capturing the total market section shrinkage effect
$p = \pi(c)$	The relation relating the level of investment c to the discovery probability of a bug p. In this paper, we use $p = \pi(c) = 1 - e^{-\theta c}$.

- **Market shrinkage** $\boxed{\tau \geq 0}$: the common loss as a result of a successful attack that affects both, even the firm that is not compromised. This is the effect of the market shrinkage after a successful attack as a result of a portion of both demand and investment moving away from (abandoning) the whole service/technology in favour of "safer" alternatives, or simply relinquishing that sector altogether.
- **Competitive loss** $\boxed{\delta \geq 0}$: when *only one* firm is compromised by attackers, the compromised firm loses δ while the other gains δ. This represents the shifting of demand and/or public investment (stocks) upon a successful attack.

Given the notions described above, there are four possibilities of net cost for each bug that a player may incur: (a) The bug is known by both players (either through own discovery or through the information shared by the other firm). In this case, the utility of the players is $(0, 0)$, as neither one of the players loses anything.[5] (b) The bug is known by player i, but not player j. In this case, the utility pair is $(\delta - \tau, -\delta - \tau - l)$: the bug will be exploited at firm j, which causes its direct loss l and a competitive advantage δ for firm i, while both of them will lose τ due to market shrinkage. (c) The bug is known by player j, but not player i. This is the mirror situation to case-b: the utility pair is $(-\delta - \tau - l, \delta - \tau)$. (d) The bug is known by neither one of the players. Here, there is no competitive advantage of one over the other, but there is still the market shrinkage effect, besides the direct losses to both. Hence, the utilities are $(-\tau - l, -\tau - l)$.

To facilitate the computation of the expected utilities, we define the following auxiliary random variables (as also depicted by a Venn diagram in Fig. 1): Let $B_{i,j}$, $B_{i,\neg j}$, $B_{\neg i,j}$ and $B_{\neg i,\neg j}$ represent the number of bugs that, respectively, both players, only player i, only only player j, and neither player knows about. Let the (expected) utility of players be denoted by u, which is a function from the strategy profile of the players and the state of the world to the set of real

[5] The assumption is that once the bug is discovered, its "fix" is immediate and costless.

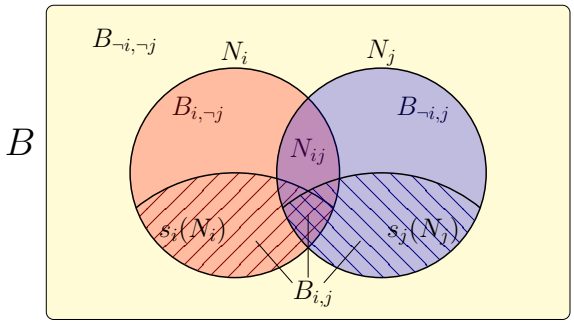

Fig. 1: Venn diagram illustration of the sets of bugs.

numbers. The expectation is taken with respect to the realisation of $B_{i,j}$, $B_{i,\neg j}$, $B_{\neg i,j}$ and $B_{\neg i,\neg j}$ given B, N_i, N_j and N_{ij}, and the sharing strategies. We are now ready to compute the expected utility of player i given a realisation of the state of the world $\omega = (b, n_i, n_j, n_{ij})$, and $\sigma_i = (p_i, s_i)$, $\sigma_j = (p_j, s_j)$:

$$u_i(\omega, \sigma_i, \sigma_j) = -c_i(p_i) + 0 \cdot \mathbb{E}(B_{i,j}) + (\delta - \tau) \cdot \mathbb{E}(B_{i,\neg j}) \\ + (-\delta - \tau - l) \cdot \mathbb{E}(B_{\neg i,j}) + (-\tau - l) \cdot \mathbb{E}(B_{\neg i,\neg j}) \quad (1)$$

In what follows we analyse further the structure of this utility function and derive the "outcome" of the game and study its properties.

3 Analysis of the Game

When dealing with strategic entities with inter-dependent utilities, investigating equilibria, most notably Nash Equilibria, is a method of predicting their decisions. Our game contains sequential moves, and thus an ordinary Nash equilibrium concept would potentially cause the problem of "non-credible threats". Also note that our game contains simultaneous actions in each stage, and hence is of "imperfect information". We therefore examine possible perfect Bayesian equilibria (PBE), a solution concept that effectively eliminates non-credible threats in sequential games with incomplete and imperfect information.

Informally, a PBE is a profile of strategies such that, given any belief about the game history that is consistent with that profile, the induced strategy profile must be a Nash equilibrium for the induced subgame (the game from the belief in an information Set onward). To find the set of PBEs, we notice that since the investment decisions are announced before sharing, each Bayesian game in the second stage is a proper subgame of the whole game. This means that we can use backward induction and first construct $((p_i, s_i), (p_j, s_j))$ such that s_i and s_j form a Bayesian Nash equilibrium (BNE) of the Bayesian game of sharing induced by choices of p_i and p_j. This in turn determines the utility of the players for each choice of (p_i, p_j), which allows us to build a simple strategic-form game with actions p_i and p_j corresponding to the first stage of the game. The remaining

task will be to find a Nash equilibrium for this game, which will lead to a proper PBE for the whole two-stage game. We thus proceed by studying the second stage of the game (information sharing), and then proceed to analyse players' investments given their equilibrium sharing strategies.

3.1 Second stage: Sharing the bug discoveries

To study the Bayesian game of the second stage, we first compute the utility functions of the players from the basic description in (1). Since $\mathbb{E}(B_{i,j})$ is multiplied by zero, we can safely ignore it. For the rest, we have:

$$\mathbb{E}[B_{i,\neg j}|\omega,\sigma_i,\sigma_j] = (n_i - n_{ij})(1 - \frac{s_i(p_j,n_i)}{n_i}) \tag{2a}$$

$$\mathbb{E}[B_{\neg i,j}|\omega,\sigma_i,\sigma_j] = (n_j - n_{ij})(1 - \frac{s_j(p_i,n_j)}{n_j}) \tag{2b}$$

$$\mathbb{E}[B_{\neg i,\neg j}|\omega,\sigma_i,\sigma_j] = b - n_i - n_j + n_{ij} \tag{2c}$$

