Cyber-warranties as a quality signal for information security products

Daniel W. Woods and Andrew C. Simpson

Department of Computer Science, University of Oxford
Wolfson Building, Parks Road, Oxford OX1 3QD, UK
firstname.lastname@cs.ox.ac.uk

Abstract. Consumers struggle to distinguish between the quality of different enterprise security products. Evaluating performance is complicated by the stochastic nature of losses. It is recognised that this information asymmetry may lead to a "market for lemons" in which suppliers face no incentive to provide higher quality products. Some security vendors have begun to offer cyber-warranties — voluntary ex-ante obligations to indemnify the customer in the event of a cyber attack — to function as a quality signal. Much like how consumer protection laws are relatively more costly to firms offering low quality products, cyber-warranties are more costly for firms developing low quality enterprise security products. In this paper, we introduce a decision-theoretic model to explore how consumers might use cyber-warranties to increase information when purchasing security products. Our analysis derives four inferences that consumers can make about a security product. We discuss the difficulties customers might face in using these inferences to make real world decisions.

Keywords: cyber warranties, decision theory, enterprise security, quality signals, cyber insurance

1 Introduction

The "market for lemons" has been used to understand how information asymmetry can degrade the quality of traded goods [1]. Akerlof illustrated the concept by considering a used-car market dominated by sellers of "lemons" (low-quality cars) in which buyers are unable to distinguish between a "lemon" and a "peach" (a high-quality car). Recent work has used this analogy to explain the cyber crime market [18] and secure software markets [2].

It may be argued that enterprise security products exhibit qualities of a market for lemons. Security firms must decide whether to invest additional resources in developing a more effective product or, alternatively, sell the less-developed product (a "lemon"). If buyers are unable to distinguish between the two products, there is no incentive for the security firm to develop a more effective product; buyers will purchase the "lemon" regardless and the seller avoids incurring additional development costs. Enterprise security products lack a signal of quality that might address the information asymmetry.

Rao et al. [28] suggest that "brand name can convey unobservable quality credibly when false claims will result in intolerable economic losses". These losses can result from damage to the reputation of a brand, which is used by consumers to identify perceived product quality [12]. Alternatively, consumer protection laws place involuntary obligations on vendors that are more costly for the sellers of low quality products. Firms may be required to replace faulty products or may even be liable for the resulting damages in the case of strict product liability, which "induces firms to improve product safety" [27].

Such signals may be inappropriate in the context of enterprise security products. Evaluating the performance of security products is complicated. Preventing all attacks (giving rise to what might be termed 'absolute security') is widely held to be impossible [3]. A cyber attack may result from misfortune rather than from a faulty or low-quality product. Consequently, it is difficult to link product performance to product quality, not least because firms are reluctant to share detailed information about breaches [22]. This undermines both the function of reputation and the ability to identify faulty products to assign liability.

Enterprise security firms have begun to use so-called *cyber-warranties* as an alternative signal of quality. For example, a managed security provider[1] has offered a $100,000 warranty and an end-point protection firm[2] offers a $1,000,000 warranty. In this paper, we consider cyber-warranties to be voluntary ex-ante obligations in which enterprise security providers promise to indemnify consumers in the event of a successful attack. The voluntary ex-ante aspect of cyber-warranties differentiates them from the concept of software liability found in tort law [15, 30, 33] (or even criminal law [31]). By accepting and publicising these obligations, security firms seek to unilaterally shape market dynamics. There are many questions regarding what consumers can infer from these signals, how cyber-warranties impact the investment in security products, and whether this reduces the expected losses for the consumer.

This paper presents an economic consideration of how cyber-warranties affect the market for enterprise information security products. Section 2 identifies related work. In Section 3, we introduce a decision-theoretic model that captures both the vendor's short-run decision of setting the warranty level while investment is fixed and the long-run decision in which investment can vary. Section 4 contains our main contribution: the derivation of four inferences the consumer can make based on the cyber-warranty level. Section 5 illustrates how these inferences depend on the information structure between the consumer and the vendor. We discuss how applicable these inferences are with regards to real world decisions in Section 6. Section 7 offers conclusions and some directions for future work.

[1] `https://www.armor.com/cyber-warranty/`

[2] `https://www.sentinelone.com/press/sentinelone-establishes-1-million-cyber-threat-protection-guarantee/`

2 Background and Motivation

Cyber-warranties blur the line between risk mitigation and risk transfer. The vendor is tasked with both setting the optimal investment in product development and transferring the optimal amount of risk from the consumer in the form of a warranty. However, there has not been an academic consideration of cyber-warranties. Consequently, in this section we highlight how the literature on information security investments and risk transfer is relevant to our model (which we introduce in Section 3).

We focus on research into cyber insurance because it is concerned with a similar phenomenon: a cyber-warranty is a promise of indemnification much like an insurance contract. Cyber-warranties may lead to greater investment in the development of security products in the same way that insurance provides incentives for organisations to better manage information security [32]. For example, an insurer considering whether to directly invest in software security [24] faces a similar incentive structure to vendors offering cyber-warranties. Further, vendors may purchase market insurance to cover the liability for cyber-warrenties, which relates to research into cyber insurance for third party providers [21].

