Should We Outlaw Ransomware Payments?

Debabrata Dey
Foster School of Business
University of Washington, Seattle
ddey@uw.edu

Atanu Lahiri
Jindal School of Management
University of Texas, Dallas
atanu.lahiri@utdallas.edu

Abstract

In recent times, there has been an upsurge in ransomware attacks, where an attacker encrypts a user's files and then demands a ransom in exchange for the decryption key. While paying the ransom allows the user to quickly unlock the locked files and avoid potentially larger losses, it also strengthens the hands of the attacker and increases the chance of a future attack. We study this dilemma of the victims and the externality posed by their actions using a game-theoretic model on top of a Markov decision process. The resulting equilibrium leads to several interesting insights such as that legally prohibiting ransom payments may not always have the desired economic effects—in some cases, a ban may be effective in addressing the economic externality but, in others, it could reduce public welfare. Our findings have important implications for policymakers who are currently debating legislation that, if enacted, will outlaw ransom payments to attackers.

Keywords: *Ransomware, Markov decision process, information security, externality, social cost.*

1. Introduction

Ransomware, a form of computer malware, has quickly become one of the top cybersecurity threats faced by today's organizations. In the United States alone, it has cost businesses and governmental institutions more than $7.5 billion in 2019 [3]. According to the Federal Bureau of Investigations, nearly 1,500 cases were reported in 2018 and, on average, each victim suffered a loss of an eye-popping $3.6 million [10, 13]. Clearly, ransomware, which used to be a nuisance primarily to individual users, has hit the big market and has started impacting us all, directly or indirectly.

As apparent from the name itself, ransomware seeks to make money through extortion. It spreads primarily through phishing scams or by compromising unpatched systems, and once it successfully infects a target, it starts encrypting files on the victim's computer, files such as documents, images, videos, and perhaps most importantly, transactional databases. After encrypting a large number of such files, it notifies the victim that his files have been rendered inaccessible and that he must pay a ransom to recover the encrypted files. In other words, victims must pay for the decryption key, or else they will face significant data loss and business disruption.[1]

Interestingly, victims, although unwilling to admit so publicly, often end up paying substantial amounts as ransom. For example, Travelex, a well-known financial firm, has paid $2.3 million following a breach in the late December of 2019 [12]. Only a couple of months earlier, DCH Health System, a medical center operator, paid an undisclosed sum to its attackers [2]. In fact, it is quite common for companies to pay such ransoms in cryptocurrencies such as Bitcoin [10], which makes the transactions and their recipients quite difficult to track. Moreover, quite often, perpetrators of such crimes operate from a location that is outside of the legal reaches of the nation to which the victim belongs.

The list of victims consists of not just private enterprises; even governmental entities such as city, town, and county governments have been compelled to pay such ransoms. In Florida, for example, two cities, Lake City and Riviera Beach, had to pay $500,000 and $600,000 respectively, while Le Porte County in Indiana paid $130,000. As of August 2019, more than 70 state and local governments suffered ransomware attacks, and many of them eventually succumbed to the demands of attackers [14]. In the rare cases that they refused, they were hit with a much bigger loss and recovery cost; for example, when the City of Atlanta, following a breach in March 2018, refused to pay a $51,000 ransom, it had to spend a whopping $17 million to rebuild its systems [12]. In short order, the City of Baltimore suffered a similar fate when it refused to comply with a ransomware demand. Such worries for a much larger loss have created a surreal scenario where city councils, state governments, and even police departments are lining up after a breach to pay the ransom necessary to get their stuff back [8]. The lesson

[1] Of course, there is no guarantee that the attacker would honor its word upon receiving the ransom payment. As a result, a victim firm may question in the first place the value of acceding to the ransom demand. Such dilemma notwithstanding, the compelling reality of today is that a large number of firms end up trusting the attacker and paying the ransom.