In (2a),(2b), we have in part used the fact that the bugs to be shared are chosen uniformly randomly across the discovered bugs. Replacing in (1), we obtain:

$$\begin{aligned} u_i(\omega,\sigma_i,\sigma_j) = -c_i(p_i) + (\delta-\tau)(n_i - n_{ij})(1 - \frac{s_i(p_j,n_i)}{n_i}) + \\ (-\delta-\tau-l)(n_j - n_{ij})(1 - \frac{s_j(p_i,n_j)}{n_j}) + (-\tau-l)(b - n_i - n_j + n_{ij}) \end{aligned} \tag{3}$$

We are looking for strategy profiles (strategy pairs (s_i, s_j) in our two-player context) that are simultaneous best responses to each other, given the information that each player has, notably including its number of discovered bugs. In the Bayesian Nash equilibria of the game, each candidate strategy for a player must be a maximizer of its expected utility given the strategy of the other player and given its observed type (number of discovered bugs).[6] Formally, for a given p_i and p_j, we are looking for the strategy pairs (s_i^*, s_j^*), such that:

$$\forall n_i \in \mathbb{N}^+,\ s_i^*(p_j,n_i) \in \arg\max_{s_i(p_j,n_i)} \mathbb{E}[u_i(\omega,(p_i,s_i(p_j,n_i)),(p_j,s_j^*(p_i,n_j)))|n_i] \tag{4}$$

and simultaneously vice versa for j. Such pairs constitute the (pure) Bayesian Nash Equilibria of the second stage of our game. The pair (s_i^*, s_j^*) is further, a Dominant (pure) Bayesian Nash Equilibrium iff:

$$\forall n_i \in \mathbb{N}^+, \forall s_j,\ s_i^*(p_j,n_i) \in \arg\max_{s_i(p_j,n_i)} \mathbb{E}[u_i(\omega,(p_i,s_i(p_j,n_i)),(p_j,s_j(p_i,n_j)))|n_i] \tag{5}$$

and simultaneously vice versa for j. We are now ready to express the main result of this section:

[6] To analyse the game, each player must specify its actions for all of its possible types, and not just the realised (and observed) type. This is because, the expected utility of each player depends on the possible actions of the other player(s) weighted against their potential types, since the type of other player(s) are not directly observed.

Proposition 1. *Suppose $p_i, p_j < 1$. If $\delta < \tau$, the unique dominant pure Bayesian Nash Equilibrium of the second stage of the game is $(s_i^*(p_j, n_i), s_j^*(p_i, n_j)) = (n_i, n_j)$, i.e., sharing all the discovered bugs. If $\delta > \tau$, it is $(s_i^*(p_j, n_i), s_j^*(p_i, n_j)) = (0, 0)$, i.e., sharing no information at all. When $\delta = \tau$, any strategy pair becomes a Bayesian Nash Equilibrium. This proposition holds irrespective of the distribution of the total number of bugs.*

Proof. According to (5), a pair (s_i^*, s_j^*) constitutes a Dominant Bayesian Equilibrium if, for each type of a player, its corresponding action is the best (provided the knowledge of its type), irrespective of the strategy of the other player. From (3), the only term in the the expression of $u_i(\omega, \sigma_i, \sigma_j)$ that involves s_i is the second term: $(\delta - \tau)[(n_i - n_{ij})(1 - s_i(p_j, n_i)/n_i)]$. Hence, with the assumption of $p_j < 1$ in mind, the maximization of $\mathbb{E}[u_i(\omega, \sigma_i, \sigma_j)|n_i]$ with respect to $s_i(p_j, n_i)$ reduces to maximizing $(\delta - \tau)(1 - s_i(p_j, n_i))$, which yields the proposition.[7] □

Discussion The proposition makes intuitive sense: when $\delta > \tau$, each bug that is only known by a player wins it a strictly positive (expected) competitive gain of $(\delta - \tau)$, as the competitive shift in the demand and public investment outweighs the overall drop in the demand and fall in the stock market of the whole market section. Hence it rather not share any of its findings, irrespective of what the other player chooses. This is because the players have no means of making their decisions "contingent" on the decision of the other.[8] Similarly, when $\delta < \tau$, the competitive shift in the demand and capital, falls short of the whole market section shrinkage. Therefore, the players prefer to share all their findings to (selfishly) keep themselves from being hurt. Perhaps the surprising result is that the dominant strategy of the players turned out to be completely determined by the relative values of only two parameters δ and τ. This proposition fully determines the sharing strategy of the firms. Notably, aside from the special case of $\delta = \tau$, the equilibrium is unique and hence, there is no ambiguity in selection of the equilibrium. Next, we investigate how each firm invests for discovering the bugs knowing the subsequent sharing strategies.

3.2 First stage: Investment for bug discovery

In the first stage of the game, each player decides about its investment amount for the discovery of bugs, heeding the strategy of the other player in the second stage. To obtain closed-form results, we need to model the relation between

[7] Although the proposition leaves out the cases in which the condition $p_i, p_j < 1$ are not satisfied, they are not difficult to analyse: suppose $p_j = 1$, then $\mathbb{E}[(n_i - n_{ij})(1 - s_i(p_j, n_i)/n_i)|n_i] = 0$, and hence the expression for $\mathbb{E}[u_i(\omega, \sigma_i, \sigma_j)|n_i]$ will not depend on s_i at all. Hence, in any Bayesian Nash Equilibria, the choice of s_i becomes arbitrary. Similar situation happens for s_j when $p_i = 1$. Intuitively, if the other player "knows every bug for certain", then a player cannot affect its utility through its action: it cannot gain any competitive advantage if $\delta > \tau$, or help prevent market shrinkage when $\delta > \tau$. Note that realistically, we can safely assume $p_i, p_j < 1$, as no practical amount of investment leads to absolute certainly of finding all bugs.

[8] We will see in §4 how this situation can be altered in the presence of a mediator.

investment decision and the chance of finding bugs. A simple candidate for such relation is the following: $p = \pi(c) = 1 - e^{-\theta c}$, where θ represents a measure of the efficiency of the investment: a larger θ corresponds to a higher efficiency of the investment. As the level of investment increases to infinity, the probability of discovery of each bug asymptotically approaches unity. The two firms may be different in how "efficient" they are in their investment. A firm with more prepared talents can expect higher chances of discovery with less investment. To capture the potential heterogeneity in the investment efficiencies, we consider two potentially different θ_i and θ_j. Our investment-discovery probability relation has the extra property that the relative efficiency of the investment stays constant for all investment values, specifically: $(\partial\pi_i/\partial c)/(\partial\pi_j/\partial c) = \theta_i/\theta_j$. This relation can also be equivalently represented in its inverse form: $c_i(p_i) = -\ln(1-p_i)/\theta_i$ for $p_i \in [0, 1)$, and likewise for j. Note that the condition of Proposition 1 $p_i, p_j < 1$ is automatically satisfied when $\lim_{p\to 1} c(p) \to \infty$, as is the case in our example.