Böhme and Schwartz [10] introduce a framework to describe how different cyber insurance models approach this problem. A common approach [14, 19, 26] characterises the risk to the consumer by: a fixed loss l_i; insurance coverage $\beta_i \in [0, 1]$ that indemnifies a fraction of the loss; and a defence function D_i representing the probability of suffering a loss. Our model broadly adopts this framework to describe cyber-warranties, although it diverges on some specifics. We opt for simplicity, rather than trying to incorporate considerations such as secondary losses [5] made in other models.

The defence function D_i links the probability of suffering a loss to the "security investment s_i" [10]. In game-theoretic approaches, D_i has been assumed to have linear returns on investment in [14] and diminishing marginal returns on investment in [26], while Johnson et al. [19] assume that the other players' defensive investments influence D_i [19]. The seminal decision-theoretic work of Gordon and Loeb [16] introduces two probability breach functions analogous to D_i, which were corroborated using data on e-local governments in Japan [35].

Although representing the warranty level as a fraction of a set loss has precedent in the insurance literature [10, 14, 19, 26], doing so abstracts away from the myriad challenges of transferring so-called cyber risks. Empirical work reveals a more legalistic reality in which coverage is delimited into first and third party losses, with losses related to reputation damage and intellectual property loss not covered [29]. Policymakers have suggested that standardised policy wordings may help consumers understand what exactly they are purchasing [37]. These results from cyber insurance suggest there are significant real world problems in defining what a warranty covers.

Risk transfer leads to principal–agent problems such as adverse selection and moral hazard [4]. The first, adverse selection, occurs when riskier consumers purchase insurance at a greater frequency in the knowledge they are more likely to make a claim. Insurers attempt to better understand an applicant's risk by

Symbol	Description
V_i	The i-th vendor
S_i	The product offered by the i-th vendor
c_{f_i}	The fixed costs incurred in offering product S_i
z_i	The amount of investment into security during development of S_i
P_i	The price of S_i
Ψ_i	The proportion of realised losses the i-th vendor will indemnify
λ	The set loss resulting from a successful attack
v_0	The consumer's vulnerability before employing a security product
$S(v_0, z_i)$	The probability of successful attack given an investment of z_i
$R_c(R_v)$	The revenue of the consumer (vendor)

Table 1. Descriptions of each parameter in the model.

collecting information about information security controls [29, 36]. Based on this decision, they may decide to refuse coverage or offer it at a higher price [13]. However, empirical work suggests that less than a third of cyber insurers price risk according to information security factors [29]. Vendors might reflect on how to prevent warranties being purchased by the riskiest consumers. The second, moral hazard, occurs when an insured engages in risky behaviour in the knowledge the insurer will cover the losses. Traditionally, insurers address this problem by offering partial coverage so that the insured also suffers some financial consequences resulting from losses [38]. Another method involves exclusions, whereby insurers are no longer liable if certain procedures are not followed. For example, Kesan et al. [20] identify that a "failure to take reasonable steps to maintain and upgrade security" invalidates each of the policies in their study. Here we push up against the problem of defining the warranty as detailed conditions regarding risky behaviour increase contractual complexity.

Having identified a common modelling approach to cyber insurance and some of the associated real world principal–agent problems, we introduce our model for cyber-warranties in the next section.

3 Model

The model considers a number of vendors $V_1, .., V_n$, with each V_i selling a single security product S_i. Each vendor sets the amount of development investment z_i that represents costs, including developer time, training costs, participation in threat intelligence schemes and purchasing development tools.

We assume a Bertrand model of competition [6] in which a vendor can chooses a price P_i and a warranty $\Psi_i \in [0, 1]$; how this choice interacts with market demand determines the quantity supplied. The Bertrand model is relevant to software markets where quantity supplied can dynamically meet market demand [34] — unlike, for example, car manufacturers who must forecast market demand in order to begin a production process that may take months to complete.

To model the random nature of cyber attacks, we consider a Bernoulli trial in which the consumer faces a set loss λ with probability of occurrence p_i when the

consumer purchases product S_i and a probability of v_0 if no purchase is made. This realisation of losses is in line with the common approach to modelling other forms of risk transfer [10, 14, 19, 26]. As the set loss is fixed, the security products mitigate the probability of successful attack without affecting the impact of the attack. Consequently, our analysis will be less relevant to security products that seek to reduce the impact of losses.

As identified in Section 2, there are many functions relating p_i to the investment z_i in the security product S_i. Gordon and Loeb's seminal paper [16] established three core assumptions that such a function should fulfill in the context of protecting an information set. These are listed below.

A1: $S(z_i, 0) = 0$ for all $z_i \in \mathbb{R}$
A2: $S(0, v_0) = v_0$ for all $v_0 \in [0, 1]$
A3: $\frac{\delta S}{\delta z}(z_i, v_0) < 0$ and $\frac{\delta^2 S}{\delta z^2}(z_i, v_0) > 0$ for all $v_0 \in [0, 1]$ and $z_i \in \mathbb{R}$. Furthermore, for all $v_0 \in [0, 1]$ we have,

$$\lim S(z_i, v_0) \to 0 \text{ as } z_i \to \infty$$

The third assumption ensures that further investment reduces the probability of attack, but does so at a diminishing rate. Further, no finite investment results in perfect security.

