Dynamic Pricing for Ransomware

Tristan Caulfield[1], Christos Ioannidis[2], and David Pym[1,3]

[1] University College London
t.caulfield@ucl.ac.uk, d.pym@ucl.ac.uk
[2] Aston Business School
c.ioannidis@aston.ac.uk
[3] The Alan Turing Institute

Abstract Ransomware encrypts files on a target's computer, demanding a payment in return for decrypting the files. However, ransomware authors do not know the value of the data their malware has encrypted. We present a dynamic pricing approach using a logit model to learn the willingness to pay of a population of companies, allowing better pricing. We then explore how price perturbations, the reliability of ransomware, and the amount of information sharing between companies affects ransomware revenue. Finally, we draw insights from the model about ransomware prevention.

1 Introduction

Ransomware is a type of malware that denies access to files on an infected device until a ransom has been paid. Typically, the files are encrypted and payment of the ransom results in the return of the cryptographic key that allows their decryption. Ransomware has been growing as a threat; recent varieties, such as *WannaCry*, *Petya*, and *BadRabbit* show increasing sophistication and ability to infect devices. Many businesses choose to pay the ransom in order to recover their data [5], with average ransom demands of over $500 (in 2017) and over $1000 (in 2016) [10].

In this paper we model a dynamic pricing strategy for ransomware, where the ransomware authors — unaware of the value of the data they have encrypted or whether backups of the data exist — learn the willingness to pay of a population of companies through the companies' decisions to pay or not pay the ransom. Understanding how ransom demands can be priced can give insights into the effects of businesses' decisions to employ backups or to share information about payments on the profits of ransomware authors.

In Section 2, next, we discuss related work. Then, in Section 3 we present the structure of the model. Section 4 explores the results of simulating the model under a number of different parameters. Finally, Section 5 concludes the paper with insights from the model about ransomware prevention.

2 Related Work

There has been a limited, but growing, amount of research into economic aspects of ransomware. In [4], the authors discuss ransomware profitability and pricing, including how ransomware authors should adjust the ransom price to maximise profit. The paper also discusses how price discrimination and bargaining can affect ransomware profits, along with the determinants of victims' willingess to pay.

Other research has taken game theoretic approach: [6] looks at the interactions of the attacker and defender, and focuses particularly on the investment in mitigations (such as backup technologies). Some variants of ransomware allow the ransomware authors and victims to communicate, which makes it possible to bargain about the price of the ransom. This case, in which the victim can make a counter-offer is explored in [1], where the authors relate ransomware to kidnapping and build a game theoretic model.

Other work looks at the payments made to ransomware authors. Bitcoin transactions, which are commonly used for ransomware payments are publibly visible; [8] traces the payments made to known ransomware addresses using bitcoin, and finds that a minimum of nearly $12.8 million was paid.

In this paper, we consider dynamic pricing that uses a logit model to learn the population's willingess to pay, which has has been studied extensively. The basic model, introduced in work such as [9], looks at a company making pricing decisions as it sells to a sequence of customers; there is a trade-off between learning the demand curve and maximising revenue. Finding optimal policies is a very difficult problem, so much work explores the effectiveness of different heuristics for myopic policies. In [3], heuristics for learning between two possible demand models with a myopic Bayesian policy are developed. Similarly, [2] finds upper and lower bounds for the pricing problem, and finds near-optimal heuristics. In [7], dithering strategies are analysed, in which perturbations of myopically optimal price can lead to increased rates of learning.

3 The Model

There are many different types of ransomware that operate in many different ways. For example, some ransomware allows communication between the ransomware authors and the owner of the affected device, meaning that bargaining about the price is possible; other types increase the payment amount after some amount of time. Similarly, companies may employ many different defenses and countermeasures to prevent or reduce the cost of ransomware infection.

This model looks at a simplified picture of ransomware. We consider only ransomware variants where communication is not possible, and the price is fixed once announced to the target — there is no bargaining or negotiation. We also restrict the model to consider just one type of countermeasure against ransomware: the use of backups, which allow the hostage data to be restored without paying the ransom, for a small cost (representative of time or effort).

3.1 Companies

We model a population of N companies. Each company has a value for the data that gets encrypted by the ransomware. Some proportion of companies use backup technology, which can restore the lost data for a small cost.

The reward, $r_{p,i}$, for a company i that pays the ransom is defined as

$$r_{p,i} = d_i \, \mathbb{E}[\rho]$$

where $d_i \sim N(1000, 150)$ is the value of the data for the company, and $\mathbb{E}[\rho]$ is the expected reliability of the ransomware — essentially, how likely paying the ransom will result in the return of data.