URI: https://hdl.handle.net/10125/71414
978-0-9981331-4-0

is thus clear: paying a ransom is certainly not the only option, but doing so is often the cheaper way to restore and recover.[2]

However, with more frequent ransom payments has risen the controversy surrounding them. The reason is clear—there seems to exist an economic externality involved in such payments. As organizations start paying their attackers, they also end up encouraging these "bad" guys to come back for more [9], amplifying in the process the risk of future attacks on everyone, including themselves. According to Shi [14], "Ransom payments fuel the efforts of the cybercriminals. Hackers use that money to become more capable, commit more crimes, and expand their operations." Criticizing the shortsightedness of ransom-payers in this regard, Shi actually calls for a legal prohibition on all such payments. Shi is not alone in this demand. Outraged by increasingly large ransoms and an increased frequency of attacks, many others are raising their voice in favor of banning ransom payments altogether [15]. In fact, some policymakers have already started to heed this advice. For example, two bills have been introduced in the New York State Senate to prohibit municipalities from paying ransomware attackers, one to ban the practice of paying a ransom with taxpayer dollars and the other to ban it entirely [11, 15].

In this backdrop, several questions arise naturally: Should policymakers indeed outlaw ransom payments? Will doing so actually alleviate the situation, by trading off short-term benefits for a long-term gain? In other words, can a ban be effective in addressing the economic externality involved in ransom payments? These questions are not only relevant, but they are also unanswered.

The literature on economics of information security has grown considerably in recent years, with game-theoretic approaches gradually rising to prominence [e.g., 4, 5]. Interestingly, the issue of externality, which is highly relevant to this work, has also caught the attention of economists [16]. Further, recognizing the importance of ransomware and its broad implications, has emerged a new substream of research that focuses on the economics of ransomware. Among the key works in this substream, August et al. [1] consider the perspective of a software vendor, in particular how the underlying economic externalities affect the vendor's pricing strategy. In contrast, Hernandez-Castro et al. [6] examine the perspective of the ransomware attackers and how price-discrimination strategies employed by the attackers would impact social welfare. Laszka et al.[7] investigate the decision of the potential victims to invest in backup technologies and to what extent such technologies can serve as a deterrent. Although these papers examine important economic aspects related to ransomware attacks, they do not specifically address the issue of whether or not a ban on ransom payments can be economically beneficial in the long run, an issue eminently central to this work.

We address the issue by setting up a multi-period game involving two firms, where one firm's decision to pay a ransom in any period increases the probability of future attacks on both. The question we then ask is how these firms would fare with or without a ban. Our answers happen to be interesting. We find that, contrary to common wisdom, a ban is effective in mitigating the economic externality only in limited circumstances. Specifically, there are only two situations when a ban might work: (i) when there is an asymmetric equilibrium in which one firm resists and the other accedes to the ransom demand, (ii) when the multi-period game takes the form of a prisoner's dilemma, causing both firms to pay ransoms even when doing so is not mutually beneficial. In all other cases, a ban on ransom payments could be counterproductive. Policymakers, therefore, need to be careful before they institute laws to prohibit ransom payments.

2. Model

To illustrate the effect of how firms may react to ransomware attacks, we consider an infinite-horizon multi-period model in which future payoffs are discounted appropriately. For simplicity, we consider a game with two firms. At the beginning of each period, each firm experiences a breach with a positive probability. If breached, a firm has two options. It can make a ransom payment of r to the attacker to quickly restore its normal operations or, alternatively, it can refuse to pay the ransom and lose $c \geq r$. This loss of c may result from business disruptions as well as costs incurred towards recovery. Typically, ransoms provide firms a quick way out in the short term, which is why $c \geq r$. Let $\gamma = \frac{c-r}{r} \geq 0$ be the normalized additional cost associated with non-payment of ransom. It can also be viewed as the *loss ratio* or the defiance premium—the larger the loss ratio, the less likely is a firm to resist payment.

