Security Audits Revisited

Rainer Böhme

Department of Information Systems, University of Münster, Germany
rainer.boehme@uni-muenster.de

Abstract. Security audits with subsequent certification appear to be the tool of choice to cure failures in providing the right level of security between different interacting parties, e. g., between an outsourcing provider and its clients. Our game-theoretic analysis scrutinizes this view and identifies conditions under which security audits are most effective, and when they are not. We find that basic audits are hardly ever useful, and in general, the thoroughness of security audits needs to be carefully tailored to the situation. Technical, managerial, and policy implications for voluntary, mandatory, unilateral, and bilateral security audits are discussed. The analysis is based on a model of interdependent security which takes as parameters the efficiency of security investment in reducing individual risk, the degree of interdependence as a measure of interconnectedness, and the thoroughness of the security audit.

1 Introduction

Information technology has spurred innovation and productivity gains [8], but the flip side is the emergence of cyber risk. A characteristic feature that distinguishes this "new" kind of risk from traditional risks is the sensitivity to interdependence between decisions of individual actors [16]. Therefore, a profound understanding of the particularities of cyber risk is essential to guide the design of secure systems as well as supporting organizational measures. Security audits belong to the set of organizational tools to manage and regulate risk-taking in the internet society. This paper sets out to rigorously analyze why and under which conditions security audits can be most effective.

1.1 Interdependent Cyber Risk

In the context of cyber risk, *interdependence* means that the success of risk mitigation does not only depend on the actions of the potentially affected party, but also on actions of others. In economics jargon, interdependence can be described as an instance where security investment generates *externalities*.

Examples for interdependent security risks exist on various levels of abstraction. For example, modern software engineering relies on the composition of reusable components. Since the security of a system can be compromised by a single vulnerability, the overall security of a system does not only depend on the effort of the first-tier developer, but also on the effort of the developers of

components, libraries, development tools, and the transitive closure thereof (i. e., libraries used by components, development tools used to build libraries, etc.). Also service-oriented architectures provide a wide range of examples for interdependence. In supply chains and other kinds of outsourcing relations, including all types of cloud computing, the confidentiality, availability, and oftentimes also the integrity of business-relevant data depends on the security level of all involved parties. As a final example, take the internet as a whole. Botnets are the backbone for a range of threat scenarios. Their existence and growth is only possible because not all nodes connected to the network make sufficient efforts to secure their systems. So the insecurity of a victim partly depends on the inaptitude of others—that is a clear case of interdependence.

The issues arising from interdependent security, notably free-riding and lack of incentives to invest in security, are not always reflected in the literature discussing the potentials of interconnection for the sake of sharing information, services, and other resources. If not fully ignored, these issues are often described as open yet solvable (e. g., [3]). Or the problem is deferred informally to security audits and certification (e. g., [23, 18]). These organizational measures would ensure high enough security standards. Such claims motivate us to take a closer look at security audits and interdependence to see if the hopes are realistic.

1.2 Security Audits

Generally, it is hard to directly measure the security level of products, systems, services, or organizations [14]. This has mainly two reasons: first, the difficulty of specifying all security requirements—the bug versus feature problem. And secondly, threats neither occur deterministically nor is their occurrence observable in realtime. Hence the conclusion a system be secure because no attacks were observed in the past is obviously invalid. One might just have been lucky that no attacks occurred, or the consequences of successful attacks—for instance loss of confidentiality—will only be observable at a later point in time. These difficulties impede measuring the security level of one's *own* systems. It is easy to see that the problems aggravate for systems owned by *others*, as it is the case in the context of interdependence. Therefore, security almost always has the properties of a credence good [2].

As direct measurement is hard, one can resort to examine all security-relevant attributes of an object to *estimate* its latent security level. This involves considerable effort, because these examinations are not fully automatizable and they require special knowledge and experience of the examiner. Moreover, the effort will often grow disproportionately to the complexity of the object under investigation because more and more dependencies need to be checked. We refer the reader to the literature (e. g., [28]) for an overview of different types of security assessments and their process models. According to this literature, semi-standardized examinations can at least help to identify weaknesses against specific known threats, and to fix the weaknesses thereafter.

Our notion of *security audit* in this paper goes beyond a mere examination. It also includes certification by the examiner who is trusted by third parties. This

way, the result of an examination is verifiable and can serve as a credible signal to other market participants. *Pure security examinations without certification are not subject of this paper* because they cannot contribute to solve the problems arising from interdependent cyber risk.

In practice, security audits with subsequent certification are very common and cause substantial costs to the industry. Examples include the Common Criteria (where audits may last up to one year and cost up to a million dollars) and their predecessors Orange Book (U. S.) and ITSEC (Europe). These standards were designed for public procurement. Other security audits are laid down in industry standards, such as PCI DSS for payment systems or ISO 17799, respectively ISO 27001. In addition, there exists a market for a variety of quality seals issued by for-profit and non-profit organizations alike. Examples include VeriSign, TrustE, or the European data protection seal EuroPriSe.

1.3 Economics of Security Audits

Economic theory suggests two channels through which security audits can generate positive utility:

1. **Overcoming information asymmetries.** From the fact that security is a credence good follows a lemon market problem [2]. The demand side lacks information about the quality of goods. In the simplest case, this quality information can be thought of a binary attribute: secure versus insecure. It can be shown that the equilibrium price for goods of unknown quality drops to the price of insecure goods. As a result, no market exists for secure products. Security audits can help to signal quality and fix this market failure.
2. **Solving coordination problems.** If credible signals are available, additional strategies emerge in the game-theoretic models of interdependent security. The players not only decide about their own security investment, but also whether or not to signal information about their own security level. This can generate new welfare-maximizing equilibria or stabilize existing ones. Security audits are the means to generate credible signals in practice.