To analyse this stage, we note that Proposition (1) fully determines (s_i^*, s_j^*) for each profile of (p_i, p_j). This allows us to treat the first stage as a "one-shot" game of investment with action profiles of the form (p_i, p_j).

3.3 The case of $\delta < \tau$

For the case of $\delta < \tau$, from Proposition 1, the dominant strategy of both players is to share *all* of their findings, i.e., $s_i(p_j, n_i) = n_i$ and $s_j(p_i, n_j) = n_j$ for all $n_i, n_i \in \mathbb{N}^+$. Then, the second and third terms in (3) become zero, and we get:

$$\begin{aligned}\mathbb{E}[u_i(\omega, (p_i, s_i^*), (p_j, s_j^*))] &= -c_i(p_i) + (-\tau - l)\mathbb{E}[B - N_i - N_j + N_{ij}] \\ &= -c_i(p_i) + (-\tau - l)\lambda(1-p_j)(1-p_i)\end{aligned}$$

The best response p_i^{BR} as a relation over p_j is hence:

$$p_i^{BR}(p_j) = [c_i'^{-1}(\kappa(1-p_j))]^+,^9 \qquad \text{where } \kappa := \lambda(\tau + l). \tag{6}$$

Note that when $p^{BR} > 0$, $\partial p_i^{BR}/\partial p_j = -\kappa/c_i''(p_i^{BR}) < 0$, i.e., more investment by the other player leaves *less* incentive for a player to invest. Similarly, we have: $\mathbb{E}[u_i(\omega, (p_i, s_i^*), (p_j, s_j^*))] = -c_j(p_j) + (-\tau - l)(1-p_i)\lambda(1-p_j)$, and hence: $p_j^{BR}(p_i) = [c_j'^{-1}(\kappa(1-p_i))]^+$. The fixed points of the best response correspondence $(p_i, p_j) \rightrightarrows ([c_i'^{-1}(\kappa(1-p_j))]^+, [c_j'^{-1}(\kappa(1-p_i))]^+)$ constitute the outcome of the first stage. For our example cost function $c = -\ln(1-p)/\theta$, the simultaneous best response must hence satisfy the following (Fig. 2a):

$$p_i^{BR}(p_j) = [1 - \frac{1}{\theta_i\kappa(1-p_j)}]^+, \qquad p_j^{BR}(p_i) = [1 - \frac{1}{\theta_j\kappa(1-p_i)}]^+.$$

This, together with Proposition 1, lead to the following result:[10]

[10] The exact values of the investments depend on the cost function adopted, however, the qualitative observations hold for a wide class of such functions.

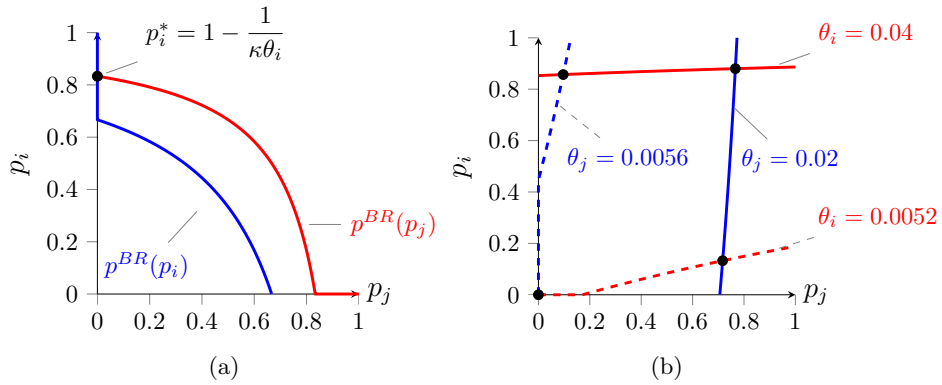

Fig. 2: (a) Example best response curves for the case of $\delta < \tau$, investigated in §3.3. In the figure $\theta_i > \theta_j$. The intersection gives the simultaneous best response pair in the first stage of the game as: $(p_i^*, p_j^*) = ([1 - (\kappa\theta_i)^{-1}]^+, 0)$. The parameters used are: $\lambda = 100$, $\tau = 0.5$, $l = 1$, $\theta_i = 0.04$, $\theta_j = 0.02$. (b) Example best response curves for the case of $\delta < \tau$ and different θ_is and θ_js.

Proposition 2. *If $\delta < \tau$ and $\theta_i > \theta_j$, the Perfect Bayesian Equilibrium (PBE) of the two-stage game is $((p_i^*, s_i^*(p_j, n_i)), (p_j^*, s_j^*(p_i, n_j))) = (([1 - \frac{1}{\kappa\theta_i}]^+, n_i), (0, n_j))$ for all $n_i, n_j \in \mathbb{N}^+$ and all $p_i, p_j \in [0, 1)$, where $\kappa := \lambda(\tau + l)$. That is, only the more efficient firm invests in discovery of the bugs – to achieve discovery probability of $[1 - (\kappa\theta_i)^{-1}]^+$ – and all the findings are then shared.*[11]

Discussion The less efficient firm free-rides on the bug discovery investment of the more efficient company, knowing that all the findings will be shared. This might leap the reader to the conclusion that the PBE outcome is socially inefficient simply because of the existence of "free-riding". However, a social planner may also prefer that the investment is done by the more efficient firm as opposed to distributing the investment among both, hence garnering a higher social return on the aggregate investments. In what follows, we will evaluate the social utility and the socially efficient outcome and compare the two.

Investigating social welfare: Let W represent the expected (utilitarian) *social utility*, defined simply as the sum of the expected utilities of the two firms, i.e., $W := u_i + u_j$.[12] First off, it is straightforward to argue that in the socially optimal outcome, all the findings are shared (the social utility can only be improved by sharing the findings, as the investment decisions are now disentangled from the

[11] When $\theta_i = \theta_j = \theta$, i.e., the two firms are homogeneous in terms of their efficiencies of bug discovery investments, the equilibrium point is not unique and becomes the set: $\{(p_i^*, p_j^*) \in [0, 1]^2, p_i^* = [1 - (\theta\kappa(1 - p_j))^{-1}]^+\}$.

