2024 IEEE 63rd Conference on Decision and Control (CDC)
December 16-19, 2024. MiCo, Milan, Italy

How to Pass an Audit: Decisions on Sequential Investment

Ziyuan Huang, Gergely Biczók, Mingyan Liu

Abstract— **We study the following sequential decision problem: a vendor with a product or process needs to pass a mandatory audit in order to be able to release the product onto the market; it is allowed to go through the audit repeatedly, and thus the vendor needs to determine what level of effort to put into the product (e.g., to enhance its quality/performance) following each failed audit. We examine the vendor's optional decision process and fully characterize its properties under mild technical assumptions. We next examine what happens if the audit is optional, and to incentivize the vendor to voluntarily subject itself to the audit, the auditor offers a waiver from future product liabilities provided the audit is successful. We examine what type of audit might be the most effective in not only incentivizing participation but also more desirable effort from the vendor.**

I. Introduction

In this paper, we study a sequential decision problem faced by a vendor with a product or process going through an audit: it needs a successful audit in order to bring the product to the market and can go through the audit repeatedly following each failure. The audit mechanism is known to the vendor, and it thus needs to determine what level of effort to put into the product (e.g., to enhance its quality/performance) following each failure. The audit is assumed to be informative but not perfect, with some built-in randomness. We examine this problem in two settings: the audit is mandatory or optional; in the latter case, to incentivize the vendor to voluntarily subject itself to the audit, the auditor offers a waiver from future product liabilities provided the audit is successful. In the first case, we examine the vendor's optional decision process and fully characterize its properties under mild technical assumptions. In the second case, we examine what type of audit might be the most effective in incentivizing not only participation but also the desirable effort of the vendor.

This study is primarily motivated by the pressing issue of software security facing the software industry. Making software products more secure is arguably one of the most important and challenging elements in securing our overall computer network ecosystem. For one, security features in a software product can be hard to monetize; for another, a software vendor's exposure to potential security risk is limited by the fact that the vast majority of the cost incurred in a security incident is borne by the consumer of the software. All these facts diminish the vendor's motivation to minimize software vulnerabilities a priori, leading to the problem of moral hazard [1]. This also motivates a series of national security policy directives that aim to allow liability claims against insecure software products. Specifically, the United States National Cybersecurity Strategy (USNCS), released in April 2023 [2], has proposed the possibility of a liability waiver mechanism tied to government-mandated security audits, serving as a financial incentive for software companies to improve their product security practices.

Existing literature on audit assumes strategic auditors. [3]–[6] model auditing process with simultaneous-move games between an auditor and an auditee. The auditor inputs costly auditing design efforts aiming to achieve the highest accuracy net design costs, while the auditee tries to pass the audit with minimum investment costs. Similar approaches were adopted by [7], [8] on carefully designed finite sequential audit games. The auditing problem has also been studied in the insurance literature [9]–[11] where the insurer acts as the auditor trying to maximize its profit less the audit cost.

Our model relaxes the budget requirement on auditing efforts, enabling indefinitely repeated interactions between the auditor and the auditee (vendor). We also do not assume strategic auditors in this work so as to focus on understanding the impact of the audit structure. This paper is the first attempt at modeling enforced repeated audits in the software industry, which are more suitably implemented by regulatory authorities such as governments as suggested by the USNCS [2], in contrast to profit-maximizing audit models adopted by the vast insurance and accounting literature.

Our goal in this paper is to understand: (1) from the vendor's perspective, what is the optimal sequence of investment when it is under audit, (2) from the auditor's perspective, how to maximally incentivize vendors to opt into such an audit. Our main findings are as follows:

- The optimal policy for the vendor under audit is, in general, non-unique but enjoys some very interesting properties. An optimal policy falls into two broad categories: the "one-and-done" type and the "incremental" type (Section III). Under the first type, the vendor invests in one installment at the beginning of the process, an amount from a well-defined optimal set, prior to the initial audit, and then waits to pass the audit, even if it takes multiple rounds. Under the second type, the vendor invests multiple times with each amount from the same optimal set, but the timing of these investments can be arbitrary.
- We show how the audit quality (accuracy and hardness) influences the vendor's participation incentive and how to adjust these parameters to increase participation (Section IV).

Ziyuan Huang and Mingyan Liu are with the Electrical and Computer Engineering Department, University of Michigan, 1301 Beal Avenue, Ann Arbor, MI 48109, USA; e-mail: {ziyuanh, mingyan} @umich.edu; Gergely Biczók is with the CrySyS Lab, Department of Networked Systems and Services, Budapest University of Technology and Economics, 1111 Budapest, Műegyetem rkp. 3, Hungary; e-mail: biczok@crysys.huu.