Understanding both channels is certainly relevant. However, only the second channel is directly linked to interdependent security. Therefore, we concentrate our attention in this paper on the solution of the coordination problem and refer the reader to the relevant literature [1, 21, 2] on the role of audits in fixing information asymmetries (cf. Sect. 4 for comments on that literature).

Note that in practice, security audits are commissioned also—if not primarily—because of legal or contractual obligations. Another reason can be liability dumping: a CIO might find it easier to repudiate responsibility after a successful attack by referring to regular security audits, no matter how sound they actually are. Both motivations can generate individual utility. In the following, we will not directly deal with these motivations. The focus of our analysis is economic in nature, that means, in the long run, uninformative audits will not help the CIO in the above example. With regard to mandatory audits, we start one step ahead.

The very objective of our analysis is to scrutinize the economic justification of existing or future legal and contractual obligations *because of* their potential to prevent market failures.

1.4 Research Question and Relevance

Now we can formulate the research question: *Under which conditions do security audits (defined in Sect. 1.2) generate positive utility by solving the coordination problems (see Sect. 1.3), which would otherwise hinder the reduction of interdependent cyber risks?*

The response to this question is relevant for security managers who decide whether commissioning security audits is profitable[1] given:

a. the *security productivity*, a property of the organization and its business,
b. the *thoroughness* of the security audit, and
c. the *degree of interdependence*, a property of the organization's environment.

Our contribution in this paper is a new analytical model to answer this research question. The model can also be employed for decision support whenever a change in one of the conditions (a–c) is anticipated. The latter mainly concerns decisions to increase interconnectivity, for instance by supporting more interfaces or integrating new services. Each affects the degree of interdependence.

Solving the coordination problem not only increases individual utility, but also leads to improvements in social welfare. Therefore, our model and its analysis is equally relevant for regulators. For example, regulations requesting mandatory audits should be designed such that audits are only required when it is economical. Moreover, the model can help to formulate high-level requirements for security audits such that audits have a welfare-maximizing effect.

Everything that has been said for the regulator can be applied to market situations in which one market participant defines the standards for an industrial sector. This can be an industrial organization (such as in the case of PCI), or a blue chip company orchestrating its supply chain.

1.5 Roadmap

The next section presents our model, which is designed parsimoniously without omitting properties necessary for the interpretation. The model is solved and all pure strategy equilibria identified. Section 3 analyzes the equilibria with regard to the utility generated by security audits. We will explain under which condition security audits are helpful, and when they are not needed to solve the coordination problem. Section 4 discusses relations to prior art, both in terms of the subject area and the analytical methodology. A critical discussion and our outlook precede the final conclusion (Sect. 5).

[1] We are agnostic about defining a price for the security audit. Hence "profitable" should be read in the sense of strictly positive utility.

2 Model

The analytical model consists of three components: a formalization of the security audit process, a model of security investment, and a model of interdependent security. Each component includes exactly one free parameter, that is one for each of the three properties (a–c) described informally in Section 1.4. To the extent possible, we combine established modeling conventions. However, the resulting model as a whole is novel and specific to the analysis of security audits.

2.1 Stylized Audit Process

To capture security audits in an economic model, it is essential to reduce them to their most relevant features. In particular, at our level of abstraction, it does not matter *how* a security audit is conducted technically and organizationally. The only relevant outcome is its result.

For this we assume that every examinable object X has a latent—i. e., not directly observable—attribute $s_X \in \mathbb{R}^+$ describing its security level. Objects X of interests can include products, systems, services, or entire organizations. The probability of loss due to security incidents decreases monotonically with increasing security level s_X.

Now we can model a security audit as function which takes object X as input, compares its security level s_X to an internal threshold t, and returns one bit,

$$\mathsf{SecAudit}(X) = \begin{cases} 1 \text{ if } s_X \geq t \\ 0 \text{ otherwise.} \end{cases} \tag{1}$$

The result of the audit shall be verifiable by third parties. In practice this can be ensured by issuing a (paper) certificate or by having the auditor sign the result cryptographically. In any case, the result is just a snapshot in time and has to be annotated with a time stamp if state changes of X are of interest. Our analysis in this paper is limited to one-shot games with fixed states.

The assumption of a threshold t can be justified with the common practice to conduct security audits along semi-standardized checklists where the thoroughness of the audit has to be defined beforehand. It is certainly conceivable to consider a family of functions $\mathsf{SecAudit}_t$ from which the appropriate function is selected depending on the situation. A real-world example for this are the seven Evaluation Assurance Levels (EAL) specified in the Common Criteria. However, it most cases the number of different thresholds will be small and countable. So we cannot assume that t can be chosen from a continuous interval.

Note that we simplify the audit problem to a single summative measure of security level. In practice, different aspects (e. g., protection goals, security targets in the Common Criteria terminology, etc.) or components of a system can have different levels of security. This view is compatible with our approach if one considers each system as a bundle of objects X and a given security audit as a collection of functions, one for each property of the bundle.

Our abstraction ignores that practical audits may cause side effects. Audits impose costs, which typically depend on X and t. There is also a risk of hidden information leakage as the auditor and its staff may get to know sensitive information about X. In a dynamic setting, there might be a non-negligible lag between the time when the audit decision is taken and the time when the output is available. All these side effects are not considered in this paper. Therefore, our simplifications may let security audits appear more useful in our analysis than they actually are in practice. Conversely, we err on the side of caution in cases where security audits turn out useless in our analysis. The reader is advised to keep this bias in mind when interpreting our results.