The reward, $r_{n,i}$ for a company that does not pay the ransom is defined as

$$r_{n,i} = \begin{cases} 0 & \text{if } b_i = 0 \\ d_i - c_i & \text{if } b_i = 1 \end{cases}$$

where b_i indicates whether the company uses backup technology that can allow it to recover the data for a cost $c_i \sim N(150, 50)$.

The *willingness to pay* of a company is then given by

$$w_i = r_{p,i} - r_{n,i}$$

which describes how much the are willing to pay to a ransomware author in order to recover their data. Figure 1 shows how the willingness to pay values of the companies are distributed. It is bimodal, split between the companies using backups and those that are not.

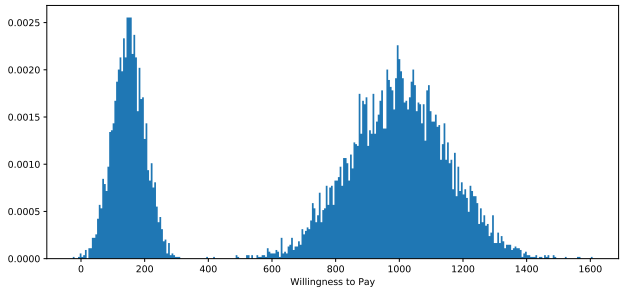

Figure 1: Distribution of the ransom values companies are willing to pay, when 30% of the population is using backups.

3.2 Ransomware Author

We assume that the ransomware author does not know the value of the data that the ransomware has encrypted on the company's machine, nor do they know whether or not the company employs some sort of backup mechanism which would allow the company to recover its data relatively inexpensively. They should select a price that gives the greatest expected revenue given their belief about the population's willingness to pay, and can learn from the prices and results (whether or not the company paid the ransom) in order to improve their ability to predict the best price.

We model this as a sequence of price decisions and feedback. At each time period t, the ransomware infects a company i and the ransomware author decides a ransom price, x_t; the company then chooses to either pay the ransom, if the price is less than their willingness to pay, or not pay, if it is greater. The ransomware author then uses this binary feedback, s_t, to update their model of the population:

$$s_t = \begin{cases} 1 & \text{if } x_t < w_i \\ 0 & \text{if } x_t \geq w_i \end{cases}$$

and receives revenue v_t:

$$v_t = \begin{cases} x_t & \text{if } s_t = 1 \\ 0 & \text{if } s_t = 0 \end{cases}$$

The goal of the malware author is to maximise total revenue V:

$$V = \sum_{t=1}^{T} v_t$$

The population of companies is represented (from the ransomware author's perspective) using a logit model, which describes how the binary purchase decision of companies depends on the price. The model estimates the probability of a decision to pay the ransom for a given price. The parameters of this model are determined using a Bayesian logistic regression, which gets updated after each time period. We assume that the ransomware author has some knowledge of the willingness to pay distribution, and so we create a prior by generating a small number of companies and random price points, and fitting the logit model to this data.

Figures 2a and 2b show the probability of a successful ransom payment for each price, along with the expected revenue at each price, according to the logit model after a number of time period. In Figure 2a, the use of backups in the population is low, at 20%; accordingly, the maximum expected price is higher than that of Figure 2b, where a higher number (80%) of companies use backups.

We model the ransomware author's pricing decisions using a myopic policy: the price selected has the highest expected revenue given the current state of the logit model. After the model is updated, a new best price is selected for the next time period. Figure 3 shows the history of pricing over time as the model learns the willingnes to pay distribution of the population.

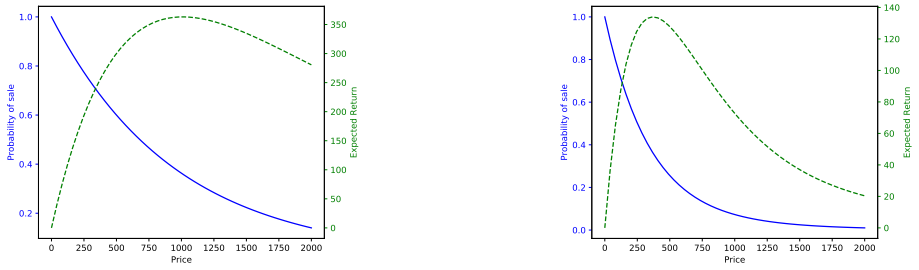

(a) Low use of backups.