II. Model and Preliminaries

The basic problem consists of a neutral auditor with a pre-determined, publicly known audit rule and a utility-maximizing (software) vendor responding optimally to the audit rule. The process plays out in discrete time and over multiple periods, as the vendor may need to be audited repeatedly to pass. followed by the solution approach.

A. Problem Description

Suppose a vendor has to satisfy a certain auditing requirement to enter the market. The audits are executed repeatedly and stop only when the vendor passes the audit or quits the market entirely. The rationale for the latter assumption is that a product with known defects/vulnerabilities should not be allowed to enter the market. The vendor is utility maximizing and potentially risk averse, optimizing over its (successive) investment/effort levels (e.g., to enhance the security of the product). We shall assume that the audit service is free but not perfect, i.e., the audit outcome may contain both false positives and false negatives. More importantly, we assume that successive audits in the process are independent.

The Vendor: From the vendor's point of view, the auditing process can be modeled as an infinite-horizon discrete dynamical system. Index the time steps of this process by $t \in \{0,1,\dots\}$. The vendor must determine and commit to an investment $x_0 \in \mathbb{R}_+$ at $t = 0^+$. The value x_0 is private to the vendor; however, since its utility function is assumed public knowledge, the vendor's optimal strategy, including the value x_0, is ultimately known to the auditor (following the same computation). The effort x_0 goes into the software development over the first time step, incurring a cost of $c_0 = C(x_0)$ by $t = 1^-$, when the product is submitted for audit, with its outcome revealed at time $t = 1$. $C(\cdot)$ is assumed positive and continuous. Audit outcome at t is denoted by $s_t \in \{0,1\}$. If the outcome is positive/successful, i.e., $s_1 = 1$, the process terminates: the vendor is granted market entry at $t = 1^+$, earning a reward/revenue r_1, assumed to be R. This is the terminal reward for passing the audit with discounting applied through the utility function given shortly below.

If the audit outcome is negative/fail, i.e., $s_1 = 0$, then the vendor is temporarily denied market entry. It can either choose to quit the process (exit the market) at $t = 1^+$, thereby receiving zero revenue but incurring no further cost, or opt for re-auditing. In the latter case, the vendor must decide a new (cumulative) effort level $x_1 \in [x_0, \infty)$ at $t = 1^+$, thereby committing to an *additional* investment of $x_1 - x_0$ over the next time step. This incurs a cost of $c_1 = C(x_1) - C(x_0)$ at $t = 2^-$. The process then proceeds to stage $t = 2$ and repeats indefinitely until the vendor either passes the audit or decides to quit. For any non-terminal stage, it generates a reward of $r_t = 0$. Let q_t denote the binary continuation decision with $q_t = 1$ indicating a quit. The sequence of actions and decisions is illustrated in Fig. 1.

Formally, define $\underline{x} := \{x_t\}_{t=0}^{\infty}$ as the increasing sequence of total (cumulative) investments by the vendor at time t^+. Let $\tau_q \in \{1,2,\dots\}$ denote the quitting time and $\tau_s \in \{1,2,\dots\}$ the time at which the vendor first passes the audit; both are in

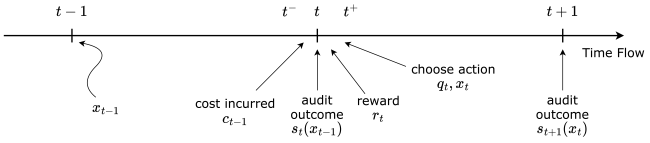

Fig. 1: Timeline of the decision process, assuming the process has not stopped by $t+1$.

general random (stopping) times of the processes $\{q_t\}_{t\geq 0}$ and $\{s_t\}_{t\geq 0}$ respectively. The vendor's utility, given its decision on $\underline{x}$ and τ_q and the discount factor α, is:

$$U(\underline{x}, \tau_q) = \mathbb{E}_{\tau_s}\left[\sum_{t=0}^{\min\{\tau_s,\tau_q\}-1} \alpha^t \left(r_{t+1} - c_t\right)\right]. \tag{1}$$

Note that $r_{\tau_s} = R$ and $r_{\tau_q} = 0$ as described earlier. The vendor's goal is to maximize its utility in Eqn (1). Denote the maximal utility by $U^* := \sup_{\underline{x},\tau_q} U(\underline{x}, \tau_q)$.

The Auditor: The auditor is modeled as a neutral (without its own utility function) party defined by the quality of its audit: the sequence of functions $\underline{p}(\underline{x}) := \{p_t(x_{t-1})\}_{t=1}^{\infty}$, where $p_t(x_{t-1})$ is the probability of the product passing audit at time t given cumulative effort x_{t-1}.

Assumption 1: The audit is static, defined by $p_t(x) = p(x)$ $\forall t$ where $p : \mathbb{R}_+ \to [0,1]$ is continuous and increasing.