In [16], Gordon and Loeb propose two classes to which the security breach probability function may belong. These will be used going forward and may be expressed in the form

$$S^I(z_i, v_0) = \frac{v_0}{(\alpha z_i + 1)^\beta} \tag{1}$$

and

$$S^{II}(z_i, v_0) = v^{\alpha z_i + 1} \tag{2}$$

These assumptions and the corresponding functions were introduced in the context of protecting an information set. It can be argued that there is relevance to enterprise security products, particularly when the products out-source the task of protecting an information set.

The vendor incurs total cost $c_{f_i} + z_i$, where c_{f_i} represents the costs unrelated to security in offering the product, which we assume to be fixed. While z_i is the variable describing investment in security, which a vendor can may set. Each vendor seeks to maximise their profit Π_i by setting P_i, z_i and Ψ_i:

$$\Pi_i = P_i - S(z_i, v_0)(\lambda \cdot \Psi_i) - (c_{f_i} + z_i) \tag{3}$$

The consumer only has knowledge of the price P_i, warranty Ψ_i and set loss λ. The investment z_i is assumed to be unobservable due to information asymmetry. The consumer chooses the security product S_i that minimises

$$R_c = P_i + S(z_i, v_0) \cdot \lambda(1 - \Psi_i) \tag{4}$$

We assume that customers are homogeneous and all demand the same product, leading to the kind of winner-takes-all market dynamics that have been observed in many other software markets [2, 34].

4 Analysis

We consider a market without security warranties ($\Psi_i = 0$) to illustrate the market for lemons. Using Equation 3, the vendor receives

$$P_i - (c_{f_i} + z_i)$$

while the consumer's expected security expenditure is

$$P_i + S(z_i, v_0) \cdot \lambda$$

The vendor has no incentive to increase the development investment beyond $z_i = 0$ because the consumer cannot observe ex-ante the resulting decrease in vulnerability. In a competitive market without warranties, the market equilibrium is $P_i = c_{f_i}$ with $z_i = 0$. Clearly vendors still invest in product development without offering warranties in spite of this result and we discuss why they might do so in Section 6. The rest of this section identifies four inferences consumers can make regarding security products, as well as the information they need to do so.

First, we consider a vendor V_i with a fixed investment of $c_{f_i} + z_i'$ in the product. Each vendor can offer the product at a price P_i with warranty Ψ_i. Equation 3 shows that the vendor's profit at the price P_i is as follows.

$$\Pi_i(P_i) = P_i - S(z_i', v_0)(\lambda \cdot \Psi_i) - (c_{f_i} + z_i') \tag{5}$$

In the short-run, the vendor may incur losses up to the value of the fixed costs of operation $(c_{f_i} + z_i')$. This observation leads to the constraint

$$\Pi_i(P_i) \geq -(c_{f_i} + z_i')$$

from which we derive Inference 1.

Inference 1 *Vendor V_i can offer S_i in the short-run with a warranty level of $\Psi_i \in [0, 1]$ at any price*

$$P_i \geq S(z_i', v_0)\lambda \cdot \Psi_i$$

The left-hand side represents the expected value of the indemnification payment to the consumer. It is reasonable to assume that no risk-neutral vendor would offer the warranty unless they receive at least this value as an up-front payment. This provides an upper bound of $\frac{P_i}{\Psi_i}$ for the expected loss a consumer faces — dividing by Ψ_i adjusts for the proportion of the loss that the vendor pays. The inference can be made in the presence of information asymmetry regarding the vendor's security efficiency (α, β), the shape of the probability breach function $S(\cdot, \cdot)$, or their security investment during development z_i.

The consumer seeks to minimise Equation 4 despite having incomplete information about z_i. The consumer can use Inference 1 to calculate a lower bound z_{min_i}, which represents the smallest investment value such that vendor i can break even in offering offer a product with warranty Ψ_i at price P_i. This value

may be used to calculate the worst-case expected loss $R_{c_{min}}$ resulting from purchasing the product S_i:

$$R_{c_{min}}(S_i) = P_i + S(z_{min_i}, v_0) \cdot \lambda(1 - \Psi_i) \tag{6}$$

The consumer is assumed to be indifferent between purchasing the product S_i and the product S_j if

$$R_{c_{min}}(S_i) = R_{c_{min}}(S_j) \tag{7}$$

From this, we can construct a (worst-case) indifference curve for the consumer.

Calculating the (worst-case) indifference curve involves finding the smallest z_i such that

$$\Pi(S_i) \geq -(c_{f_i} + z_i') \tag{8}$$

Using Equation 3 and the formulae for each class of probability breach function, we derive Inference 2.

Inference 2 *If the product S_i has been offered at price P_i and the warranty level is Ψ_i we have that*

$$z_{min_i} = \begin{cases} \frac{(\frac{\lambda\Psi' v_0}{P_i})^{\frac{1}{\beta}} - 1}{\alpha} & \text{if } S(\cdot,\cdot) \text{ is Class I} \\ \frac{ln(P_i) - ln(\Psi\lambda v_0)}{\alpha ln(v_0)} & \text{if } S(\cdot,\cdot) \text{ is Class II} \end{cases}$$

It is worth noting that Inference 2 may provide an under-estimate of the product investment. A profit-making vendor analysed as if the vendor was breaking even would appear to have invested less than they did in actuality.