(b) High use of backups.

Figure 2: Probability of sale and expected revenue for low (20%) and high (80%) use of backups, according to the logit model.

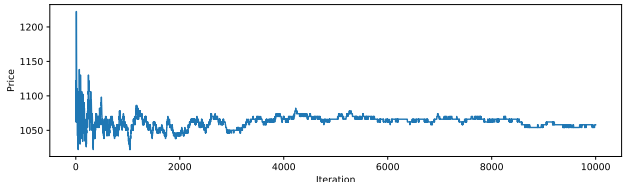

Figure 3: Price offered each iteration

4 Results

In this section we present the results from a number of simulations of the model using different parameters. We start with the basic model and then extend it to explore different price perturbation strategies, different reliabilities, and the impact of information sharing among the population of companies. For each simulation, we use a sequence of 2000 companies, with 1000 iterations for each parameter combination.

4.1 Dynamic Pricing

First, we consider the basic model. We fix $\mathbb{E}[\rho] = 1$, meaning that the ransomware is perfectly reliable — any company that pays the ransom gets their data back. The ransomware author picks the best price in each time period, according to the logit model, and updates the model after learning whether the company does or does not pay.

Figure 4 shows the revenue earned by the ransomware author for a range of different levels of backup use in the company population. Increasing the percentage of companies using backups has only a small effect on the total revenue until 70% of the population are using backups, after which it drops significantly. Until

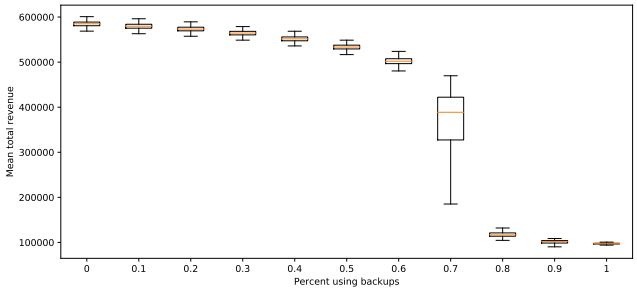

Figure 4: Revenue for ransomware author for different levels of backup use in the population.

this point, it is more profitable to continue charging higher fees, even though much of the population does not pay.

4.2 Price Perturbation

Setting a ransom price has consequences for the ransomware author. The first is, naturally, to create revenue. The second is to learn more about the willingnes to pay of the population. The basic policy is to use the myopic best price at each time period. However, dithering, or perturbing the price may increase the amount of information learned, leading to greater longer-term revenue [7].

We try three different perturbation policies. The first two are drawn from skew-normal distributions (with scale parameter 50), using skew parameter values of 4 or -4; this affects the price either positively or negatively. The final policy uses normally-distributed perturbations, affecting the price in either direction.

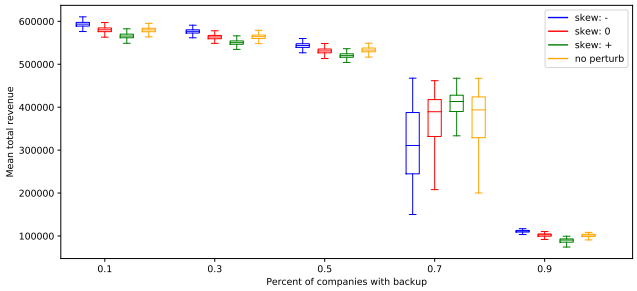

Figure 5: Revenue using different skew values for perturbing the offered price, for different levels of backup use in the population.

Figure 5 shows the results for these three perturbation policies at different backup percentages, compared to the basic, non-perturbed policy. In most cases, using a negative perturbation improves the total revenue over the original policy; the neutral perturbations result in roughly similar performance; and, the positive skew results in a decrease in performance compared to the basic policy. The exception is at the 70% use of backups. Here, the best expected price is between the two modes of the willingness to pay distribution, meaning that increasing the price means that extra revenue is earned with a small chance of exceeding a firm's willingness to pay. Similarly, reducing the price does not make it more likely that backup-users will pay, but still reduces overall revenue.

4.3 Changing Reliability

We now introduce into the model the notion of ransomware reliability. As with all software, ransomware sometimes has bugs that affect performance; here, it means that companies that pay the ransom might not get their data back. A company that knows a particular ransomware variant is unreliable will be less willing to pay the ransom. For example, *WannaCry* and *Petya*, despite their sophistication, were not capable of—or perhaps not designed to allow—decrypting the files after payment

This expected reliability, $\mathbb{E}[\rho]$, affects the expected reward companies get for paying the ransom. As reliablity goes down, so does the reward, and thus so does the company's willingness to pay.