B. Preliminaries

In light of Fig. 1, the vendor's decision process can be formulated as a discounted-reward Markov decision process (MDP). Let $e_t \in \{0,1\}$ denote the *continuation state* of the process: $e_t = 1$ if the process has terminated by time t (inclusive), and $e_t = 0$ if the process proceeds into t^+. Define states $z_t := (e_t, x_{t-1}) \in \mathscr{Z} := \{0,1\} \times \mathbb{R}_+$ for $t = 1,2,\dots$ where x_{t-1} is the cumulative investment over the first $t-1$ steps. Given the current state z_t, the vendor chooses an action $u_t := (q_t, a_t) \in \mathscr{U} := \{0,1\} \times \mathbb{R}_+$ where $q_t = 1$ when the vendor decides to quit at this stage and a_t represents the vendor's additional investment in case of continuation.

Define an alternative instantaneous reward function as

$$\begin{aligned}\rho(z_t, u_t) = \big[&p(x_{t-1}+a_t)R - \\ &(C(x_{t-1}+a_t) - C(x_{t-1}))\big]\mathbf{1}_{\{0\}}(e_t)\mathbf{1}_{\{0\}}(q_t).\end{aligned} \tag{2}$$

This is the expected payoff (reward minus cost) that the vendor earns at time t. It is zero when either the process has stopped ($e_t = 0$) or the vendor decides to quit ($q_t = 1$).

The state at time $t+1$ can be updated by the tuple (z_t, u_t):

$$x_t = x_{t-1} + a_t \;\text{ and }\; e_{t+1} = \begin{cases} w_t & e_t = 0 \text{ and } q_t = 0 \\ 1 & e_t = 1 \text{ or } q_t = 1 \end{cases}$$

where $w_t \sim \text{Bernoulli}(p(x_{t-1}+a_t))$. So, this system is a valid MDP by construction.

Let $\pi := \{u_t\}_{t=1}^{\infty}$ denote an arbitrary policy. The *expected total discounted reward* with initial state z under π is

$$\begin{aligned} V^{\pi}(z) &= \mathbb{E}_{\tau_s}\left[\sum_{t=0}^{\tau_s \wedge \tau_s^{\pi}-1} \alpha^t \rho(z_t^{\pi}, u_t^{\pi}) \,\middle|\, z_0^{\pi} = z\right] \\ &= \mathbb{E}\left[\sum_{t=0}^{\infty} \alpha^t \rho(z_t^{\pi}, u_t^{\pi}) \,\middle|\, z_0^{\pi} = z\right], \end{aligned} \tag{3}$$

where the superscript π emphasizes the dependence of relevant variables on the policy π. The second equality holds because $\rho(z_t^\pi, u_t^\pi) = 0$ if the process stopped before time t, i.e., $e_t = 1$. The goal of the MDP is to find the optimal policy π that maximizes the objective in Eqn (3). Denote the *optimal reward function* as $V^*(z) := \sup_\pi V^\pi(z)$ and the optimal policy as $\pi^*(z) \in \arg\sup_\pi V^\pi(z)$.

We will only focus on non-terminal states as $V^*(1,x) \equiv 0$ by Eqn (2) and (3). With a slight abuse of notation, we will denote $V^*(x) := V^*(0,x)$. Additionally, without loss of generality, we will only consider stationary policies, i.e., state-dependent and time-invariant, that depend on the state z_t only through x_{t-1}. In other words, we are only interested in functions of the form $g : \mathbb{R}_+ \to \mathscr{U}$ s.t. $u_t = g(x_{t-1})\ \forall t \geq 0$.

Using the notation above, we can express the vendor's optimal utility as follows:

$$U^* = V^*(0) - C(0). \tag{4}$$

The methodology used to compute V^* is the Bellman equation. By Theorem 2.2 in [12], V^* is the unique solution to the following fixed-point (Bellman) equation,

$$V^*(x) = \sup_{u \in \mathscr{U}} \rho(e=0,x,u) + \alpha \mathbb{E}\Big[V^*(z') \Big| e=0,x,u\Big]. \tag{5}$$

where z' represents the next state.

III. The Optimal Strategy Under Audit

We now characterize the optimal strategy of the vendor. We begin by determining when it is optimal to quit and then derive its optimal investments using the Bellman equation.

A. Optimal Quitting Time

We expand $V^*(x)$ using the quitting decision q_t as

$$V^*(x) = \max\Big\{0,\ \sup_{y \geq x} C(x) - C(y) + p(y)R + \alpha(1-p(y))V^*(y)\Big\}, \tag{6}$$

where the first term in the *max* operator is the maximum reward-to-go for quitting, i.e., $q = 1$, and the second term is for continuation. Notice,

$$\sup_{y \geq x} C(x) - C(y) + p(y)R + \alpha(1-p(y))V^*(y) \geq\ p(x)R + \alpha(1-p(x))V^*(x) \geq 0,$$

where the last inequality utilized $V^*(x) \geq 0$ by Eqn (6). This directly leads to the following lemma.