2.2 Security Investment

Consider for now a single firm[2] making security investments to reduce the probability of incurring a loss of unit size $l = 1$ due to security incidents. We adopt the functional relationship between security investment s and the probability of loss $p(s)$ from the well-known Gordon–Loeb model of security investment [11],

$$p(s) = \beta^{-s}. \tag{2}$$

This function reflects a decreasing marginal utility of security investment, a property that has been confirmed empirically [17], by practitioners [11], and can be justified theoretically [7]. Parameter $\beta \geq e^2$ represents the firm-specific *security productivity*. The range of s is limited to the interval $[0, 1]$. This is so because risk-neutral firms prefer $s = 0$ over all alternatives $s > l = 1$. To keep the number of parameters manageable, we fix the parameter for vulnerability in [11] at $v = 1$: without security investment, every realized threat causes a loss.

Our model shares another simplification with most analytical models of information security investment. It does not distinguish between security investment and security level. This implies the assumption that all security investment is effective. By contrast, practitioners often observe the situation of security overinvestment (from a cost perspective) still leading to a suboptimal security level [6]. Hence caution is needed when transferring conclusions on security overinvestment or under-investment from analytical models to the real world.

The firm's expected cost can be expressed as sum of the security investment and the expected loss,

$$c(s) = s + p(s) = s + \beta^{-s}. \tag{3}$$

This model is good enough to find optimal levels of security investment for a single firm. However, without interdependence, this is not of interest here.

2.3 Modeling Interdependence

The simplest possible case to model interdependence is to assume two symmetric a priori homogeneous firms who act as players in a game. Security investments s_0

[2] For consistency and didactic reasons, we use the term "firm" to refer to a single rational decision maker. This does not limit the generality of the model. Firm stands as placeholder for any entity conceivable in a given context, e. g., "organization", "defender", "nation state", "player", or "user".

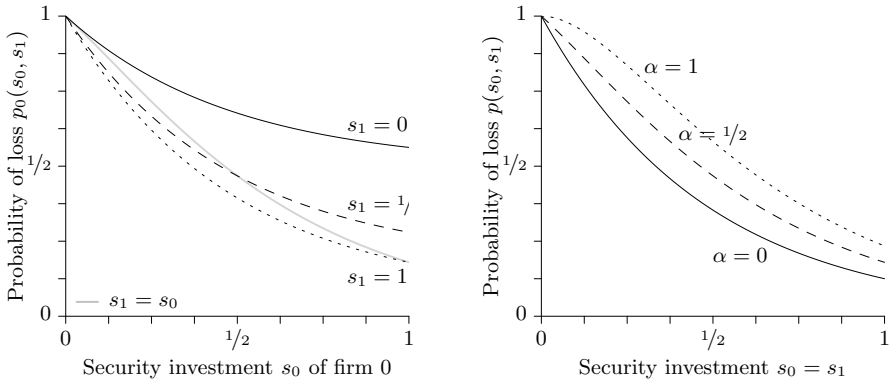

(a) $\beta = 8$, $\alpha = 1/2$ (b) $\beta = 8$

Fig. 1: Interdependent security: firm 0's probability of loss partly depends on firm 1's security investment (left); the probability of loss increases with the degree of interdependence α even if both firms invest equally (right)

and s_1 are the only choice variables. Symmetry implies that both firms share the same security productivity β. This can be justified by generalizing Carr's argument [9] to security technology. Security technology has become a "commodity" which is rarely a factor of strategic differentiation between firms.

Consider the following function for the probability of loss p_i of firm $i \in \{0, 1\}$,

$$p_i(s_i, s_{1-i}) = 1 - (1 - \beta^{-s_i})(1 - \alpha\beta^{-s_{1-i}}). \tag{4}$$

This reflects the intuition that a firm evades a loss only if neither it falls victim to a security breach, nor a breach at an interconnected firm is propagated. Parameter $\alpha \in [0, 1]$ is the *degree of interdependence*, a property of the environment of both firms. For $\alpha = 0$ (no interdependence), Eq. (4) reduces to Eq. (2).

Figure 1a illustrates the effect of interdependence described informally in the introduction. We set $\alpha = 1/2$ for moderate interdependence. The black curves show that the probability of loss of firm 0, for every choice of its own security investment $s_0 > 0$, also depends on the choice of s_1 by firm 1. By contrast, the gray intersecting curve shows the probability of loss if both firms make equal security investments. This setting prevails in Figure 1b. Here we show curves for different settings of the degree of interdependence α. Observe that the probability of loss grows with the degree of interdependence for every fixed security investment $s_0 = s_1 > 0$.

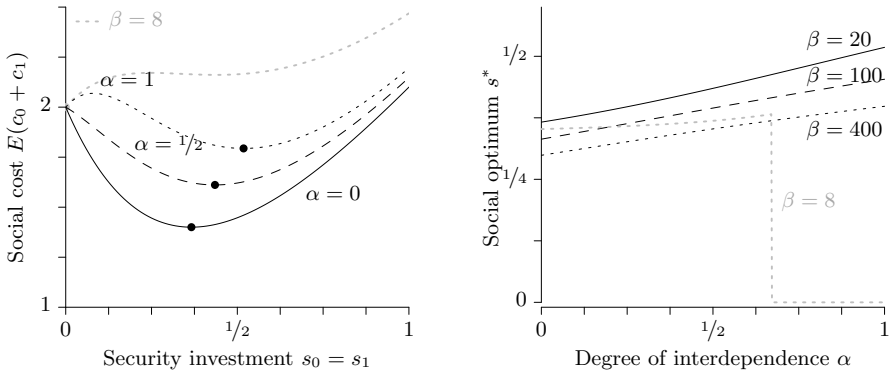

(a) $\beta = 20$ (except topmost curve) (b) Location of social optima

Fig. 2: Security investment s^* minimizes the expected "social" cost of both firms