For simplicity, we assume that all companies share the same estimation of reliability. We model this using a beta distribution $\rho \sim \beta(a, b)$ where a and b are the number of successes (data restored) or failures (data lost) respectively, with initial values $a = 1, b = 1$. The expected value of this initial state is 0.5: companies have no knowledge of the reliability. The parameters for ρ are updated after each successful or unsuccessful ransom payment, where success or failure is stochastic, based on the reliability assigned to the ransomware in the simulation.

Figure 6 shows the revenue for high-reliability ransomware ($P(\text{success}) = 0.95$). The total revenues are similar to, but slightly lower than those from the previous example (Figure 5), which had fixed perfect reliability. The reason for the decrease is that the initial belief among the population about the reliability is .5, depressing the willingness to pay until the high reliability is learned.

Figure 7 shows the revenue for low-reliablity ransomware ($P(\text{success}) = 0.2$). The total revenue in all cases is much lower here, reflecting the populations low willingnes to pay for an unreliable return of data. For low levels of backup in the population, the positively skewed price policy seems to have a slight advantage. This is possibly because the updates about reliability occur only after a firm decides to pay; by increasing the ransom price, the revenue collected will be greater before information about the poor reliability spreads.

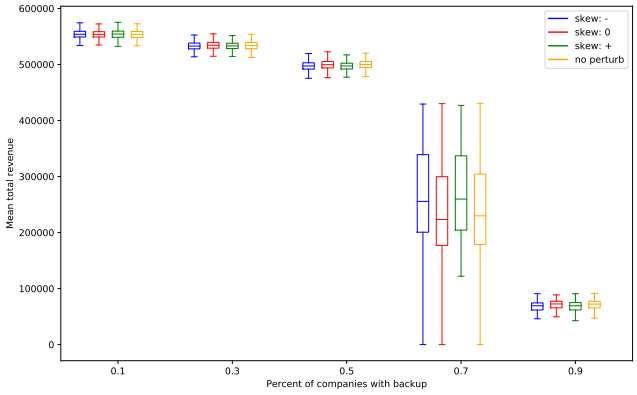

Figure 6: Revenue for high-reliability ransomware.

4.4 Information Sharing

Finally, we look at the effects of the rate at which information about reliability is shared. Many firms do not publicly announce that they have been hit by ransomware, nor do they announce whether the payment was successful in restoring their data. We introduce a connectivity parameter to model the rate at which firms share information, which determines the probability of the reliability estimate being updated after a payment. A small probability models a less-connected network where companies are less willing to share data; a large probability models a more-connected network, where firms are happy to disclose.

Figure 8 shows the effects of different levels of connectivity at different ransomware reliability levels for a population with a low (20%) backup rate. Figure 9 shows the same, but for a 70% backup rate.

There are two interesting observations here. First, at low reliability levels, ransomware revenue is boosted by low connectivity: if news about the poor reliability does not spread, more companies will likely pay the ransom and not recieve their data. Second, more reliable ransomware tends to be boosted by higher connectivity: greater information sharing leads to increased awareness of the high reliability, increasing the willingness to pay.

5 Conclusions

Although this paper is focused on policies for pricing ransomware ransoms, there are useful conclusions and insights to draw for the purposes of preventing ransomware.

First, there will be no herd immunity or ability to freeride, at least until extremely high levels of backup use in the population. The profitability of

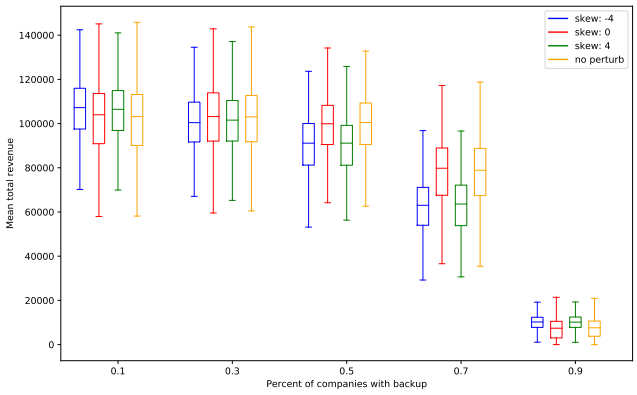

Figure 7: Revenue for low-reliability ransomware

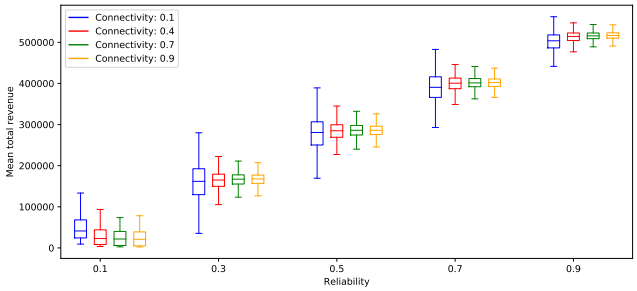

Figure 8: Information sharing: low percent of backup use (%20).