Lemma 3.1: The vendor never quits in an optimal strategy.

B. Optimal Continuation Investment

Given that a vendor never quits, we can remove the *max* operator in Eqn (6) and express $V^*(x)$ more concisely as

$$V^*(x) = \sup_{y \geq x}\ C(x) - C(y) + p(y)R + \alpha(1-p(y))V^*(y).$$

Define $W(x) := V^*(x) - C(x)$. The above equation implies

$$W(x) = \sup_{y \geq x}\ -C(y) + p(y)R + \alpha(1-p(y))V^*(y). \tag{7}$$

The optimal additional investment a^* given the cumulative investments x thus satisfies

$$x + a^* \in \arg\sup_{y \geq x}\ -C(y) + p(y)R + \alpha(1-p(y))V^*(y). \tag{8}$$

There exists an optimal additional investment if the *arg sup* in the second term yields a non-empty set.

Lemma 3.2: W is decreasing in x.

We provide all proofs in the online version [13]. The monotonicity of W has a very interesting implication on the vendor's behavior. Suppose the cumulative investments up to t is x_{t-1} and the vendor chooses an a_t^* optimally according to Eqn (8), resulting in a new cumulative investment $x_t = x_{t-1} + a_t^*$. If the vendor fails the audit at this level, its optimal additional investment now becomes zero since the maximum in Eqn (7) is already obtained with x_t over $[x_t, \infty)$ by the monotonicity of W. Therefore, one of the vendor's optimal strategies given any cumulative investment (sunk cost) is to immediately invest at the optimal additional level and wait indefinitely until it passes the audit.

Lemma 3.3: The function W can be expressed as

$$W(x) = \sup_{y \geq x}\ G(y), \quad \forall x \geq 0, \tag{9}$$

where

$$G(y) := -C(y) + \frac{p(y)R}{1 - \alpha + \alpha p(y)}. \tag{10}$$

Suppose the optimal additional investment $a^*(x)$ exists for any cumulative investment x. Then, the set $\mathscr{G}_x := \arg\sup_{y \geq x} G(y)$ is non-empty for every x and the optimal additional investment satisfies

$$x + a^*(x) \in \mathscr{G}_x. \tag{11}$$

Comparing the maximum utility value in Eqn (4) and the definition of W in Eqn (7), we see that $U^* = W(0)$. Therefore, we can directly calculate the optimal (sequential) investments by evaluating the function $G(x)$.

Theorem 3.4: Suppose the optimal additional investment exists for any cumulative investment x. The vendor's optimal strategy under audit has to satisfy the following properties:

(1) the vendor will never quit;
(2) it is given by any non-decreasing sequence of cumulative investments $\{x_t\}_{t \geq 0}$, where $x_t \in \mathscr{G}_0$;
(3) the optimal utility is given by $\max_{x \geq 0} G(x)$.

Theorem 3.4-(2) says that the optimal strategy is in general non-unique, but it must fall into two broad categories. The first type is such that the vendor invests any amount $x \in \mathscr{G}_0$ at stage 0 followed by nothing else in subsequent stages, essentially waiting for the audit to return a positive outcome (which is guaranteed to occur with high probability given our assumptions on the audit process). This class of strategies can be referred to as the "patient" type, deciding on a total expenditure and then waiting it out. In particular, those that invest in the smallest amount (the smallest x in $\mathscr{G}_0$) are foregoing revenue (due to the long expected time it takes

to pass the audit, the eventual revenue would be severely discounted) in exchange for a small initial investment.

The other class of optimal strategies involves investing at two or more different times, each time reaching some cumulative amount $x \in \mathscr{G}_0$. When these investments are made is arbitrary, provided the first occurs at time 0. This type of strategy is more impatient or opportunistic: they invest a small amount initially, hoping to pass the audit on good luck; when that doesn't happen for some time, they decide to up the game and invest more hoping to pass the audit this time, and so on.

It is important to emphasize that both types of strategies yield the same utility under our model; they essentially reflect different tradeoffs between willingness to invest vs. willingness to wait for return on investment.

Those who invest the largest amount $x \in \mathscr{G}_0$ at time 0 necessarily belong to the first type, as there is no more feasible action left given the non-decreasing nature of the sequence. These are the "ideal" or most desirable vendors from a public interest or social welfare perspective – they invest the maximum amount in one go, resulting in the highest quality product. In the next section, we will discuss which configurations of the auditing process can help lead to this type of strategy.

The following example shows more concretely the property of the optimal strategies given in Theorem 3.4-(2) and discussed above.