[12] Other notions of social welfare exist, e.g., the egalitarian objective $W := \min(u_i, u_j)$.

sharing decisions). The expected social utility is hence as follows:

$$\mathbb{E}W = -c_i - c_j - 2(\tau + l)\mathbb{E}B_{\neg i,\neg j} = -c_i(p_i) - c_j(p_j) - 2\kappa(1-p_i)(1-p_j) \quad (7)$$

For our example cost function, maximizing $\mathbb{E}W$ hence yields: $(\hat{p}_i, \hat{p}_j) = ([1 - (2\kappa\theta_i)^{-1}]^+, 0)$. Comparing the socially optimal solution with the PBE outcome, we have $\hat{p}_j = p_j^* = 0$, and when $2\kappa\theta_i > 1$, we have: $\hat{p}_i > p_i^*$. That is, to maximize the social utility (sum of the expected utilities of the two firms), the less efficient firm, as in the PBE outcome, makes no investment free-rides on the investment of the more efficient firm. However, compared to the PBE outcome, the more efficient firm invests more. This makes intuitive sense: the less efficient firm offers a lower return on investment (offers less "return" in turning investment into probability of bug discovery) and hence should not invest at all. Instead, the investments must be made by the more efficient firm and all the findings be shared. Moreover, the more efficient firm must consider the aggregate losses and invest more carrying the burden of the two, compared to the PBE, where it only considers the effect of its investment on its own losses. Note that even when the players are homogeneous in terms of their efficiencies, i.e., when $\theta_i = \theta_j$, the socially optimal investment turns out to choose only one of the firms to invest. This is because it will prevent from discovery of the same bugs by both players. The value of the optimum social welfare is:

$$W(\hat{p}_i, \hat{p}_j) = -\ln(2\kappa\theta_i)/\theta_i - 1/\theta_i \text{ for } \kappa\theta_i > 1/2, \text{ and: } -2\kappa \text{ for } \kappa\theta_i \leq 1/2. \quad (8)$$

The social welfare that is achieved at the equilibrium outcome of the game is:

$$W(p_i^*, p_j^*) := -\ln(\kappa\theta_i)/\theta_i - 2/\theta_i \text{ for } \kappa\theta_i > 1, \text{ and: } -2\kappa \text{ for } \kappa\theta_i \leq 1. \quad (9)$$

An example comparison between the two is depicted in Fig. 3a.

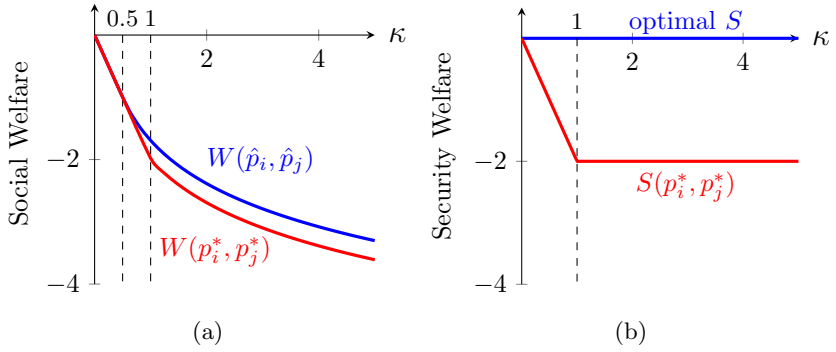

Fig. 3: (3a): Example depiction of the optimal and achieved social welfare (3a) and security utility (3b) for the case of $\delta < \tau$ as functions of $\kappa = \lambda(\tau + l)$.

Here, we define another metric of social welfare in the context of economics of network security. Let the *security utility* u^S of a player be the negative of the

costs of security attacks. Security utility, such defined, is related to the utility of a player as $u^S = u + c$: it includes all the secueity damages but excludes the investment cost. Now, let the *security welfare* S, as a metric of the aggregate security of the two firms, be the sum of their security utilities: $S := u_i^S + u_j^S$. The security utility is related to the utilitarian social welfare in the following way: $S = W + c_i(p_i) + c_j(p_j)$. The optimal S is achieved by picking $p_i = 1$ and sharing all the findings, which yields $S = 0$. Fig. 3b illustrates a comparison between the achieved security utility at the equilibrium and the optimal S.

Comparative statics [13] Recall from Proposition (2), that for $\delta < \tau$, in part we have: $(p_i^*, p_j^*) = ([1 - 1/(\kappa\theta_i)]^+, 0)$. Hence, as long as $\delta < \tau$, $\theta_i > \theta_j$ and $p_i^* > 0$ (i.e., for $1 < \kappa\theta_i$), we have the following straightforward observations:

$$\frac{\partial p_i^*}{\partial \tau}, \frac{\partial p_i^*}{\partial l}, \frac{\partial p_i^*}{\partial \lambda}, \frac{\partial p_i^*}{\partial \theta_i} > 0, \qquad \frac{\partial p_j^*}{\partial \tau}, \frac{\partial p_j^*}{\partial l}, \frac{\partial p_j^*}{\partial \lambda}, \frac{\partial p_j^*}{\partial \theta_j} = 0.$$

We also have $\partial p_i^*/\partial \theta_j = 0$, and perhaps most interesting of all $\partial p_i^*/\partial \delta = 0$; intuitively, player i shares all of its findings and thus removes any dependence of its utility (and hence its best strategy) on δ. Also, note that even though $\partial p_i^*/\partial \theta_i > 0$, i.e., more efficiency in investment means higher choice of probability of discovery, this does not necessarily translate to higher choice of investment. In fact, we have: $\partial c_i(p_i^*)/\partial \theta_i < 0$ for $1 < \kappa\theta_i < e$, and $\partial c_i(p_i^*)/\partial \theta_i > 0$ for $\kappa\theta_i > e$. Moreover, from (9), for $p_i^* > 0$ we have: $W^* := W(p_i^*, p_j^*) = -\ln(\kappa\theta_i)/\theta_i - 2/\theta_i$ and $S^* := S(p_i^*, p_j^*) = -2/\theta_i$. Hence, when $\delta < \tau$, $\theta_i > \theta_j$ and $1 < \kappa\theta_i$, we have:

$$\frac{\partial W^*}{\partial \tau}, \frac{\partial W^*}{\partial l}, \frac{\partial W^*}{\partial \lambda} < 0, \ \frac{\partial W^*}{\partial \theta_i} > 0, \qquad \frac{\partial S^*}{\partial \tau}, \frac{\partial S^*}{\partial l}, \frac{\partial S^*}{\partial \lambda} = 0, \ \frac{\partial S^*}{\partial \theta_i} > 0.$$