The long-run decision reduces to first selecting a warranty level and then determining the optimal investment as setting the investment first would reduce to the short-run analysis.

Suppose that the vendor unilaterally sets the warranty level at $\Psi_i' > 0$. The vendor will make the long-run investment z_i^* that optimises profit $\Pi_i(P_i)$ for all values of P_i. The marginal net benefit of investment is given by

$$\frac{\partial \Pi_i}{\partial z_i} = -\frac{\delta S}{\delta z}(z_i, v_0)(\lambda \cdot \Psi') - 1 \tag{9}$$

Using the convention that investment is non-negative ($0 \leq z_i$), we can derive the following.

Inference 3 *If the vendor has committed to the warranty level Ψ_i', the optimal choice of product investment is*

$$z_i^* = \begin{cases} \frac{(\alpha\beta\lambda\Psi' v_0)^{\frac{1}{\beta+1}} - 1}{\alpha} & \text{if } S(\cdot,\cdot) \text{ is Class I and } \alpha\beta\lambda\Psi' v_0 > 1 \\ \frac{-ln(-\alpha\lambda\Psi' v ln(v))}{\alpha ln(v)} & \text{if } S(\cdot,\cdot) \text{ is Class II and } \alpha\lambda\Psi' v ln(v) > -1 \\ 0 & \text{otherwise} \end{cases}$$

Inference 3 allows the consumer to infer the exact level of investment providing the warranty level was decided in the long-run and investment was optimised for this decision. If the investment z_1^* can be inferred, the consumer can expect revenue $R_c(S_i)$ if they purchase S_i, where

$$R_c(S_i) = P_1 + S(z_i^*, v_0)(1 - \Psi_i)\lambda \tag{10}$$

In a fully competitive market, we can expect that

$$P_i = S(z_i^*, v_0)\Psi'\lambda + c_{f_i} + z_i^* \tag{11}$$

However, Inference 2 and Inference 3 both rely on the consumer knowing the shape of the probability breach function and the vendor's security productivity.

In both the short-run and the long-run, the price P_i must increase to compensate for any increase in the warranty Ψ_i at a rate equal to the risk-transfer rate of substitution (RTRS) $\frac{\partial \Pi_i}{\partial \Psi_i}$ in order to keep profits constant.

Inference 4 *The risk-transfer rate of substitution for the vendor V_i is equal to the consumer's expected loss when the security product S_i is in place.*

$$\frac{\partial \Pi_i}{\partial \Psi_i} = S(z_i', v_0)\lambda$$

The consumer can discover the expected loss if the risk-transfer rate of substitution is observed. This inference can be made with knowledge of only the price and warranty level regardless of whether the warranty has been offered in the short-run or the long-run. This inference might be considered the most powerful as it can be made with information asymmetry regarding the vendor's technological constraints.

The price and warranty offered by each vendor will depend on the market environment in both the short-run and the long-run. If the vendors have perfect information about the competitors' investments in product development, there may exist one vendor who can extract a supplier surplus by setting (P_i, Ψ_i) such that any competitor would suffer an economic loss in offering a competing product. However, this will depend on the particular values of both investments z_i, existing vulnerability v_0 and breach probability function $S(z_i, v_0)$. The relative risk aversion of the vendor and the consumer will determine the optimal pair (P_i, Ψ_i).

5 Numerical Illustration

In this section we illustrate each of the inferences in turn.

Firms will only shut down in the short-run if price exceeds average cost. Figure 1 shows how the minimum price is determined by the warranty level and the investment. We define the shutdown-isoprofits to be the lines with a loss equal to fixed costs; the curve for $z = z_j'$ represents the possible pairs (P_j, Ψ_j) for which the j-th vendor's profit $(\Pi_j(S_j))$ is equal to $c_{f_j} + z_j'$. Although vendors

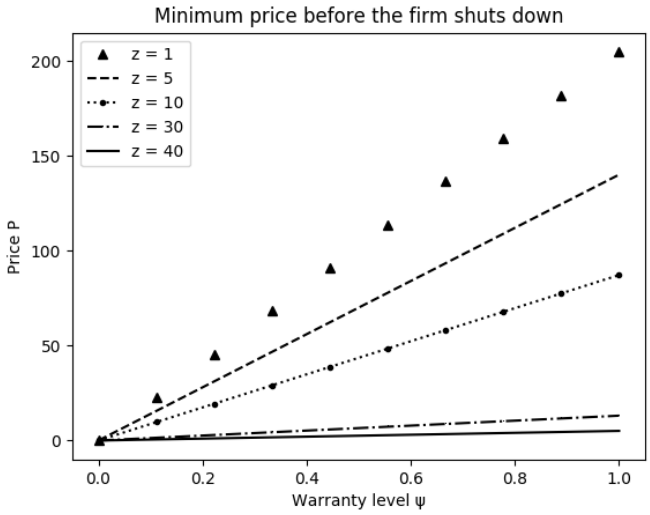

Fig. 1. The price at which the vendor would shut down if price fell any further, for different investment levels z and a Class I probability breach function with: $\alpha = 0.9, \beta = 1, \lambda = 500, v_0 = 0.5$ and $c_f = 5$.