Example 1: (Property of the Optimal Sequence of Cumulative Investments) Suppose $G(x)$ contains two local maximizers x_L and x_H with $x_L < x_H$, illustrated in Fig. 2. If $G(x_L) \neq G(x_H)$, then there is a single global maximum, and the optimal investment strategy is unique: investing at the global maximum level at time 0 and nothing else thereafter.

If $G(x_L) = G(x_H)$, then the set $\mathscr{G}_0$ in Theorem 3.4 contains exactly 2 values x_L and x_H. By Theorem 3.4-(2), every optimal investment sequence should start with $x_0 \in \{x_L, x_H\}$. In other words, the initial action $a_0^* \in \{x_L, x_H\}$.

If an optimal strategy starts with $x_0 = x_H$, then all subsequent cumulative investments remain at x_H, i.e., no additional investment in the future. In this case, the optimal cumulative investment sequence and the optimal action sequence are,

$$\{x_t\}_{t\geq 0} = \{x_H, x_H, x_H, \dots\} \text{ and } \{a_t^*\}_{t\geq 0} = \{x_H, 0, 0, \dots\}.$$

As x_H is the largest element in $\mathscr{G}_0$, $\{x_H, x_H, x_H, \dots\}$ is the unique non-decreasing sequence of cumulative investments in $\mathscr{G}_0$ given $x_0 = x_H$. This strategy minimizes the vendor's expected time for passing the audit.

If an optimal strategy starts with $x_0 = x_L$, then this can lead to either type of optimal sequence. Under the first type, the vendor invests nothing more beyond the initial amount:

$$\{x_t\}_{t\geq 0} = \{x_L, x_L, x_L, \dots\} \text{ and } \{a_t^*\}_{t\geq 0} = \{x_L, 0, 0, \dots\}.$$

Under the second type, the vendor invests an additional $x_H - x_L$ at some arbitrary future time $t \geq 1$, i.e.,

$$\{x_t\}_{t\geq 0} = \{x_L, x_L, \dots, x_L, x_H, \ x_H, \dots\},$$
$$\{a_t^*\}_{t\geq 0} = \{x_L, 0, \dots, 0, x_H - x_L, 0, \dots\}.$$

As mentioned, the exact time when the additional investment $x_H - x_L$ is made has no impact on the strategy's optimality.

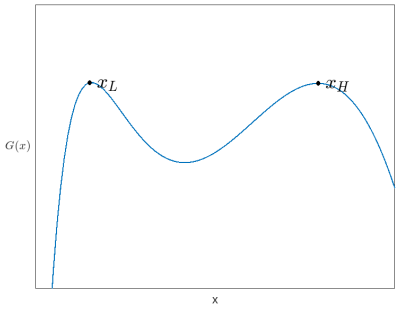

Fig. 2: Shape of $G(x)$ with two local maxima x_L and x_H with $x_L < x_H$; both attain the same global maximum.

IV. Incentivizing Audit Participation

The previous section focuses on the vendor's optimal strategy to pass the audit. We now turn to the question of how the audit mechanism should be designed to encourage participation by the vendo. This is the *voluntary participation* problem extensively studied in the mechanism design literature [14], [15]. We will continue to assume that once the vendor chooses to go through the audit, it cannot release the product until it passes the audit. In return, passing the audit relieves the (vendor of the) product from potential future liabilities (i.e., it obtains a liability waiver). One might equivalently view the audit as a certification process, whereby passing the audit earns the vendor a stamp of approval that carries certain reputational or pricing benefits.

On the other hand, if the vendor decides to opt out, then it will bear the cost of any potential liability on its product in the case of an adverse event (or equivalently, it will not obtain the reputational benefit of certification). Thus the availability of such an audit service may be viewed as a type of mechanism aimed at incentivizing a vendor to increase its effort and subject its product to audit. The central question is then under what conditions would a vendor voluntarily participate in such a mechanism, and whether the audit can induce better/higher effort from the vendor.

If the vendor opts out of the auditing mechanism, its optimal action is to choose some $x \in \mathbb{R}_+$ at $t = 0^+$ that maximizes the following *expected* opt-out utility which is attained at time $t = 1$ when the product is put on the market:

$$U^{\text{out}}(x) := R - C(x) - C_L(x), \tag{12}$$

where $C_L(x)$ represents the potential liability loss. While the reward and development cost terms are assumed deterministic, losses are random in general. Thus, $C_L(x)$ denotes the expected loss perceived by the vendor, with potential risk aversion built in. We discussed this next.

Specific Functions Used in the Analysis: Denote the actual monetary liability loss (in USD) by the random variable $Z(x)$, assumed to follow a normal distribution with mean $\mu_Z(x)$ and standard deviation $\sigma_Z(x)$. We will assume $\mu_Z(x)$ and $\sigma_Z(x)$ are both positive and decreasing in x, i.e., higher effort reduces the expected loss and the uncertainty in the loss.