3.4 The case of $\delta > \tau$

Following Proposition 1, the dominant strategy of the players in the second stage is to share *none* of their findings, i.e., $s_i(p_j, n_i) = 0$ and $s_j(p_i, n_j) = 0$ for all $n_i, n_i \in \mathbb{N}^+$ and all $p_i, p_j \in [0, 1)$. Then from (3), we obtain:

$$\begin{aligned}\mathbb{E}[u_i(\omega, (p_i, s_i^*), (p_j, s_j^*))] = &-c_i(p_i) + (\delta - \tau)\lambda p_i(1 - p_j) \\ &+ (-\delta - \tau - l)p_j\lambda(1 - p_i) + (-\tau - l)(1 - p_j)\lambda(1 - p_i) \quad (10)\end{aligned}$$

The best response relation for player i is therefore:

$$p_i^{BR}(p_j) = [c_i'^{-1}(\lambda(\delta + l + p_j\tau))]^+.$$

A point to observe is that for $p_i^{BR} > 0$, we have: $\partial p_i^{BR}/\partial p_j = \lambda\tau/c_i''(p_i^{BR}) > 0$, i.e., more investment by the other player leads to *more* investment by a player. This is in sharp contrast to the the previous case of $\delta < \tau$. Similarly: $p_j^{BR}(p_i) =$

[13] In economics, *comparative statics* is the study of the change in the "equilibrium" outcome when a change in a parameter is/would be introduced.

$[c_j'^{-1}(\lambda(\delta + l + p_i\tau))]^+$. For our example cost function, the simultaneous best response is therefore the solution the following system (Fig. 2b):

$$p_i^{BR}(p_j) = [1 - \frac{1}{\theta_i\lambda(\delta + l + p_j\tau)}]^+, \qquad p_j^{BR}(p_i) = [1 - \frac{1}{\theta_j\lambda(\delta + l + p_i\tau)}]^+. \tag{11}$$

Straightforward algebraic investigation reveals that the solution is unique and given as follows:

$$\text{If } \Delta \geq 0\text{: } \begin{cases} p_i^* = \dfrac{\left[-\lambda\theta_i\theta_j((\delta + l)^2 - \tau^2) + \tau(\theta_i - \theta_j) + \sqrt{\Delta}\right]^+}{2\tau\theta_i\theta_j(\delta + l + \tau)} \\ p_j^* = \dfrac{\left[-\lambda\theta_i\theta_j((\delta + l)^2 - \tau^2) - \tau(\theta_i - \theta_j) + \sqrt{\Delta}\right]^+}{2\tau\theta_i\theta_j(\delta + l + \tau)} \end{cases}, \tag{12}$$

and if $\Delta < 0$: $(p_i^*, p_j^*) = (0, 0)$, where $\Delta := \left(\tau(\theta_i + \theta_j) - \lambda\theta_i\theta_j(\delta + l + \tau)^2\right)^2 - 4\tau^2\theta_i\theta_j$. This, along with Proposition 1, fully determines the PBE:

Proposition 3. *When $\delta > \tau$, the Perfect Bayesian Equilibria (PBE) of the security information sharing game is unique, in which (p_i^*, p_j^*) are provided in* (12)*, and $(s_i^*(p_j, n_i), s_j^*(p_i, n_j)) = (0, 0)$ for all $n_i, n_j \in \mathbb{N}^+$ and all $p_i, p_j \in [0, 1)$. That is, both of the firms may invest – to achieve discovery probabilities as given in* (12) *– and none of the consequent findings are shared.*

Discussion When $\delta > \tau$, the competitive gain outweighs the market shrinkage of not sharing the found bugs. Knowing that the found bugs will not be shared, both players, notably even the less efficient player, invest in discovery of the bugs on their own. This is because of two facts: 1- Since the findings are not shared, the firm would be exposed in its bugs if it does not discover and rectify them if it does not invest. 2- Since the other firm invests and expectedly discovers some bugs, the firm will further suffer through the competitive effect of being the sole victim of such bugs if it does not invest.

Comparison to socially optimal outcome: The social optimal outcome certainly shares the found bugs. Compared to the case of $\delta < \tau$, both players invest strictly more in discovery of the bugs. The social inefficiency of the outcome for the case of $\delta < \tau$ was due to underinvestment. Here, it is primarily due to lack of sharing of the found bugs: if a player would receive information of a bug that has not discovered itself, the social utility would have improved by preventing the potential direct losses in that player as well as the market shrinkage losses in both players. Another source of social inefficiency is the fact that "both" players make discovery investment: there is a positive probability that the same bug can be discovered independently by both firms. The investment could have been more efficient by preventing such cases of "duplicate effort", if directed to only one player and the subsequent findings are shared. Anther source of social inefficiency, which is again rooted in lack of information sharing of the players, is the possibility of "over-investment" in bug discovery. The optimal expected social

utility is the same as was computed in (8). Note in particular that it does not depend on the value of δ. Sharing the information in the social optimal removes the competitive effect of δ. However, in the case of $\delta > \tau$, the investment value of both players increases with δ. This means that the threat of competitive losses due to being the sole victim of a security attack can drive both firms to invest inefficiently large values in bug discovery, when they know the discoveries, as competitive advantages, will not be shared. A combination of all of these three effects is responsible for a high social inefficiency in this case.

Comparative Statics Given $\delta > \tau$ and our example cost functions, we note that players' best response functions as in (11) are increasing and concave. Investigating the best-response expressions in (11) further reveals:

$$\frac{\partial p_i^{BR}}{\partial \tau}, \frac{\partial p_i^{BR}}{\partial l}, \frac{\partial p_i^{BR}}{\partial \lambda}, \frac{\partial p_i^{BR}}{\partial \theta_i}, \frac{\partial p_i^{BR}}{\partial \delta} > 0, \quad \frac{\partial p_j^{BR}}{\partial \tau}, \frac{\partial p_j^{BR}}{\partial l}, \frac{\partial p_j^{BR}}{\partial \lambda}, \frac{\partial p_j^{BR}}{\partial \theta_j}, \frac{\partial p_j^{BR}}{\partial \delta} > 0.$$