2.4 Social Optima

A social optimum is reached if the sum of the expected costs of both firms is minimal, thus

$$s^* = \arg\min_s 2 \cdot c(s, s) \tag{5}$$

$$= \arg\min_s s + 1 - (1 - \beta^{-s})(1 - \alpha\beta^{-s}). \tag{6}$$

We may substitute s_i by s due to symmetry. Figure 2a shows the objective function and their minima for selected parameters. Their location can be obtained analytically from the first-order condition of Eq. (6). We get

$$s^* = -\log\left(\frac{(1+\alpha) - \sqrt{(1+\alpha)^2 - 8\alpha\log^{-1}(\beta)}}{4\alpha}\right)\log^{-1}(\beta) \tag{7}$$

for $\alpha > 0$, and

$$s^* = \log(\log(\beta))\log^{-1}(\beta) \quad \text{for the special case } \alpha = 0. \tag{8}$$

For high degrees of interdependence and low security productivity, the social optima reside at the lower end of the value range of s. The gray dotted curves in Figs. 2a and 2b visualize this case (for $\beta = 8$). We will discuss the implications of this special case on security audits below in Sect. 3.4.

Figure 2b shows the location of social optima as a function of α for selected values of β. Observe that the socially optimal security investment does not react

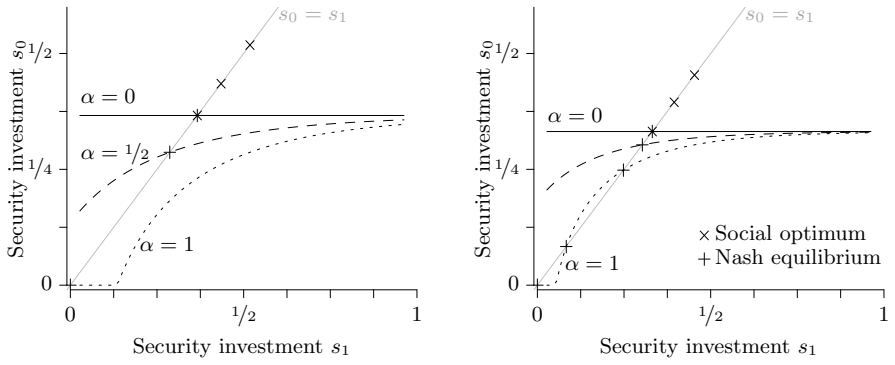

(a) $\beta = 20$ (b) $\beta = 100$

Fig. 3: Best response of firm 0 given security investment of firm 1; all fixed points of this function are pure strategy Nash equilibria

monotonously to changes in the security productivity β. Apart from the above-described discontinuity, increasing degree of interdependence α shifts the social optimum towards higher levels of security investment s. Frankly speaking, this means that an increasingly interconnected society ceteris paribus has to spend more and more on security to maintain a welfare-maximizing[3] level. This follows directly from the relation depicted in Figure 1b.

2.5 Nash Equilibria

Knowing the location of social optima does not imply that they are reached in practice. This will only happen if all players have incentives to raise their security investment to the level of s^*. The analysis of incentives—which obviously depend on the actions of the respective other firm—requires a game-theoretic perspective and the search for Nash equilibria.

Only pure strategies are regarded in this paper. Firm i's best response s^+ given s_{1-i} is the solution to the following optimization problem:

$$s^+(s_{1-i}) = \arg\min_s s + p(s, s_{1-i}) \tag{9}$$

$$= \arg\min_s s + 1 - \left(1 - \beta^{-s}\right)\left(1 - \alpha\beta^{-s_{1-i}}\right) \tag{10}$$

$$\text{s. t.} \quad s \geq 0.$$

After finding roots of the first-order condition and rearranging, we obtain:

$$s^+(s_{1-i}) = \sup\left\{\frac{\log\left(\log(\beta)\right) + \log\left(1 - \alpha\beta^{-s_{1-i}}\right)}{\log(\beta)}, 0\right\}. \tag{11}$$

[3] Welfare is defined as the reciprocal of social cost.

Figure 3 shows the best response as function of s_1 for three different degrees of interdependence ($\alpha \in \{0, 1/2, 1\}$) and two values of security productivity ($\beta \in \{20, 100\}$). Nash equilibria, defined as fixed points of the best response function, are located on the intersections with the main diagonal. For comparison, we also plot the social optima as given by Eq. (7).

Depending on the parameters, there exist up to three Nash equilibria at

$$\tilde{s}_{1,2} = \log\left(\frac{\log(\beta) \pm \sqrt{\log^2(\beta) - 4\alpha\log(\beta)}}{2}\right)\log^{-1}(\beta) \tag{12}$$

(if both expression and discriminant are positive) and

$$\tilde{s}_3 = 0 \quad \text{for} \quad \alpha > 1 - \log^{-1}(\beta). \tag{13}$$

The parameters in Figure 3 are chosen such that every case of interest is represented with at least one curve. We will discuss all cases jointly with the interpretation in Section 3. The formal conditions for the various equilibrium situations are summarized in Appendix A.1.

3 Analysis

For all strictly positive values of $\alpha > 0$, the Nash equilibria are located below the social optimum. This replicates a known result: security as a public good is under-provided in the marketplace [27, 16]. The reasons are lack of incentives, more specifically a coordination problem [24]. If firm i knew for sure that firm $1 - i$ cooperates and invests $s_{1-i} = s^*$, then it would be easier to decide for the socially optimal level of security investment as well. In practice, however, firm i can hardly observe the level of s_{1-i}.

Security audits can fix this problem. They allow a firm to signal the own security level to its peers in a verifiable way. This can convince others of the willingness to cooperate and stimulate further cooperative responses. Now we have to distinguish between the case of coordination between multiple equilibria, and the case of coordination at non-equilibrium points. The former helps to avoid bad outcomes, the latter in needed to actually reach the social optimum.