To capture the vendor's risk aversion, we will model the liability cost that enters into the vendor's utility function as $C_L(x) := \mathbb{E}\exp(\gamma Z(x))$, where $\gamma > 0$ represents the vendor's risk attitude. Then, we can write $C_L(\gamma, x)$ as

$$C_L(\gamma, x) = \exp\left(\gamma\mu_Z(x) + \frac{1}{2}\gamma^2\sigma_Z^2(x)\right). \tag{13}$$

We model the audit rule $p(\cdot)$ as an estimation process, whereby the auditor predetermines a threshold δ and estimates whether the vendor's effort exceeds it. It follows that the estimate, given the vendor's effort x, can be represented as a random variable $Y := x + W$ where $W \sim \mathcal{N}(0, \sigma^2)$. The presence of noise highlights the fact that any audit cannot be perfect. The independent normal assumption models various unknown sources of uncertainty; its variance represents the *accuracy* or *quality* of the audit: a more accurate one has higher certainty. The probability of passing the audit is $p(x) = \mathbb{P}(Y \geq \delta) = \mathbb{P}(W \geq \delta - x) = 1 - \Phi((\delta - x)/\sigma)$, where $\Phi(\cdot)$ is the CDF of the standard normal distribution.

The audit is only *meaningful* or *informative* if it is correct more than 50% of the time. The above threshold model is indeed meaningful: if $x \geq \delta$, then $p(x) \geq 1 - \Phi(0) = \frac{1}{2}$; if $x < \delta$, then $p(x) < 1 - \Phi(0) = \frac{1}{2}$.

For the cost of investment, we adopt a linear form where the marginal cost of investment for the vendor is constant, i.e., $C(x) := c \cdot x$ for some $c > 0$. We note that the specific functional form of this marginal cost is not critical to the subsequent analysis, as our results hold without the linearity or even the monotonicity of this cost function [13].

While we do not model the auditor as a strategic agent, the vendor's strategy, and moreover, its choice of participation, is indeed influenced by the audit threshold δ and audit noise σ. Below, we first examine how these audit parameters impact the vendor's strategy when it opts in, and then show how they impact the vendor's decision to opt in vs. stay out. For clarity, we will refer to terms introduced prior to Section IV with the prefix "opt-in", such as "opt-in utility" and "opt-in strategy" to distinguish them from Problem (12).

A. On the Vendor's Optimal Opt-In Strategy

Results in Section III-B suggest there can potentially be many optimal strategies for an opt-in vendor, some starting at very low investment levels depending on the global maximizers of $G(x)$. While these are equally optimal by our model, the auditor may favor earlier and higher investments. Below we show that different choices of δ and σ can reshape $G(x)$ so as to induce more desired opt-in strategies.

Fig. 3 and 4 depict the shape of $G(x)$ under different values of δ and σ respectively, while keeping the other fixed[1]. The global maximum solutions of each curve correspond to the optimal opt-in efforts in that specific parameter setup, marked by solid stars in the figures. The main observations are:

1) From Fig. 3, we see when two local maximizers exist in $G(x)$, a high threshold (more difficult audit) causes the low solution $x = 0$ to dominate (it becomes the global maxima), whereas a low threshold (easier audit) causes the high solution to dominate. Also, a lower threshold always results in higher utility for the vendor, regardless of the effort. Thus, a high threshold somtimes encourages low effort as an optimal strategy. This seems counterintuitive; the reason is that a difficult audit poses a risk of failing the audit even at decent effort levels, so the vendor invests less and instead relies on waiting for a positive audit outcome to materialize by chance. A low threshold, on the other hand, reduces the need to gamble on the outcome of the audit and encourages the vendor to invest at the optimal (high) level from the start for a speedy pass.
2) From Fig. 4 we see that the shape of the $G(x)$ function is even more sensitive to the audit noise: a low noise (greenish curves) drives the high solution to become the global maximum (thus high investment as an optimal strategy), while a high noise (reddish curves) drives the low solution to be the global maximum (thus low investment as an optimal strategy).

In short, the above observations suggest that an accurate (low noise) but not overly difficult audit is the best choice: it minimizes opportunistic behavior and reliance on chance and encourages higher levels of effort early on in the process.

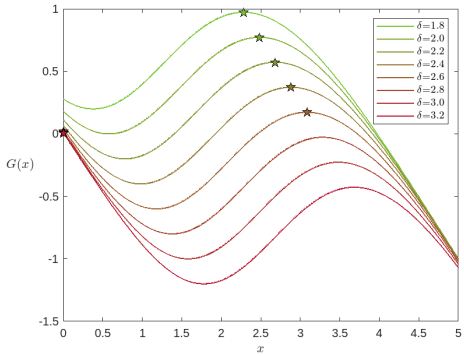

Fig. 3: $G(x)$ with varying δ under fixed $\sigma = 1$. The stars mark the (unique) global maximum of $G(x)$.