V_1 and V_5 have invested $z_1 = 1$ and $z_5 = 40$ respectively, both accept a minimum price of 0 when no warranty is offered. The difference between the size of their losses will be given by

$$\Pi_5(0) - \Pi_1(0) = (c_{f_5} + z_5) - (c_{f_1} + z_1) = -39$$

because V_5 has larger fixed costs as a result of higher fixed investment z_5.

Figure 2 illustrates the isoprofits when the firms break even. For a given warranty level Ψ, the vendor with the isoprofit curve intersecting $x = \Psi$ at the lowest point can offer the most competitive product. This provides a graphical illustration of the market for lemons as the product with investment $z = 1$ is most competitive when no warranty is offered. The downside of over-investment can be seen by considering that the vendor who invested $z = 30$ is more competitive at every warranty level than the vendor who invested $z = 40$.

If the consumer had knowledge about the shape of the probability breach function and the vendor's security efficiency, Inference 2 can provide information about the vendor's minimum investment z_{min_i} in the short-run. Figure 3 highlights the points at which the strongest inference can be made; more information is contained in a warranty as the price it is offered at decreases. Consider a duopoly with vendors V_1 and V_2 who have made investments of $z_1 = 5$ and $z_2 = 10$ respectively. If vendor V_2 sets (P_2, Ψ_2) to be equal to any pair of Figure 3 with $z_m in > 5$, then V_1 would sooner shut-down operation than offer the same contract. Offering such a contract functions as a reliable signal of product quality in this scenario.

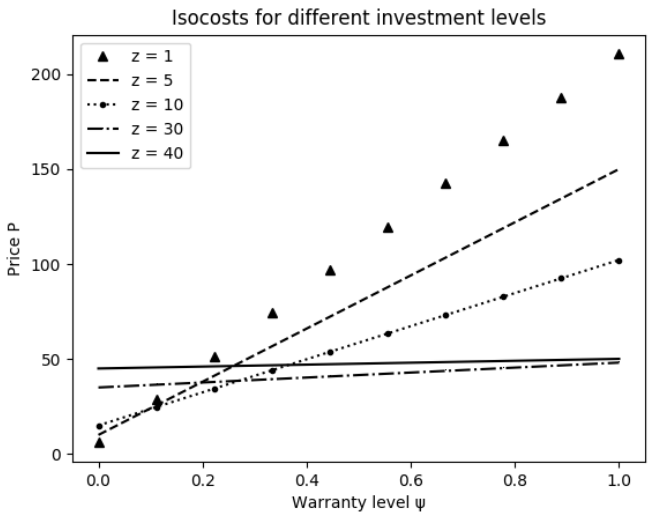

Fig. 2. The price at which the vendor would shut down if price fell any further, for different levels of security investment z and a Class I probability breach function with: $\alpha = 0.9, \beta = 1, \lambda = 500, v_0 = 0.5$ and $c_f = 5$.

For each contract (P_i, Ψ_i), Inference 2 may also be understood graphically as the smallest value of z_i for which the associated isoprofit curve intersects (P_i, Ψ_i) or falls beneath it. For z_{min_i} to be the worst case, we have to assume that the isoprofit corresponded to the points where the loss is equal to the fixed costs. Inference 2 might lead to different conclusions if we used the isoprofit curves corresponding to a different profit condition, such as breaking even as in Figure 2. If a functional form is difficult to obtain for a given profit condition, the graphical interpretation of Inference 2 may be used instead.

Turning to long-run investments, Figure 4 shows the price a vendor must charge for a given warranty level in order to break even. We have circled the optimal investment for each warranty level and can see that it is increasing in Ψ_i. Consumers may use Inference 3 to discover the optimal investment level z^* and use it to calculate their expected loss. When investment costs are fixed, the consumer can only infer a lower bound for investment whereas the consumer can now infer the optimal investment level for a given warranty level.

Figure 5 describes the expected loss (R_c) for each customer if they purchase a product with warranty Ψ_i assuming the optimal investment has occurred. As the curve is always downward-sloping we must have

$$\frac{\partial R_c}{\partial \Psi_i} < 0 \text{ for all } \Psi_i \in [0, 1]$$

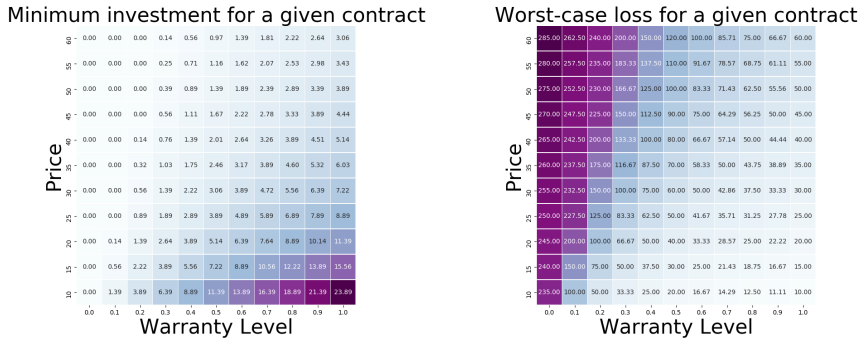

Fig. 3. The minimum investment value z_{min_i} and worst-case loss $R_{c_{min}}$ for a given price P_i and warranty level Ψ_i, for a Class I probability breach function with: $\alpha = 0.9, \beta = 1, \lambda = 500, v_0 = 0.9$ and $c_f = 5$.