This means that player i is willing to invest more as any of the following parameters increases: τ, l, λ, θ_i, and similarly for player j (with θ_i replaced by θ_j). Investigating the effect on the equilibrium point is a bit trickier. For simplicity of exposition, we illustrate the "shift" in the equilibrium pair pictorially. In Fig. 4, the effect of increasing δ is depicted. Note that, on the "p_i–p_j" plane, $p_i^{BR}(p_j)$ shifts "up" and $p_j^{BR}(p_i)$ shifts "right" as the value of δ increases. Hence, the intersection, which indicates the equilibrium, moves towards up and right. The algebraic details of the analysis is removed for brevity. Analysing the effect of each parameter in turn reveals:

$$\frac{\partial p_i^*}{\partial \tau}, \frac{\partial p_i^*}{\partial l}, \frac{\partial p_i^*}{\partial \lambda}, \frac{\partial p_i^*}{\partial \delta}, \frac{\partial p_i^*}{\partial \theta_i}, \frac{\partial p_i^*}{\partial \theta_j} \geq 0, \quad \frac{\partial p_j^*}{\partial \tau}, \frac{\partial p_j^*}{\partial l}, \frac{\partial p_j^*}{\partial \lambda}, \frac{\partial p_j^*}{\partial \delta}, \frac{\partial p_j^*}{\partial \theta_j}, \frac{\partial p_j^*}{\partial \theta_i} \geq 0.$$

In words, the above inequalities indicate that if any of the following parameters increases, then firms would invest more: τ, l, λ, and δ. Indeed, the higher these parameters, the more severe impacts of security incidents would be, and thus both firms have to secure themselves, especially when they receive no aid from the other. An interesting result is the effect of improvement in the investment efficiency of the competitor: If θ_j is improved, then firm i invests more in vulnerability research. Intuitively, this is due to the fact that an improvement in the discovery probability of the competitor firm j means more competitive pressure on firm i. This is because each bug that is discovered exclusively by firm j brings it a net advantage of $\delta - \tau$ at the cost of firm i. Thus the increase in efficiency of firm j forces firm i to also improve its probability of discovery, which happens by increasing its investment. This means that the utility of player i decreases as the result of an improvement in player j's efficiency. Specifically, $\partial u_i(p_i^*, p_j^*)/\partial \theta_j < 0$. This is while, $\partial u_j(p_i^*, p_j^*)/\partial \theta_j > 0$. Due to these opposing effects of efficiencies on individual utilities, in general, the equilibrium social welfare, $W(p_i^*, p_j^*)$, which is the sum of the two utilities at the equilibrium, may increase or decrease as θ_i or θ_j is improved. Note, however, that the equilibrium security welfare, $S(p_i^*, p_j^*)$, always improves when θ_i or θ_j increases.

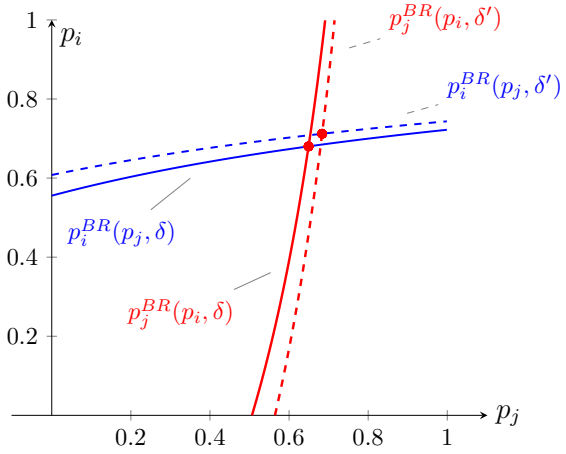

Fig. 4: Example illustration of the comparative statics for the case of $\delta > \tau$. The parameters used are $\lambda = 1.5$, $l = 0.5$, $\theta_i = 1$, $\theta_j = 0.9, \tau = 0.9$, and the value of δ is increased from $\delta = 1$ to $\delta' = 1.2$. Notice the shift in the equilibrium value towards "up" and "right" as a result.

4 Mediation: Encouraging Information Sharing

Our analysis in the previous section characterized the players' behaviour in equilibria. For the case of $\delta < \tau$, which pertain to a the case where security acts effectively as a "common good", sharing of security findings becomes inevitable, and exactly because of that, free-riding emerges, which in turn leads to underinvestment. In contrast, when $\delta > \tau$, which represents cases where security effectively becomes a "competitive advantage", firms would individually strive for their security and refrain from sharing their findings. We observed that none of these outcomes are in line with desirable social planning.

In this section, we make a preliminary attempt to remedy one of the sources of social inefficiency, specifically, failure in information sharing in the "competitive advantage" case. We develop a mediation mechanism that partially removes the negative incentives of sharing the information while allowing the players to gain from its positive effects. Informally put, our mediation plan states that if a firm wants to be informed about n bugs that it failed but the other firm succeeded to discover, it must reveal in exchange n bugs that the other firm is not aware of. Note that this was not possible in the previous sections, as there was no means of making the sharing actions of a firm "contingent" on the action of the other. The mediator effectively ensures that no net "competitive advantage" is lost by sharing the vulnerability findings, as any leakage of an "exclusive" discovery is *matched* by an "exclusive" discovery of the competitor. We will hence refer to our mediation plan as "matched sharing".

Matched sharing operates in two steps: (i) each player/firm submit its set of found bugs to the mediator, along with a specification of a "threshold" as

the maximum number of bugs it is willing to exchange with the other firm. (ii) Subsequently, based on the reported sets and the players' thresholds, the mediator moderates the exchange of as many bugs as possible in the following manner: the mediator marks the bugs that are exclusive to each player, i.e., that the other player has not discovered them. Then the information of a bug is transferred from player i to player j iff a) there is an exclusive bug to *match*, i.e., to transfer from player j to i, and b) if the total number of bugs transferred so far does not exceed either one of the players' requested maximum threshold. Note that the mediator is not a strategic player, and its behaviour is known to and trusted by both players.