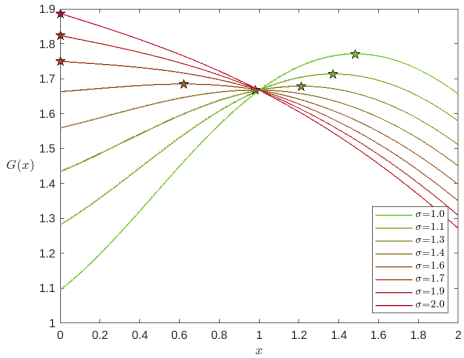

Fig. 4: $G(x)$ with varying σ under fixed $\delta = 1$. The stars mark the (unique) global maximum of $G(x)$.

B. On the Vendor's Choice of Participation

Clearly, the vendor only has the incentive to participate in the audit mechanism if $U^* \geq U^{\text{out},*}$, with ties broken in favor

[1]Other parameters are $c = 1$, $\alpha = 0.5$, $R = 4$.

of participation. To highlight the dependence of the vendor's opt-out utility on its risk attitudes, we write $U^{\text{out},*}(\gamma)$ instead of $U^{\text{out},*}$ with liability loss taking the form in Eqn (13).

Theorem 4.1: There exists $\overline{\gamma} \in [0,\infty]$ such that for $\gamma \geq \overline{\gamma}$, $U^{\text{out},*}(\gamma) \leq U^*$ and the vendor has an incentive to participate in the audit mechanism; for $\gamma < \overline{\gamma}$, $U^{\text{out},*}(\gamma) > U^*$ and the vendor prefers to stay outside. Specifically, when $\overline{\gamma} = \infty$, the vendor never participates; and when $\overline{\gamma} = 0$, the vendor always participates, regardless of the specification of the audit rule.

The value $\overline{\gamma}$ is the boundary risk attitude at which the vendor is indifferent between committing to the waiver/audit or not. Above this level, the vendor is relatively risk-averse and therefore interested in participating and transferring its risk to the auditor. Below this level, the vendor is relatively risk-seeking and does not have an incentive to participate.

As $\overline{\gamma}$ is nonnegative, participation increases with a lower $\overline{\gamma}$. Below we show how the auditor can lower $\overline{\gamma}$ by adjusting its auditing threshold and noise.

We will write the maximum opt-in utility as $U^*(\delta,\sigma)$ to emphasize its dependence on the auditing parameters. For each pair of (δ,σ), we can calculate $\overline{\gamma}$ by solving $U^{\text{out},*}(\overline{\gamma}) = U^*(\delta,\sigma)$ for $\overline{\gamma}$. We will similarly write it as $\overline{\gamma}(\delta,\sigma)$. Define the *coverage* of an audit mechanism with a fixed threshold δ as $\overline{\gamma}_\delta := \inf_{\sigma\geq 0}\overline{\gamma}(\delta,\sigma)$. Similarly, define the coverage associated with a fixed accuracy σ as $\overline{\gamma}_\sigma := \inf_{\delta\geq 0}\overline{\gamma}(\delta,\sigma)$. The coverage of the audit mechanism where both δ and σ are free variables is denoted as $\overline{\gamma}_0$.

Theorem 4.2:
(1) $\overline{\gamma}_0 = 0$, i.e., there exist audit mechanisms that ensure full coverage (for all vendor types).
(2) For fixed σ and $\delta_1 \leq \delta_2$, $\overline{\gamma}(\delta_1,\sigma) \leq \overline{\gamma}(\delta_2,\sigma)$ and $\overline{\gamma}_\sigma = \overline{\gamma}(0,\sigma)$. Thus, $\overline{\gamma}(\delta,\sigma)$ increases in δ and the maximum coverage is reached at $\delta = 0$. However, this is practically undesirable as $\delta = 0$ means a non-investing vendor; a behavior that should not be encouraged.
(3) $\overline{\gamma}_{\sigma_1} \leq \overline{\gamma}_{\sigma_2}$, $\forall \sigma_1 \leq \sigma_2$. Thus, higher accuracy increases the coverage by attracting less risk-averse vendors.

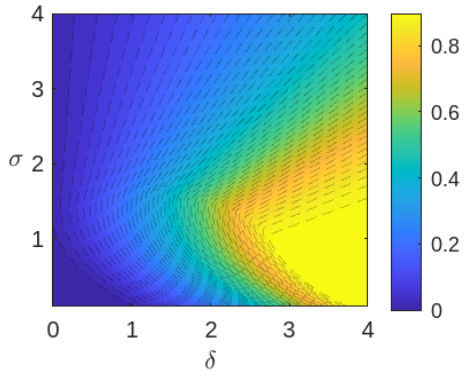

Fig. 5: Contour plot of $\overline{\gamma}(\delta,\sigma)$ against various values of δ and σ. Other parameters are $c = 1$, $\alpha = 0.5$, $R = 4$, $\mu_Z(x) = 1/x$, and $\sigma_Z(x) = 1.5/x$.

