To Mitigate Insecure IoT Devices, Regulate Manufacturers or Consumers?

Matheus Ferreira, Tithi Chattopadhyay, Nick Feamster, Matt Weinberg, Danny Yuxing Huang
Princeton University

I. Introduction

Recent years have seen the emergence and proliferation of Internet-of-Things (IoT) devices. However, many of them are found to contain security vulnerabilities that adversaries could exploit remotely to launch high-profile attacks [1]. For example, the Mirai botnet compromised a few hundred thousand Internet-connected cameras, before it launched a distributed denial-of-service (DDoS) attack on a DNS provider used by Twitter and Reddit, causing widespread disruption on these major platforms [2]. As the total number of IoT devices is projected to reach a billion in the next five to ten years [3], attacks from insecure devices are likely to cause higher damages.

These problems have led to some effort from the industry to create security standards among manufacturers of IoT devices [4]. Additionally, one consumer advocacy group offers security evaluations of devices, in the hope of helping buyers make more informed decisions [5]. For both manufacturers and consumers, however, such momentum for security efforts remains to be seen. One reason is the lack of incentives to adopt security practices. For the manufacturers, security features could lead to higher costs [6], which often translate into higher prices for consumers. For consumers with potentially insecure and vulnerable devices, improving device security, such as regularly downloading any available firmware upgrades, takes extra effort [7]. If a device is compromised and used to attack other services online, as in the case of the Mirai botnet, few manufacturers or owners suffer significant losses from the attack.

To mitigate this lack of incentives, we propose a policy that regulates the manufacturers and/or consumers of IoT devices. Specifically, we **require the manufacturers to enforce minimum security standards** for their devices — for instance, setting strong passwords or encrypting the network traffic — to reduce the probability of being compromised. Alternatively, we **fine owners** of IoT devices which are compromised and used to attack other services on the Internet, in the hope that consumers will favor more secure devices and thus drive less secure ones out of the market.

II. Methodology

One challenge we face is whether to regulate the manufacturers, the consumers, or both, as well as how to regulate them. If we require the industry to enforce minimum security standards, the cost of production is likely to increase. What kinds of security standards should we impose? How much would the cost of production increase? How would this rise in cost affect the device prices, the consumers' purchase decisions, and the manufacturers' profits? On the other hand, if we fine the consumers, device security is likely to play a bigger role in purchase decisions. How much should the penalty be? How would different levels of fines affect the manufacturers' profits and the overall risk of compromised devices?

These questions are hard to answer, especially when the manufacturers and consumers display non-uniform behaviors. For example, manufacturers have varying degrees of financial means and technical expertise. Even if we impose the same minimum security standards, the extra cost is likely to vary across device makers. For the consumers, a fine may discourage risk-averse buyers from insecure devices, but it may have limited effect on risk-tolerant consumers who believe their insecure devices are unlikely to be compromised. In general, if we consider such complexities when designing the policy, it would be difficult to isolate the key interactions between the manufacturers and consumers and provide meaningful policy guidance.

To this end, we propose a simplified model, where one Seller tries to sell identical Devices of the same type to a population of Buyers interested in the product. We assume that Seller knows some aggregate statistics about Buyers' behaviors, such as their willingness to regularly update the firmware. We also assume that Seller treats every Buyer in the same way, and each Buyer makes purchase decisions independently.

The goal of Seller and Buyers is to maximize their *utility*. Formally, we define Seller's utility, u_S, as follows:

$$u_S = (p - c)\Pr[u_B \geq 0] \tag{1}$$

where p is the price of Device set by Seller, and c is the cost as a result of the minimum security standards. In today's setting, $c = 0$, but as we impose a higher level of minimum security standards, c will increase for Seller. In addition, $\Pr[u_B \geq 0]$ is the fraction of Buyers who decide to purchase Devices at p; for each of these buyers, the utility is positive for the given p. We define Buyer's utility, u_B, as follows:

$$u_B = v - p - h - l \tag{2}$$

where v is Buyer's own valuation of Device, h is Buyer's *hygiene* — i.e., effort to secure Device — and l is the expected loss if Device is compromised. Two factors, in particular, affect l: the fine due to any compromise and the probability of compromise. As such, we define l as:

$$l = ye^{-c-hk} \tag{3}$$

where y is the fine, e^{-c-hk} is the probability of compromise, and k defines the effectiveness of Buyer's hygiene effort.

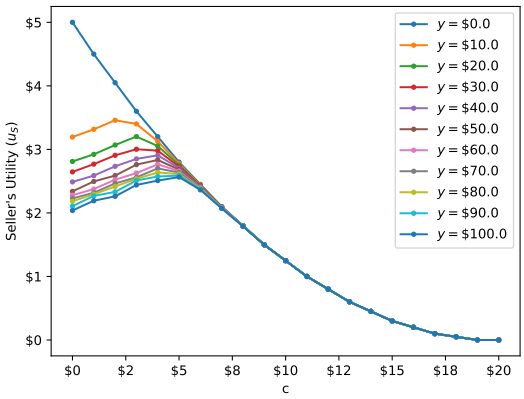

Fig. 1. How Seller's utility changes with different standards of minimum security, as reflected by c, the cost associated with these security standards, when we vary the penalty, y imposed on Buyer.

Currently, $y = 0$ as consumers are not penalized if their device is compromised. If $0 \leq k < 1$, Buyer is less effective than Seller in securing Device; if $k \geq 1$, Buyer is at least as effective as Seller. We use an exponential function when defining the probability of compromise, because when $c = 0$ and $h = 0$, there is no minimum security standards for Seller, and Buyer does not follow hygiene practices; hence the probability of compromise is 100%. In contrast, as c and h increase, both Seller and Buyer put in effort to secure the device; thus the probability of compromise is close to zero (but never reaching zero).