ransomware does not really decrease significantly until at least 70% of the population is using backups; until this point, there is still a lot of incentive to create and spread ransomware. Firms should implement their own backup strategies.

Second, once a high rate of backup use among companies has been achieved, the revenue possible will drop sharply; this will reduce the incentive to spend a lot on the development of sophisticated, reliable ransomware (assuming that increased reliability has increased development cost). Ransomware produced after this will likely be lower-cost and less reliable—and hopefully easier to prevent—aiming to target the limited revenue available from the remaining firms who have not yet implemented a backup strategy.

Finally, information sharing between firms is important. Although there is a trade-off—slow (or less) information sharing helps low-reliability ransomware, and fast information sharing helps high-reliablity ransomware—it is better to share more information and help the more reliable version. If ransomware is

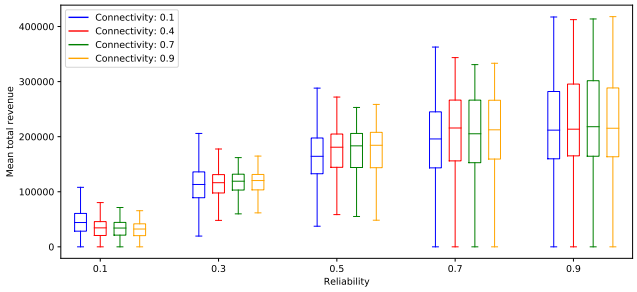

Figure 9: Information sharing: high percent of backup use (%70).

reliable, it means that companies are likely to get their data back if the ransom is paid. For low-reliability ransomware, this is not the case. It is better to increase awareness and avoid the potential twin costs of lost data and unsuccessful ransom payment.

References

[1] Edward Cartwright, Julio Hernandez-Castro and Anna Stepanova. 'To pay or not: game theoretic models of ransomware'. In: *WEIS 2018*. 2018.

[2] Li Chen. *Bayesian dynamic pricing with two-sided censored customer willingness-to-pay data*. Tech. rep.

[3] J. Michael Harrison, N. Bora Keskin and Assaf Zeevi. 'Bayesian Dynamic Pricing Policies: Learning and Earning Under a Binary Prior Distribution'. In: *Manage. Sci.* 58.3 (Mar. 2012), pp. 570–586. ISSN: 0025-1909. DOI: 10.1287/mnsc.1110.1426. URL: http://dx.doi.org/10.1287/mnsc.1110.1426.

[4] Julio Hernandez-Castro, Edward Cartwright and Anna Stepanova. 'Economic Analysis of Ransomware'. In: *CoRR* abs/1703.06660 (2017). arXiv: 1703.06660. URL: http://arxiv.org/abs/1703.06660.

[5] IBM. *IBM study: Businesses more likely to pay ransomware than consumers*. 2016. URL: http://www-03.ibm.com/press/us/en/pressrelease/51230.wss.

[6] Aron Laszka, Sadegh Farhang and Jens Grossklags. 'On the Economics of Ransomware'. In: *Decision and Game Theory for Security*. Ed. by Stefan Rass et al. Cham: Springer International Publishing, 2017, pp. 397–417. ISBN: 978-3-319-68711-7.

[7] Miguel Sousa Lobo and Stephen Boyd. 'Pricing and learning with uncertain demand'. In: *INFORMS Revenue Management Conference*. 2003.

[8] Masarah Paquet-Clouston, Bernhard Haslhofer and Benoît Dupont. 'Ransomware Payments in the Bitcoin Ecosystem'. In: *WEIS 2018*. 2018.

[9] Michael Rothschild. 'A two-armed bandit theory of market pricing'. In: *Journal of Economic Theory* 9.2 (1974), pp. 185–202.

[10] Symantec. *2018 Internet Security Threat Report: Executive Summary*. 2018. URL: https://www.symantec.com/content/dam/symantec/docs/reports/istr-23-executive-summary-en.pdf.