Coordination Between Multiple Equilibria. If multiple Nash equilibria exist, the initial conditions determine which equilibrium is chosen. So the coordination problem is to nudge the game into the equilibrium with the lowest social cost. To do this, it is sufficient if one firm unilaterally signals a security level in the basin of attraction of the best possible equilibrium. Then the other firms' rational selfish reaction is to choose a security level within that basin and the trajectory of strategic anticipation converges to the desired equilibrium solution. Therefore, in this case, it is sufficient to have *unilateral* security audits which may even be *voluntary* (if the audit costs are not prohibitive).

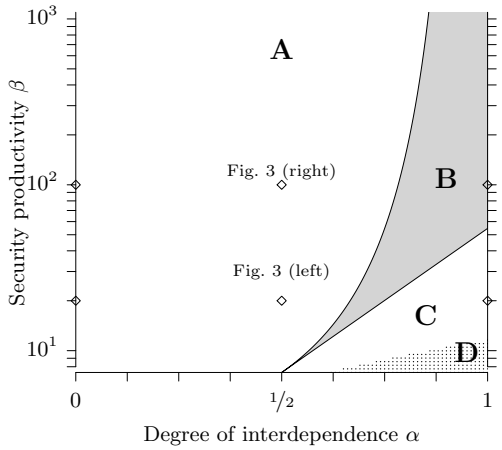

Fig. 4: Case distinction in (α, β)-parameter space

Coordination at Non-Equilibrium Points. The social optimum is not an attractor in general. Therefore, to reach it, *bilateral* security audits and additional incentives are needed. These incentives could come in the form of *mandatory* security audits and sanctions in case a claimed security level is not met. Sanctions can be enforced by regulation or be inherently embedded in the mechanism. For example, a tit-for-tat strategy of a multi-period prisoner's dilemma entails sanctions by other players [4]. A prerequisite for this strategy is the unambiguous observability of security levels in past rounds. Hence security audits are also essential in this setup.

Now we will analyze the equilibrium situations and discuss implications on the usefulness of security audits depending on their thoroughness t. For this it is useful to regard Figure 4, which identifies four equilibrium situations as regions in the (α, β)-parameter space. Six diamond marks indicate the points in parameter space for which curves are displayed in Figure 3.

3.1 Region A: Only Thorough Audits Useful

In region A, there exists exactly one Nash equilibrium (see dashed curves in Fig. 3). The best response ${s^+}_i$ on security investments $s_{1-i} < \tilde{s}_1$ below the Nash equilibrium of Eq. (12) is always larger than s_{1-i}. Therefore, firms always have incentives to invest at least $\tilde{s}_1$. Security audits with thoroughness $t < \tilde{s}_1$ below that level do not improve the situation and hence are ineffective. Thorough audits with $\tilde{s}_1 > t \geq s^*$ can improve the security level and social welfare. Since this involves a coordination at non-equilibrium points, such audits must be bilateral. To specify this, it holds that unilateral audits with thoroughness above

the social optimum for $\alpha = 0$—this is the only point ($*$ in Fig. 3) where the Nash equilibrium and social optimum concur—can never be more effective than unilateral security audits at this level. This is so because this value bounds the best response function from above.

3.2 Region B: Basic Audits Get Leverage

In region B, there exist three Nash equilibria (see dotted curve in Fig. 3b). In one of them, both firms abstain from security investments ($s_0 = s_1 = 0$). In this case, security audits can be maximally effective in solving the coordination problem between the multiple equilibria. To achieve this, the thoroughness t must be just above $\tilde{s}_2$. Then a unilateral audit is enough to move both firms into the best possible equilibrium. More thorough audits in the range $\tilde{s}_2 < t \leq \tilde{s}_1$ do not improve the situation further. In other words, a basic audit just above $\tilde{s}_2$ is leveraged to maximum outcome.

Even the best possible equilibrium is below the social optimum. To approach the optimum further, more thorough audits $t > \tilde{s}_1$ are needed. Everything said for thorough audits above in Sect. 3.1 also applies here. In particular, thorough audits must be bilateral. Superficial audits $t < \tilde{s}_2$ are moderately useful. The situation corresponds to the case discussed in the next section.

3.3 Region C: All Audits Moderately Useful

In region C, there exists exactly one Nash equilibrium in which both firms abstain from security investment (see dotted curve in Fig. 3). The distance between this equilibrium and the social optimum reaches a maximum. This case is not a coordination game in the strict sense [24]. Therefore, the effectiveness of all audits is much more limited than in region B (Sect. 3.2). Even though audits may contribute to higher security levels, more specifically, exactly at the level of the thoroughness t, if both firms perform bilateral audits. Unilateral audits are less effective in general and completely ineffective within the range where the dotted curve in Figure 3 is flat at level 0. Like in regions A and B, unilateral audits above the social optimum for $\alpha = 0$ are strictly dominated by unilateral audits of thoroughness equal to this level.

3.4 Region D: All Audits Useless

In region D, there exists exactly one Nash equilibrium in which both firms abstain from security investment. This concurs with the corner solution of the social optimum (compare the penultimate paragraph of Sect. 2.4 and see the dotted gray curve in Fig. 2a). This means all security investment is prohibitively expensive compared to the protection it promises. In other words, the firm's business is indefensible. Of course, firms would not decide to conduct audits voluntarily (attesting the absence of security investment). Manadatory audits of thoroughness $t > 0$ coupled with sanctions would induce security over-investment and destroy

social welfare. The only resorts are to improve security productivity by technological innovation or to reduce the degree of interdependence. Both measures would move the situation back to region C.

3.5 Left Edge: No Audits

Figures 4 hides the fact that the left edge ($\alpha = 0$) does not belong to region A. This edge rather represents the special case of independent firms who optimize on their own. There exists exactly one Nash equilibrium which concurs with the social optimum (see solid line in Fig. 3). No firm would conduct security audits voluntarily. Mandatory audits (with sanctions in case of failure) do not result in relevant signals if $t \leq s^* = \tilde{s}$. They are even counter-productive if $t > s^* = \tilde{s}$.