From the above description, a sharing action of a player entails the selection of the threshold on exchange number. Note specifically, that we can without loss of generality assume that both players submit all of their findings to the mediator.[14] This is because the players can restrict the sharing of their findings by specifying the threshold. For instance, no sharing corresponds to requesting a threshold of "zero". Note that due to the nature of the Bayesian game, each player must pick this bound for every realisation of bugs it discovers (given the investment decisions). Formally, we can reuse the notations $s_i(p_j, n_i)$ and $s_j(p_i, n_j)$ to represent the sharing strategies, with the different interpretation that s_i and s_j denote the threshold, i.e., the maximum number of their bugs to be shared by the mediator to the other player. Hence, the expressions in (2) in the presence of the mediator and the new interpretation of the strategies become:

$$\mathbb{E}[B_{i,\neg j}|\omega, s_i, s_j] = n_i - n_{ij} - \min\{s_i(p_j, n_i), s_j(p_i, n_j), n_i - n_{ij}, n_j - n_{ij}\}$$
$$\mathbb{E}[B_{\neg i,j}|\omega, s_i, s_j] = n_j - n_{ij} - \min\{s_i(p_j, n_i), s_j(p_i, n_j), n_i - n_{ij}, n_j - n_{ij}\}$$

and, as before, $\mathbb{E}[B_{\neg i,\neg j}|\omega, s_i, s_j] = b - n_i - n_j + n_{ij}$. In words, the term represented by the min function determines the number of bugs that are exchanged between the players, which should be no more than the bounds set by both firms, as well as what each firm individually has to offer. This in turn gives:

$$\begin{aligned} u_i(\omega, \sigma_i, \sigma_j) = &- c_i(p_i) + \delta(n_i - n_j) - \tau(b - n_{ij}) - l(b - n_i) \\ &+ (2\tau + l)\min\{s_i(p_j, n_i), s_j(p_i, n_j), n_i - n_{ij}, n_j - n_{ij}\} \end{aligned} \tag{13}$$

As we can see, the only term that involves $s_i(p_j, n_i)$ is the last term. Maximization of the expected utility of player i given the strategy of player j therefore translates to maximizing $\min\{s_i(p_j, n_i), s_j(p_i, n_j), n_i - n_{ij}, n_j - n_{ij}\}$. Hence, we have the following result:

Proposition 4. *Suppose $p_i, p_j < 1$. The weakly dominant pure Bayesian Nash Equilibrium of the second stage of the game is $(s_i^*(p_j, n_i), s_j^*(p_i, n_j)) = (n_i, n_j)$ for all $n_i, n_j \in \mathbb{N}^+$ and $p_i, p_j \in [0, 1)$, i.e., asking the mediator to share the maximum number of exclusive bugs. This proposition holds irrespective of the distribution of the total number of bugs, or correlation in the discovery of bugs.*

Proof. First, note that irrespective of the choice of s_j, $s_i(p_i, n_i) = n_i$ maximizes the expression $\min\{s_i(p_j, n_i), s_j(p_i, n_i), n_i - n_{ij}, n_j - n_{ij}\}$, and likewise

[14] Assuming that both parties have established trust with the mediator.

for $s_j(p_i, n_j) = n_j$. Hence $(s_i(p_j, n_i), s_j(p_i, n_j)) = (n_i, n_j)$ for all $n_i, n_j \in \mathbb{N}^+$ and $p_i, p_j \in [0, 1)$ belongs to the set of pure Bayesian Nash equilibria of the second stage of the game. To see the weak dominance, consider the cases where $n_j > n_i > 0$ and $n_{ij} = 0$. Note that $\Pr[N_j > n_i \ \wedge \ N_{ij} = 0 \mid N_i = n_i] > 0$. Consider the strategy of player j as $s_j(p_i, n_j) = n_j$ for all $n_j \in \mathbb{N}^+$. Then $u_i(\omega, (p_i, n_i), (p_j, s_j)) > u_i(\omega, (p_i, s'_i), (p_j, s_j))$ for any $s'_i(p_j, n_i) < n_i$, because: $\min\{n_i, s_j(p_i, n_j), n_i - n_{ij}, n_j - n_{ij}\} > \min\{s'_i(p_j, n_i), s_j(p_i, n_j), n_i - n_{ij}, n_j - n_{ij}\}$ for any $s'_i(p_j, n_i) < n_i$ when $n_j > n_i$, $n_{ij} = 0$ and $s_j(p_i, n_j) = n_j$.

4.1 Game's first stage: Investment in the presence of the Mediator

Given the weakly dominant equilibrium in Proposition 4, $\min\{s_i^*(p_j, N_i), s_j^*(p_i, N_j), N_i - N_{ij}, N_j - N_{ij}\} = \min\{N_i, N_j\} - N_{ij}$. Hence, utility of player i in (13) becomes:

$$\begin{aligned}\mathbb{E}u_i(\omega, p_i, p_j, s_i^*, s_j^*) &= -c_i(p_i) + \delta\mathbb{E}[N_i - N_j] - \tau\mathbb{E}[B - N_{ij}] - ql\mathbb{E}[B - N_i] \\ &\quad + (2\tau + l)(\mathbb{E}[\min\{N_i, N_j\}] - \mathbb{E}[N_{ij}]) \\ &= -c_i(p_i) + \lambda\delta(p_i - p_j) - \lambda\tau(1 - p_i p_j) - \lambda l(1 - p_i) + (2\tau + l)(\mathbb{E}[\min\{N_i, N_j\}] - \lambda p_i p_j)\end{aligned}$$

The term $\mathbb{E}[\min\{N_i, N_j\}]$ depends on the specific distribution of the total number of bugs. A good candidate is the Poisson distribution. The presence of this term in the utility function prevents a closed-form solutions for the best responses and the equilibrium points. Instead, we pictorially illustrate in Fig. 5 the potential usefulness of the mediator when $\delta > \tau$, i.e., when players are motivated more by competition than aggregate security. Fig. 5a depicts the equilibrium points of players' investments in two cases: sharing in the absence of the mediator (which leads to no sharing) and our "matched sharing". These are set in the context of low security damage (l) compared to competitive advantage (δ) and inefficient investment ($\theta_i = \theta_j = 0.1$). The end result is that with matched sharing, both players invest more in finding vulnerabilities, which guarantee a a better security for both. However, the social welfare, as well as the individual utilities of both players, worsens with the introduction of the matched sharing, as it exacerbates the already inefficiently high investments of the players in this example.

In contrast, Fig. 5b shows the effect of our mediator plan in situations with either a significant security damage value (large l) or efficient investments (high θ_i, θ_j), or both. In such scenarios, equilibrium points of the two cases are relatively close to each other, i.e., they make similar levels of investments. With the help of the mediator, players would share their intelligences and thus gain extra value in security, making mediation a superior solution to opportunistic sharing. This suggests the potential of our matched sharing mediation scheme, and that it should be in the interest of the social planner to monitor environment parameters and establish trusted mediation among firms whenever appropriate for players/societal benefits.