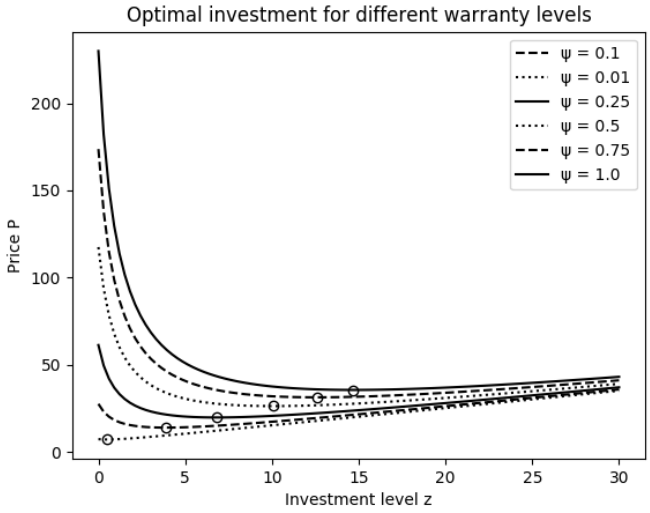

Fig. 4. The choices of price and investment level that lead a vendor to make zero profit, for different investment levels Ψ, for a Class I probability breach function with: $\alpha = 0.9, \beta = 1, \lambda = 500, v_0 = 0.9$ and $c_f = 5$.

However, the customer's expected loss falls at a diminishing rate so that

$$\frac{\partial^2 R_c}{\partial \Psi_i^2} > 0 \text{ for all } \Psi_i \in [0,1]$$

These results, derived via Inference 3, suggest that greater risk transfer to the agent deciding amount of security investment leads to a more efficient allocation of resources. As such, consumers should push to increase the warranty level over time, which can be seen in the decreasing expected loss for greater warranty levels in Figure 5. Inference 3 requires knowledge about the vendor's technological constraints, much like Inference 2.

Inference 4 states that the risk-transfer rate of substitution (RTRS) of the i-th vendor is equal to the expected loss when employing the i-th security product. The RTRS for a given vendor is equal to the slope of that vendor's isocost curve (in Figure 2 and Figure 1). The lowest investment has the steepest isoprofit curve and hence the highest expected loss. Negotiating with the vendor might reveal the RTRS if the vendor stated how much the price would have to rise for a given increase in warranty level.

In summary, Inference 1 provides a lower-bound on expected losses and Inference 2 provides a lower bound on product investment. Both of these are valid in the short-run. Inference 3 provides an exact value of the optimal investment for a given warranty but it is only valid in the long-run. However, Inference 2 and Inference 3 require knowledge about the vendor's technological constraints. Inference 4 provides an exact value of the expected loss. It is valid in both the long-run and the short-run, and requires no knowledge beyond the RTRS. The next section discusses some of the real world limitations of these inferences.

6 Discussion

The consumer can use inferences 1–4 to estimate the expected loss when implementing the security product S_i, which can be compared against the expected loss without any security product or some other security product S_j. This allows the consumer to estimate the expected benefit from the security product. More knowledge about the vendor or the risk transfer rate of substitution may allow the consumer to make stronger inferences and increase confidence in these estimates. Unfortunately, these estimates will be weakened by many complicating factors in the real world.

The warranty level will likely take the form of a contract that will not stipulate a proportion of risk the vendor will cover. The contract might instead define a selection of events for which the warranty is valid. Estimating the proportion of the expected loss that these events represent requires that risk managers understand their organisation's risk profile.

Although the model suggests that full risk transfer achieves the optimal solution for the consumer, it may not be possible in the real world. Cyber insurance policies do not cover intangible losses such as reputation damage and intellectual

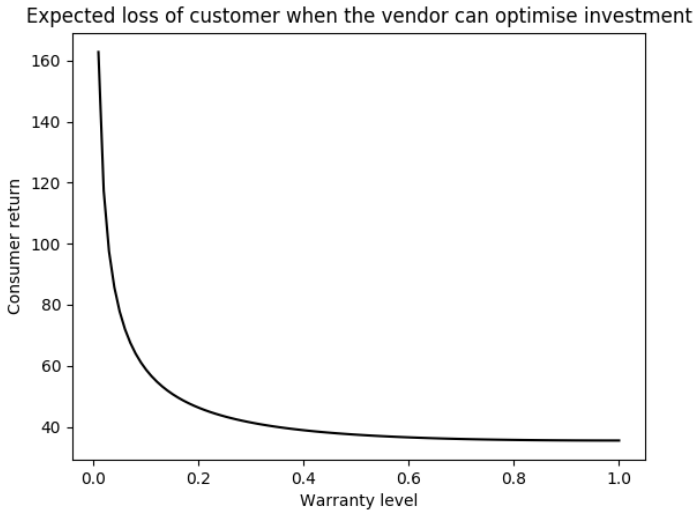

Fig. 5. The choices of price and investment level that lead a vendor to make zero profit, for different investment levels Ψ, for a Class I probability breach function with: $\alpha = 0.9, \beta = 1, \lambda = 500, v_0 = 0.9$ and $c_f = 5$.

property loss precisely because it is difficult to quantify such losses. There is no reason why vendors are better suited to offer warranties covering these risks.