Fig. 5 shows some numerical simulations highlighting the above result. We make a similar observation that in terms of maximizing the mechanism's coverage or participation, it is once again best to have a highly accurate audit but not a very strict/difficult one.

V. ACKNOWLEDGMENTS

This work has been partially funded by Project no. 138903, implemented with the support provided by the Ministry of Innovation and Technology from the National Research, Development, and Innovation Fund, financed under the FK_21 funding scheme.

REFERENCES

[1] R. J. Anderson, "Why information security is hard-an economic perspective," in *17th Annual Computer Security Applications Conference (ACSAC 2001), 11-14 December 2001, New Orleans, Louisiana, USA*. IEEE Computer Society, 2001, pp. 358–365. [Online]. Available: https://doi.org/10.1109/ACSAC.2001.991552

[2] The White House, Washington, "National cybersecurity strategy," Government Document, 2023, https://www.whitehouse.gov/wp-content/uploads/2023/03/National-Cybersecurity-Strategy-2023.pdf.

[3] J. C. Fellingham and D. P. Newman, "Strategic considerations in auditing," *The Accounting Review*, vol. 60, no. 4, pp. 634–650, 1985. [Online]. Available: http://www.jstor.org/stable/247459

[4] E. M. Matsumura and R. R. Tucker, "Fraud detection: A theoretical foundation," *The Accounting Review*, vol. 67, no. 4, pp. 753–782, 1992. [Online]. Available: http://www.jstor.org/stable/248323

[5] D. R. FINLEY, "Game theoretic analysis of discovery sampling for internal fraud control auditing*," *Contemporary Accounting Research*, vol. 11, no. 1, pp. 91–114, 1994. [Online]. Available: https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1911-3846.1994.tb00438.x

[6] F. Ben abdelaziz, S. Neifar, and M. de Bourmont, "Auditing and game theory: A survey," *Multiple Criteria Decision Making in Finance, Insurance and Investment*, pp. 249–272, 2015.

[7] E. Patterson and J. Noel, "Audit strategies and multiple fraud opportunities of misreporting and defalcation*," vol. 20, no. 3, pp. 519–549, _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1506/F9KW-QM6U-6NXF-QCUN. [Online]. Available: https://onlinelibrary.wiley.com/doi/abs/10.1506/F9KW-QM6U-6NXF-QCUN

[8] A. B. Brown, "Incentives for auditor collusion in pre-sarbanes-oxley regulatory environment." [Online]. Available: https://papers.ssrn.com/abstract=976169

[9] M. M. Khalili, P. Naghizadeh, and M. Liu, "Designing Cyber Insurance Policies: The Role of Pre-Screening and Security Interdependence," *IEEE Transactions on Information Forensics and Security*, vol. 13, no. 9, pp. 2226–2239, Sep. 2018, conference Name: IEEE Transactions on Information Forensics and Security.

[10] M. M. Khalili, M. Liu, and S. Romanosky, "Embracing and controlling risk dependency in cyber-insurance policy underwriting," *Journal of Cybersecurity*, vol. 5, no. 1, p. tyz010, Jan. 2019. [Online]. Available: https://doi.org/10.1093/cybsec/tyz010

[11] T. P. Nugrahanti, "Analyzing the evolution of auditing and financial insurance: Tracking developments, identifying research frontiers, and charting the future of accountability and risk management," *West Science Accounting and Finance*, vol. 1, no. 02, p. 59–68, Jul. 2023. [Online]. Available: https://wsj.westsciencepress.com/index.php/wsaf/article/view/119

[12] O. Hernández-Lerma, *Adaptive Markov Control Processes*, ser. Applied Mathematical Sciences, F. John, J. E. Marsden, and L. Sirovich, Eds. New York, NY: Springer, 1989, vol. 79. [Online]. Available: http://link.springer.com/10.1007/978-1-4419-8714-3

[13] Z. Huang, G. Biczók, and M. Liu, "Incentivizing secure software development: The role of liability (waiver) and audit." [Online]. Available: http://arxiv.org/abs/2401.08476

[14] P. Naghizadeh and M. Liu, "Exit equilibrium: Towards understanding voluntary participation in security games," in *IEEE INFOCOM 2016 - The 35th Annual IEEE International Conference on Computer Communications*, 2016, pp. 1–9.

[15] T. Furusawa and H. Konishi, "Contributing or free-riding? Voluntary participation in a public good economy: Contributing or free-riding?" *Theoretical Economics*, vol. 6, no. 2, pp. 219–256, May 2011. [Online]. Available: http://doi.wiley.com/10.3982/TE567