5 Conclusion

In this work, we focused on the problem of sharing cybersecurity information, as an envisioned pillar of cybersecurity planning for a more secure infrastruc-

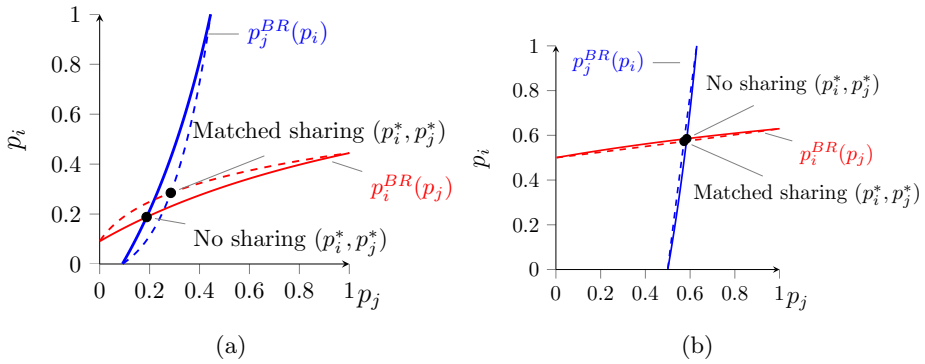

Fig. 5: Illustration of opportunistic sharing vs. matched sharing when $\delta > \tau$, with $\delta = 10$, $\tau = 1$, $\theta_i = \theta_j = 0.1$, with (a) $l = 1$ and (b) $l = 10$.

ture. We analysed the strategic decisions of two competing firms with regards to investment for discovery of security vulnerabilities (generating valuable cyber-intelligence) and subsequently, to share their findings. We showed that sharing becomes a dominant strategy when security tends to behaves as a common good, i.e., when the common losses as a result of security attacks outweigh the competitive gains of being protected. We analysed how in turn this leads to free-riding of less efficient firm and the under-investment of the more efficient firm. We also established that when security effectively becomes a competitive advantage, i.e., when there is a net positive gain when a competitor is a sole victim of an attack, then sharing no information becomes the dominant strategy, with negative implication on the social efficiency. Finally, we provided a monetary-free light-weight mediation mechanism that (partially) enables sharing of the found vulnerabilities in cases where they fail to achieve any sharing on their own.

Future Research This work has the potential to be extended in many directions. We have already made some grounds in extending our results to the multi-player situation. An interesting addition is considering "features" for the found bugs, such as severity (seriousness of the potential damage), sophistication (exploitability), etc., and hence letting the sharing strategies depend on the type of the found bug as well. Investigating the behaviour of risk-averse players – as opposed to risk-neutral in this work – is another problem. Identifying other types of "security information" to share is another interesting direction, for instance, revealing past incidents of successful attacks and resultant losses carries some market implications that sharing merely discovered security vulnerabilities does not. Also, we assumed that both firms use a common implementation (the "platform"). If instead, for instance, the firms are using a common protocol but with their private implementations of it, then "some" of the discovered bugs may be just exclusive to that party's implementation. Sharing found bugs now requires a modified analysis. Investigating other means of encouraging sharing is another important direction. An example is "bargaining": A player starts by

sharing one bug, then the other player matches with a bug of its own findings, and so on, until one stops. Another example is a generalisation of the "matched sharing" mechanism in this work by allowing unequal number of matching that may involve some randomisation as well. An exchange market of vulnerabilities is another idea, although it may suffer from adverse selection and moral hazard.

References

1. Press-Release: Government launches information sharing partnership on cyber security (March 2013) www.gov.uk[Online; 27-March-2013].
2. of Homeland Security, D.: National cybersecurity and communications integration center www.us-cert.gov/nccic[Online; Accessed June-2014].
3. Cavusoglu, H., Mishra, B., Raghunathan, S.: The effect of internet security breach announcements on market value: Capital market reactions for breached firms and internet security developers. Intl. J. of Electronic Commerce **9**(1) (2004) 70–104
4. Campbell, K., Gordon, L.A., Loeb, M.P., Zhou, L.: The economic cost of publicly announced information security breaches: empirical evidence from the stock market. Journal of Computer Security **11**(3) (2003) 431–448
5. Goel, S., Shawky, H.A.: Estimating the market impact of security breach announcements on firm values. Information & Management **46**(7) (2009) 404–410
6. Lovells, H.: DOJ and FTC clarify antitrust implications of cybersecurity information sharing (April 2014) http://www.hoganlovells.com/ [Online; 22-April-2014].
7. Netcraft: Half a million widely trusted websites vulnerable to heartbleed bug (April 2014) news.netcraft.com[Online; 08-April-2014].
8. Gordon, L.A., Loeb, M.P., Lucyshyn, W.: Sharing information on computer systems security: An economic analysis. Journal of Accounting and Public Policy **22**(6) (2003) 461–485
9. Gal-Or, E., Ghose, A.: The economic incentives for sharing security information. Information Systems Research **16**(2) (2005) 186–208
10. Hausken, K.: Income, interdependence, and substitution effects affecting incentives for security investment. J. of Accounting and Public Policy **25**(6) (2006) 629–665
11. Hausken, K.: Information sharing among firms and cyber attacks. Journal of Accounting and Public Policy **26**(6) (2007) 639–688
12. Liu, D., Ji, Y., Mookerjee, V.: Knowledge sharing and investment decisions in information security. Decision Support Systems **52**(1) (2011) 95–107
13. Liu, C.Z., Zafar, H., Au, Y.A.: Rethinking fs-isac: An it security information sharing network model for the financial services sector. Communications of the Association for Information Systems **34**(1) (2014) 2
14. Xiong, Q., Chen, X.: Incentive mechanism design based on repeated game theory in security information sharing. In: 2nd International Conference on Science and Social Research (ICSSR 2013), Atlantis Press (2013)
15. Gao, X., Zhong, W., Mei, S.: Security investment and information sharing under an alternative security breach probability function. Inf. Systems Frontiers 1–16
16. Gal-Or, E.: Information sharing in oligopoly. Econometrica: Journal of the Econometric Society (1985) 329–343
17. Shapiro, C.: Exchange of cost information in oligopoly. The review of economic studies **53**(3) (1986) 433–446
18. Vives, X.: Trade association disclosure rules, incentives to share information, and welfare. RAND Journal of Economics **21**(3) (1990) 409–430
19. Katz, M.L., Shapiro, C.: R and d rivalry with licensing or imitation. The American Economic Review (1987) 402–420