In summary, the most salient new result of this analysis is that even in this stylized model, the usefulness of security audits and the required thoroughness highly depends on the situation. We deem this an important insight for the design of audit standards and policies, which in practice are applied in contexts with many more potentially influential factors.

4 Related Work

We are not aware of any prior work addressing this specific or closely related research questions. The same holds for the combination of elements used in our analytical model. Consequently, we structure the discussion of related work broadly into two categories: works which address similar questions, and works that use similar methods for different research questions.

Anderson [2] belongs to the first category. He notes perverse incentives for suppliers of security certifications. This leads vendors who seek certification to shop for the auditor who has the laxest reading of a standard. Baye and Morgan [5] study certificates as an indicator of quality in electronic commerce. They propose an analytical model of strategic price setting in a market where certified and uncertified products compete. They find support for their model using empirical data. In another empirical study, Edelman [10] argues that less trustworthy market participants have more incentives to seek certification (and obtain it). He could show that this adverse selection inverts the intended function of TrustE seals as indicators of quality. The appearance of the seal on a representatively drawn website actually *increases* the posterior probability that the site is shady. This is largely driven by the fact that TrustE certification is voluntary, leading to self-selection. Rice [21], by contrast, recommends mandatory certification of software and services. His proposal is clearly inspired by similar efforts in the area of food and traffic safety. A similar proposal is brought forward by Parameswaran and Whinston [20], yet with a tighter focus on network intermediaries, such as Internet Service Providers. Telang and Yang [26] empirically compare the number of vulnerabilities in software products with and without Common Criteria certification. They find that certified software fixes more old vulnerabilities, but

also contains more new vulnerabilities so that the net effect is neutral. All this literature has in common that audits and certification are regarded as tools to overcome information asymmetries. Interdependent security is not reflected. Since our work exclusively deals with solutions to the coordination problem in the presence of interdependence, it complements this strand of literature.

Modeling interdependent risks has quite some tradition in the field of security economics. Varian [27] as well as Kunreuther and Heal [16] belong to the second category of related work. Both teams promoted the view of information security as a public good, suggested formal models, and thus coined the notion of interdependent security. Our model is closer to Kunreuther and Heal. Varian's approach is richer if more than two firms interact. He adopts three types of aggregation functions from the economics literature of public goods [13]: weakest link, total effort, and best shot. Grossklags et al. [12] take up this idea and extend it in a series of works. The key difference to our model is the assumption of two kinds of security investments, one that generates externalities and another one that does not. Most models of interdependence are designed with the intention to find ways to internalize the externalities. This has led to literature for different contexts, including for instance cyber-insurance [19, 15] or security outsourcing with [29] and without [22] risk transfer. The availability of security audits is sometimes assumed (e. g., in [25]), but their effectiveness is never scrutinized.

5 Discussion

Few analytical models with three parameters can quantitatively predict outcomes in reality. Nevertheless, the interaction of security productivity, degree of interdependence, and thoroughness of security audits in our model allows to draw new conclusions. These conclusions can be transferred to practical situations at least qualitatively using the insights about the underlying mechanics.

5.1 Technical Implications

Region A covers more than half of the parameter space, including all settings with low or moderate degree of interdependence ($\alpha < 1/2$). Even if the parameters are not exactly measurable in practice, the conclusions for region A can serve as rules of thumb. A relevant insight is that security audits and certifications at very low security levels are often ineffective. This stands in stark contrast to a plethora of (largely commercial) security seals that certify the "lowest common denominator". Engineers who develop audit standards and supporting tools should rather focus on the possibility to extract verifiable information about high and highest security levels.

Another result of our analysis is that the effectiveness of security audits is very sensitive to the situation. A practical conclusion is that security standards and audit procedures should best be designed in a modular manner to allow tailored examinations. At this point we can only speculate if, say, the seven Evaluation Assurance Levels laid down in the Common Criteria are sufficient, or

whether a more granular choice of audit thoroughness is needed. Tailored audits may also require technical prerequisites which need to be considered in the design of the system to be audited. Last but not least, if auditability matters, then technical measures which imply changes to the parameters α and β (e. g., change of architecture, security technology, or interconnectivity) should be evaluated with regard to the availability of appropriate audit procedures.

5.2 Managerial and Regulatory Implications

Interdependent security risks exhibit a special and non-trivial mechanic. This mechanic prevents that individually rational risk management decisions also lead to socially optimal outcomes. A first important step is to explain this mechanic to managers and regulators. This way, they can adapt their decisions and refrain from blindly commissioning or requesting security audits. Our analysis has shown that security audits with bad fit to the situation are often inefficient or useless. For example, voluntary (i. e., unilateral) security audits certifying a very basic level of security ($s > 0$) are unnecessary in the large majority of cases. By contrast, audits can be very effective if they require relatively little thoroughness—and thus presumably little cost—to stabilize an equilibrium at a substantially higher level of security. This is the case in region B (see Sect. 3.2). Another insight is that very thorough security audits, which attest highest security levels, should only be conducted bilaterally in mutual agreement between partners. This is the only way to effectively prevent free-riding.

Regulators should analyze carefully in which situations they require mandatory security audits of what thoroughness. Most importantly, mandatory audits seem unnecessary in situations where the firms have own incentives to conduct security audits. It goes without saying that security audits should not be required when they are useless. To prevent this, it might be reasonable to replace general audit requirements with more specific sets of rules that consider factors of the firm and its environment. If these criteria are transparent, market participants can choose, say, whether they reduce the degree of interdependence *or* be subject to more thorough security audits.

A challenge remains with the definition and measurement of practical indicators to guide decision support. Neither the degree of interdependence nor the security productivity is observable on the scales that appear in the model. Since this task requires comparable and partly sensitive data of many market participants, we see this task in the responsibility of the government.