III. Validating the Model

Of all the parameters in the model, our proposed policy changes only c and y. In particular, if we impose minimal security standards on device makers, then Seller's c will increase from the current level of zero. On the other hand, if we fine consumers whose devices are compromised and used for attacks, Buyer's y will also rise from the current level of zero.

To see how changing y and/or c affects the utilities of both parties, we run simulations to validate our simple model.[1] First, we simulate a policy that requires the manufacturers to implement some minimum standards of security but which does not impose any fines on consumers. As shown in Figure 1, as c increases while $y = 0$, u_S monotonically decreases. This observation is consistent with our intuition. The cost due to minimum security standards increases Device's price; Buyers are thus less likely to make the purchase.

Using the model, we further explore another policy that fines consumers but which does not require minimum security standards from device makers. As shown in the figure, increasing y while $c = 0$ leads to a drop in u_S. As expected, Buyer is less likely to purchase Device that lack any minimum security standards, due the risk of compromise and the possibility of getting fined. However, when $y > 0$ and $c = 0$, u_S is not the highest at any given value of y. Rather, Seller achieves maximum u_S when c increases to a certain value, although any further rise in c reduces u_S again. In other words, **if we penalize Buyers, Seller can achieve a higher utility if it voluntarily enforces just the right amount of security standards**. Security standards that are too low add to Buyer's risk of being penalized, whereas excessive security standards result in a higher price and lower willingness to buy.

[1] In our simulation, we assume that $k < 1$ with more than 50% probability. We will not discuss the case where $k \geq 1$ with 100% probability, as Buyer is more effective in securing Device than Seller. Implementing any security features, i.e. increasing c, always results in a lower u_S regardless of y values.

IV. Next Steps

Our initial observation of the interplay among y, c, and u_S is consistent with our intuition and thus offers a preliminary validation of the model. However, we are yet to determine if our model is still accurate in more general cases. For instance, we can change the values of v and k to simulate consumer segments with different valuations of Device and different effectiveness of hygiene practices. It is an open question whether the behaviors in Figure 1 will also be observed for these consumer segments. Moreover, our current model assumes a single seller with a single type of device. Having multiple sellers and devices in the model, while more realistic, will introduce competition among sellers. Under this setting, it remains to be seen whether the introduction of fines on consumers will still incentivize device makers to voluntarily adopt some levels of security standards.

We stress that imposing minimum security standards on manufacturers and/or fines on consumers is one of the many policies that could potentially mitigate the risk of insecure IoT devices. We plan to extend our model to simulate other similar policies. For instance, rather than have consumers pay a penalty if their devices are compromised, we can allow consumers to purchase an insurance. We will determine if fines or insurances are more effective in incentivizing consumers to buy more secure devices. Alternatively, instead of requiring minimum security standards for device makers, we can certify devices that have achieved these standards. Similar to how nutritional labels inform consumers of the risks of unhealthy ingredients, device certification informs consumers of the risks of using insecure devices. Again, it is an open question which option — whether imposing minimum security standards or certifying devices — can more effectively lead to an overall increase in device security in the industry.

In addition to using models, we plan to conduct empirical measurements to observe human behaviors under different policies. As the first step, we will measure how much consumers are willing to pay for devices of various levels of security (i.e., varying c) under different levels of perceived risks (i.e., effectively varying y). Similar to previous studies [8], we will ask subjects in a laboratory setting to purchase IoT devices with the variables above. Our hypothesis is that, for each IoT device, if we tell the subjects the level of security and the implication on their privacy, they are likely to purchase more secure devices at higher prices than if we hide the risk-related information (i.e., $y = 0$). We plan to measure how much extra subjects are willing to pay for security across different types of IoT devices (e.g. smart light bulbs *versus* home security cameras, which have more privacy-related risks than light bulbs). Our hope is that the price premium will be consistent with the observation in Figure 1, especially for devices with higher perceived risks.

REFERENCES

[1] Brian Krebs. IoT Reality: Smart Devices, Dumb Defaults. *Krebs on Security Blog, https://krebsonsecurity.com/2016/02/iot-reality-smart-devices-dumb-defaults/*, 2017.

[2] Nicky Woolf. DDoS attack that disrupted internet was largest of its kind in history, experts say. *The Guardian, https://www.theguardian.com/technology/2016/oct/26/ddos-attack-dyn-mirai-botnet*, 2016.

[3] James Manika et al. The Internet of Things: Mapping the Value Beyond the Hype. *McKinsey Global Institute, http://www.mckinsey.com/business-functions/business-technology/our-insights/the-internet-of-things-the-value-of-digitizing-the-physical-world*, 2015.

[4] Open Connectivity Foundation. OCF Security. *https://openconnectivity.org/business/ocf-security*, 2017.

[5] Consumer Reports. Consumer Reports to Begin Evaluating Products, Services
for Privacy and Data Security. *https://www.consumerreports.org/privacy/consumer-reports-to-begin-evaluating-products-services-for-privacy-and-data-security/*, 2017.

[6] Terrence August, Duy Dao, and Kihoon Kim. Market segmentation and software security: Pricing patching rights. *Workshop on the Economics of Information Security (WEIS)*, 2016.

[7] Jay Pil Choi, Chaim Fershtman, and Neil Gandal. Network security: Vulnerabilities and disclosure policy. *The Journal of Industrial Economics*, 58(4):868–894, 2010.

[8] Serge Egelman, Janice Tsai, Lorrie Faith Cranor, and Alessandro Acquisti. Timing is everything?: the effects of timing and placement of online privacy indicators. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 319–328. ACM, 2009.