A further complicating factor is that prices must reflect principal–agent problems such as adverse selection and moral hazard. These problems have presented a major problem for cyber insurance, as we observed in Section 2. Solutions to these problems, such as monitoring the consumer's security practices to prevent moral hazard or performing an in-depth assessment to prevent adverse selection, come at a cost that may be reflected in the price. However, consumers may accept a higher price because they need to invest less resources in evaluating product quality; cyber warranties incentivise the vendor to invest in the product regardless of whether the consumer can observe these investments.

The risk-transfer rate of substitution is only equal to the expected loss if the vendor is risk neutral. Otherwise the consumer would have to correct for the vendor's discomfort with holding greater liability associated with a higher warranty level. Vendors should also be concerned by the possibility of a "cyber hurricane" in which interdependent events trigger multiple indemnification claims [7, 8].

Furthermore, the insolvency risk to vendors grows as they hold more liability. The risk may be managed via self-insurance or market insurance to ensure that vendors have funds available for indemnification. An equilibrium between the cost of these risk management techniques and consumer demand will determine the warranty level available to the consumer at a given price.

7 Conclusion and Future Work

Customers face information asymmetry when deciding which information security product to purchase. Cyber-warranties can overcome this information asymmetry by creating a separating equilibrium in which the vendor reveals the level of product investment to the consumer. Vendors selling information security products face lower costs in offering cyber-warranties if they invest in developing more effective products.

Our model identifies four inferences that customers can make about a potential information security purchase based on the warranty offered. In general, more information is gained when there is more prior knowledge about the vendor. However, these inferences are likely to be weaker in the real world. Consumers must adjust for factors including the extent of the vendor's risk aversion, costs incurred to mitigate principal–agent problems, and risk-loading to deal with the variability of (potentially correlated) losses.

Future work could explore how the balance of risk aversion between vendors and consumers affects the supply and demand for cyber-warranties. Another factor to consider is the vendor's costs in terms of mitigating (via self insurance or market insurance) the insolvency risk when increasing the warranty level. Identifying an empirical basis for the parameter choices may increase relevance for practitioners. Further, future work could reflect some active research topics in information security investments including:

- the benefits from adaptive security in which a defender makes defensive investments in response to observed losses [9];
- the role of investments in recovery as opposed to just mitigation, particularly in light of recent developments in cyber crime [23];
- the interaction with the vulnerability disclosure process, particularly the role of government policy [11];
- the relevance of the stage in the development process at which investments are made [17]; and
- approaches applying game theory to consider strategic interactions between vendors and consumers [25] (for example, vendors and consumers might dishonestly avoid indemnification or fraudulently claim indemnification respectively).

8 Acknowledgements

The authors thank the anonymous reviewers for their helpful and constructive comments. Participants in the "Effect of Software Warranties on Cyber Security" workshop run by the University of Bristol's Cyber Security Group provided useful feedback for the ideas developed in this paper. Daniel Woods' research is funded by the EPSRC via the Centre for Doctoral Training in Cyber Security at the University of Oxford.