5.3 Conclusion

We have presented a novel analytical model to study the effectiveness of security audits as tools to incentivize the provision of security by private actors at a socially optimal level. The model takes parameters for the efficiency of security investment in risk reduction (security productivity), the exposure to risk from other peers in a network (degree of interdependence), and the thoroughness of the security audit. The solution of this model reveals that security audits must

be tailored to the very situation in order to avoid that they are ineffective. Moreover, "lightweight" security audits certifying a minimum level of security are not socially beneficial in the large majority of cases. Our results call for the revision of policies that require mandatory and undifferentiated security audits.

References

1. R. Anderson, R. Böhme, R. Clayton, and T. Moore. *Security Economics and the Internal Market*. Study commissioned by ENISA, 2008.
2. R. J. Anderson. Why information security is hard – An economic perspective, 2001.
3. M. Armbrust et al. Above the clouds: A Berkeley view of cloud computing. Technical Report EECS–2009–28, University of California, Berkeley, 2009.
4. R. Axelrod. *The Evolution of Cooperation.* Basic Books, New York, 1984.
5. M. R. Baye and J. Morgan. Red queen pricing effects in e-retail markets. Working Paper, 2003.
6. R. Böhme. Security metrics and security investment models. In I. Echizen, N. Kunihiro, and R. Sasaki, editors, *Advances in Information and Computer Security (IWSEC 2010)*, LNCS 6434, pages 10–24, Berlin Heidelberg, 2010. Springer-Verlag.
7. R. Böhme and T. W. Moore. The iterated weakest link: A model of adaptive security investment. In *Workshop on the Economics of Information Security (WEIS)*, University College London, UK, 2009.
8. E. Brynjolfsson and L. Hitt. Computing productivity: Firm-level evidence. *The Review of Economics and Statistics*, 85(4):793–808, 2003.
9. N. G. Carr. IT doesn't matter. *Harvard Business Review*, 81(5):41–49, 2003.
10. B. Edelman. Adverse selection in online "trust" certifications. In *Workshop on the Economics of Information Security (WEIS)*, University of Cambridge, UK, 2006.
11. L. A. Gordon and M. P. Loeb. The economics of information security investment. *ACM Trans. on Information and System Security*, 5(4):438–457, 2002.
12. J. Grossklags, N. Christin, and J. Chuang. Secure or insure? A game-theoretic analysis of information security games. In *Proc. of the Int'l Conference on World Wide Web (WWW)*, pages 209–218, Beijing, China, 2008. ACM Press.
13. J. Hirshleifer. From weakest-link to best-shot: The voluntary provision of public goods. *Public Choice*, 41:371–386, 1983.
14. A. Jacquith. *Security Metrics: Replacing Fear, Uncertainty, and Doubt.* Addison-Wesley, 2007.
15. B. Johnson, R. Böhme, and J. Grossklags. Security games with market insurance. In J. S. Baras, J. Katz, and E. Altmann, editors, *Decision and Game Theory for Security*, volume 7037 of *LNCS 7037*, pages 117–130, Berlin Heidelberg, 2011. Springer-Verlag.
16. H. Kunreuther and G. Heal. Interdependent security. *Journal of Risk and Uncertainty*, 26(2-3):231–49, March–May 2003.
17. W. Liu, H. Tanaka, and K. Matsuura. An empirical analysis of security investment in countermeasures based on an enterprise survey in Japan. In *Workshop on the Economics of Information Security (WEIS)*, University of Cambridge, UK, 2006.
18. D. Molnar and S. Schechter. Self hosting vs. cloud hosting: Accounting for the security impact of hosting in the cloud. In *Workshop on the Economics of Information Security (WEIS)*, Harvard University, Cambridge, MA, 2010.

19. H. Ogut, N. Menon, and S. Raghunathan. Cyber insurance and it security investment: Impact of interdependent risk. In *Workshop on the Economics of Information Security (WEIS)*, Harvard University, Cambridge, MA, 2005.
20. M. Parameswaran and A. B. Whinston. Incentive mechanisms for internet security. In H. R. Rao and S. Upadhyaya, editors, *Handbooks in Information Systems*, volume 4, pages 101–138. Emerald, 2009.
21. D. Rice. *Geekonomics – The Real Cost of Insecure Software*. Addison-Wesley, New York, 2007.
22. B. R. Rowe. Will outsourcing IT security lead to a higher social level of security? In *Workshop on the Economics of Information Security (WEIS)*, Carnegie Mellon University, Pittsburgh, PA, 2007.
23. S. Sackmann, J. Strüker, and R. Accorsi. Personalization in privacy-aware highly dynamic systems. *Communications of the ACM*, 49(9):32–38, 2006.
24. T. Schelling. *The Strategy of Conflict*. Oxford University Press, Oxford, UK, 1965.
25. N. Shetty, G. Schwartz, M. Felegyhazi, and J. Walrand. Competitive cyber-insurance and internet security. In *Workshop on Economics of Information Security (WEIS)*, University College London, UK, 2009.
26. R. Telang and Y. Yang. Do security certifications work? Evidence from Common Criteria certification. In *IEEE International Conference on Technologies for Homeland Security*, 2011.
27. H. R. Varian. System reliability and free riding. In *Workshop on the Economics of Information Security (WEIS)*, University of California, Berkeley, 2002.
28. S. Winkler and C. Proschinger. Collaborative penetration testing. In *Business Services: Konzepte, Technologien, Anwendungen (9. Internationale Tagung Wirtschaftsinformatik)*, volume 1, pages 793–802, 2009.
29. X. Zhao, L. Xue, and A. B. Whinston. Managing interdependent information security risks: A study of cyberinsurance, managed security service and risk pooling. In *Proceedings of ICIS*, 2009.

A Proof Sketches

A.1 Formal Conditions of Equilibria

Border Between Region A and B

Idea: Take fixed point from Eq. (12) and set it to zero,

$$\alpha = 1 - \log^{-1}(\beta).$$

Border Between Region B and C

Idea: Take determinant of Eq. (12) and set it to zero,

$$\beta = e^{4\alpha}.$$

Border Between Region C and D

Idea: Set $c(s^*) = c(0)$ (from Eqs. (3) and (7)),

$$s^* - \left(1 - \beta^{-s^*}\right)\left(1 - \alpha\beta^{-s^*}\right) = 0.$$

A.2 Social Optima

Start with Eq. (6):

$$s^* = \arg\min_s s + 1 - (1 - \beta^{-s})(1 - \alpha\beta^{-s})$$

Root of first-order condition of s:

$$\begin{aligned}
1 &= \log(\beta)\left(1 - \alpha\beta^{-s^*}\right)\beta^{-s^*} + \alpha\log(\beta)\left(1 - \beta^{-s^*}\right)\beta^{-s^*} \\
1 &= \log(\beta)\beta^{-s^*} + \alpha\log(\beta)\beta^{-s^*} - 2\alpha\log(\beta)\beta^{-2s^*} \\
1 &= (1+\alpha)\log(\beta)\beta^{-s^*} - 2\alpha\log(\beta)\beta^{-2s^*}
\end{aligned}$$

Case 1: $\alpha = 0$

$$\begin{aligned}
1 &= \log(\beta)\beta^{-s^*} \\
s^* &= \log(\log(\beta))\log^{-1}(\beta)
\end{aligned}$$

This expression corresponds to Eq. (8).

Case 2: $\alpha > 0$. We obtain the root by substituting $u = \beta^{-s^*}$, solving the quadratic equation, and subsequent resubstitution:

$$s^* = -\log\left(\frac{(1+\alpha) - \sqrt{(1+\alpha)^2 - 8\alpha\log^{-1}(\beta)}}{4\alpha}\right)\log^{-1}(\beta)$$

This expression corresponds to Eq. (7).

A.3 Best Response

Start with Eq. (10):

$$\begin{aligned}
s^+(s_{1-i}) = \arg\min_s\ & s + 1 - \left(1 - \beta^{-s}\right)\left(1 - \alpha\beta^{-s_{1-i}}\right) \\
\text{s.t.}\quad & s \geq 0
\end{aligned}$$

Root of first-order condition of s:

$$\begin{aligned}
0 &= 1 - \log(\beta)\beta^{-s^+}\left(1 - \alpha\beta^{-s_{1-i}}\right) \\
1 &= \log(\beta)\beta^{-s^+}\left(1 - \alpha\beta^{-s_{1-i}}\right) \\
\beta^{s^+} &= \log(\beta)\left(1 - \alpha\beta^{-s_{1-i}}\right) \\
s^+\log(\beta) &= \log\left(\log(\beta)\right) + \log\left(1 - \alpha\beta^{-s_{1-i}}\right)
\end{aligned}$$

Rearrangement subject to constraints:

$$s^+ = \sup\left\{\frac{\log\left(\log(\beta)\right) + \log\left(1 - \alpha\beta^{-s_{1-i}}\right)}{\log(\beta)}, 0\right\}$$

This expression corresponds to Eq. (11).

A.4 Nash Equilibria

Fixed points of the best response $\tilde{s} = s^+(\tilde{s})$ without considering constraints:

$$\begin{aligned}
\tilde{s} &= \frac{\log\left(\log(\beta)\right) + \log\left(1 - \alpha\beta^{-\tilde{s}}\right)}{\log(\beta)} \\
\tilde{s}\,\log(\beta) &= \log\left(\log(\beta)\right) + \log\left(1 - \alpha\beta^{-\tilde{s}}\right) \\
\log(\beta^{\tilde{s}}) &= \log\left(\log(\beta)\right) + \log\left(1 - \alpha\beta^{-\tilde{s}}\right) \\
\log(\beta^{\tilde{s}}) &= \log\left(\log(\beta)\left(1 - \alpha\beta^{-\tilde{s}}\right)\right) \\
\log(\beta^{\tilde{s}}) &= \log\left(\log(\beta) - \alpha\beta^{-\tilde{s}}\log(\beta)\right) \\
\beta^{\tilde{s}} &= \log(\beta) - \alpha\beta^{-\tilde{s}}\log(\beta) \\
\beta^{\tilde{s}} - \log(\beta) &= -\alpha\beta^{-\tilde{s}}\log(\beta) \\
\beta^{2\tilde{s}} - \log(\beta)\beta^{\tilde{s}} &= -\alpha\log(\beta) \\
\beta^{2\tilde{s}} - \log(\beta)\beta^{\tilde{s}} + \alpha\log(\beta) &= 0
\end{aligned}$$

We obtain the root by substituting $u = \beta^{\tilde{s}}$, solving the quadratic equation, and subsequent resubstitution:

$$\tilde{s}_{1,2} = \log\left(\frac{\log(\beta) \pm \sqrt{\log^2(\beta) - 4\alpha\log(\beta)}}{2}\right)\log^{-1}(\beta)$$

This expression corresponds to Eq. (12).

Fixed points are Nash equilibria if they fulfill the constraint $\tilde{s} > 0$. Because of the constraint in Eq (10), there exists another corner equilibrium at $\tilde{s}_3 = 0$ if $s^+(0) = 0$:

$$\begin{aligned}
0 &\geq \frac{\log\left(\log(\beta)\right) + \log\left(1 - \alpha\right)}{\log(\beta)} \\
1 &\geq \log(\beta)\left(1 - \alpha\right) \\
\alpha &\geq 1 - \log^{-1}(\beta)
\end{aligned}$$

This expression corresponds to Eq. (13).