References

1. Akerlof, G.A.: The market for lemons: Quality uncertainty and the market mechanism. In: Diamond, P., Rothschild, A. (eds.) Uncertainty in Economics, pp. 235–251. Elsevier (1978)
2. Anderson, R., Moore, T.: The economics of information security. Science 314(5799), 610–613 (2006)
3. Anderson, R.J.: Security engineering: a guide to building dependable distributed systems. John Wiley & Sons (2010)
4. Arrow, K.J.: Uncertainty and the welfare economics of medical care (American Economic Review, 1963). Journal of Health Politics, Policy and Law 26(5), 851–883 (2001)
5. Bandyopadhyay, T., Mookerjee, V.S., Rao, R.C.: Why IT managers don't go for cyber-insurance products. Communications of the ACM 52(11), 68–73 (2009)
6. Bertrand, J.: Theorie mathematique de la richesse sociale. Journal des Savants pp. 499–508 (1883)
7. Biener, C., Eling, M., Wirfs, J.H.: Insurability of cyber risk: An empirical analysis. The Geneva Papers on Risk and Insurance-Issues and Practice 40(1), 131–158 (2015)
8. Böhme, R.: Cyber-insurance revisited. In: Proceedings of The 4th Workshop on the Economics of Information Security (WEIS 2005) (2005)
9. Böhme, R., Moore, T.: The iterated weakest link model of adaptive security investment. Journal of Information Security 7(2), 81–102 (2016)
10. Böhme, R., Schwartz, G.: Modeling cyber-insurance: Towards a unifying framework. In: Proceedings of The 9th Workshop on the Economics of Information Security (WEIS 2010) (2010)
11. Caulfield, T., Ioannidis, C., Pym, D.: The US vulnerabilities equities process: An economic perspective. In: Proceedings of the 8th Conference on Decision and Game Theory for Security. pp. 131–150. Springer (2017)
12. Dodds, W.B., Monroe, K.B., Grewal, D.: Effects of price, brand, and store information on buyers' product evaluations. Journal of marketing research 28(3), 307–319 (1991)
13. Franke, U.: The cyber insurance market in Sweden. Computers & Security 68, 130–144 (2017)
14. Fultz, N., Grossklags, J.: Blue versus red: Towards a model of distributed security attacks. In: Dingledine, R., Golle, P. (eds.) International Conference on Financial Cryptography and Data Security. pp. 167–183. Springer (2009)
15. Gemignani, M.C.: Product liability and software. Rutgers Computer & Technology Law Journal 8, 173 (1980)
16. Gordon, L.A., Loeb, M.P.: The economics of information security investment. ACM Transactions on Information and System Security (TISSEC) 5(4), 438–457 (2002)
17. Heitzenrater, C., Simpson, A.C.: A case for the economics of secure software development. In: Proceedings of the 2016 New Security Paradigms Workshop. pp. 92–105. ACM (2016)
18. Herley, C., Florêncio, D.: Nobody sells gold for the price of silver: Dishonesty, uncertainty and the underground economy. In: Moore, T. (ed.) Economics of Information Security and Privacy, pp. 33–53. Springer (2010)
19. Johnson, B., Böhme, R., Grossklags, J.: Security games with market insurance. In: Proceedings of the 2nd Conference on Decision and Game Theory for Security. pp. 117–130. Springer (2011)

20. Kesan, J., Majuca, R., Yurcik, W.: Cyberinsurance as a market-based solution to the problem of cybersecurity: a case study. In: Proceedings of The 4th Workshop on the Economics of Information Security (WEI 2005) (2005)
21. Khalili, M.M., Liu, M., Romanosky, S.: Embracing and controlling risk dependency in cyber-insurance policy underwriting. In: Proceedings of The 17th Workshop on the Economics of Information Security (WEIS 2018) (2018)
22. Kotulic, A.G., Clark, J.G.: Why there arent more information security research studies. Information & Management 41(5), 597–607 (2004)
23. Laszka, A., Farhang, S., Grossklags, J.: On the economics of ransomware. In: Proceedings of the 8th Conference on Decision and Game Theory for Security. pp. 397–417. Springer (2017)
24. Laszka, A., Grossklags, J.: Should cyber-insurance providers invest in software security? In: European Symposium on Research in Computer Security. pp. 483–502. Springer (2015)
25. Manshaei, M.H., Zhu, Q., Alpcan, T., Bacşar, T., Hubaux, J.P.: Game theory meets network security and privacy. ACM Computing Surveys (CSUR) 45(3), 25 (2013)
26. Pal, R., Golubchik, L.: Analyzing self-defense investments in internet security under cyber-insurance coverage. In: Proceedings of the IEEE 30th International Conference on Distributed Computing Systems (ICDCS2010). pp. 339–347. IEEE (2010)
27. Polinsky, A.M., Shavell, S.: The uneasy case for product liability. Harvard Law Review 123, 1437–1491 (2009)
28. Rao, A.R., Qu, L., Ruekert, R.W.: Signaling unobservable product quality through a brand ally. Journal of Marketing Research 36(2), 258–268 (1999)
29. Romanosky, S., Ablon, L., Kuehn, A., Jones, T.: Content analysis of cyber insurance policies: How do carriers write policies and price cyber risk? In: Proceedings of The 16th Workshop on the Economics of Information Security (WEIS 2017) (2017)
30. Rustad, M.L., Koenig, T.H.: The tort of negligent enablement of cybercrime. Berkeley Tech Law Journal 20, 1553 (2005)
31. Ryan, D.J., Heckman, C.: Two views on security software liability. let the legal system decide. IEEE Security & Privacy 99(1), 70–72 (2003)
32. Schneier, B.: Insurance and the computer industry. Communications of the ACM 44(3), 114–114 (2001)
33. Scott, M.D.: Tort liability for vendors of insecure software: Has the time finally come. Maryland Law Review 67, 425 (2007)
34. Shapiro, C., Varian, H.R.: Information rules: a strategic guide to the network economy. Harvard Business Press (1998)
35. Tanaka, H., Matsuura, K., Sudoh, O.: Vulnerability and information security investment: An empirical analysis of e-local government in japan. Journal of Accounting and Public Policy 24(1), 37–59 (2005)
36. Woods, D., Agrafiotis, I., Nurse, J.R., Creese, S.: Mapping the coverage of security controls in cyber insurance proposal forms. Journal of Internet Services and Applications 8(1), 8 (2017)
37. Woods, D., Simpson, A.C.: Policy measures and cyber insurance: A framework. Journal of Cyber Policy 2(2), 209–226 (2017)
38. Zweifel, P., Eisen, R.: Insurance Economics. Springer Science & Business Media (2012)