Even though $\gamma \geq 0$ and it is less costly in the short-run to simply accede to the ransom demand, it may

[2] Seals [13] tells an interesting story in this regard. A firm called Proven Data Recovery was claiming to make use of technology tools to clean up ransomware breaches. Secretly, however, they were actually paying the ransom while collecting a premium from their clients.

not be the best policy from a long-term perspective, with respect to private profit or public welfare. For, ransom payments can strengthen the hands of the attacker by providing encouragement as well as resources to orchestrate attacks and may, therefore, result in a higher threat level in the future [14]:

ASSUMPTION 1. *The breach probability faced by a firm in any period t is $\beta_n = \beta(1+n\alpha) < 1$, where $n \in \{0,1,2\}$ is the number of firms that paid ransom in period $t-1$, and $\alpha, \beta > 0$ are model parameters. Further, the event of a firm facing a breach is independent of other breaches.*

The parameter α in Assumption 1 represents the amplification fraction of the breach probability and captures the externality effect within this context; the higher the α, the larger is the risk a firm induces on the other as well as on itself by complying with a ransom demand. The parameter β, on the other hand, represents the inherent risk within the context; it captures the fact that, even when the externality effect is absent, firms will continue to face a residual level of risk. The last part of Assumption 1—stochastic independence of breach events—is not critical; it simply makes the exposition clutter-free.

Assumption 1 tells us that each firm can be in one of two states—breached (1) or safe (0)—in any given period, and that the probability of being breached in the next period depends on the total ransom payment to the attacker but not on the firm's current state. The state transition diagram for a firm can be shown in Figure 1. To elaborate, if both firms are breached in period $t-1$ and they both decide to pay a ransom in that period—that is, if $n=2$—they both risk a breach in period t with a probability of $\beta_2 = \beta(1+2\alpha)$. If only one of them pays a ransom while the other one refuses, the attacker is less encouraged and, accordingly, they both face a lower breach probability of $\beta_1 = \beta(1+\alpha)$ in the next period. Finally, if neither pays any ransom, the breach probability is just $\beta_0 = \beta$ in the next period. Thus, while the payment of a ransom increases the threat level in the following period, the refusal to pay has the opposite effect of retaining a lower threat level.

Since each firm can be in only one of the two states, 0 or 1, one of the four states, $\{00, 01, 10, 11\}$, is possible in any period t. Let v_{00} denote the total expected cost to firm 1 over the infinite time horizon when neither firm experiences a breach in the current period. Likewise, we can define v_{10} to denote the expected cost to firm 1 when only firm 1 experiences a breach, v_{01} to denote the expected cost to firm 1 when only the second firm experiences a breach, and v_{11}, when both firms experience a breach.

In general, firms may observe each other's state (breached or not breached), but they are unlikely to observe each other's actions (ransom paid or not). Although firms do not enjoy telling the world that they have experienced a breach, they usually end up doing so to alert their customers, investors, and other stakeholders.[3] Even then, ransom payments are done secretly, perhaps to avoid media and public scrutiny. Now, every period, in the event of a breach, each firm chooses between: (P) pay a ransom and (N) do not pay any ransom. Therefore, the strategy profile for the two firms can be denoted by $\mathcal{S} = \{(s_1, s_2) | s_1, s_2 \in \{\text{P,N}\}\}$. In general, firms may make their decisions simultaneously upon observing their states in the beginning of each period. However, to keep the exposition straightforward, we consider only a static game in which each firm takes the same action every period. Interestingly, all our results extend to the situation in which firms make their decisions in a dynamic fashion.

Let $\delta \in (0,1)$ be the per-period discount factor, implying that a dollar in the next period is worth only δ dollars today. We are now ready to estimate a firm's expected cost in this Markov decision process. We only estimate the expected cost for firm 1; that for the other firm can be obtained from the symmetry of the problem.

We start with the case where there is no breach in the current period. In that case, by definition, the expected cost to firm 1 is $v_{00}(s)$, $s \in \mathcal{S}$. Since neither incurring r nor c is necessary in this period, $v_{00}(s)$ is just the sum of costs incurred in the future, after appropriate discounting is applied. Therefore, for all $s \in \mathcal{S}$, we can write:

$$v_{00}(s) = \delta V(0), \text{ where}$$

$V(n) = (1-\beta_n)^2 v_{00}(s) + \beta_n(1-\beta_n)\left(v_{10}(s) + v_{01}(s)\right) + \beta_n^2 v_{11}(s)$. $V(n)$ simply represents the expected cost to firm 1 starting period $t+1$ if n firms pay the ransom in period t.

Similarly, when only firm 2 gets breached, but firm 1 does not, we get:

$$v_{01}(s) = \begin{cases} \delta V(0), & \text{if } s_2\text{=N} \\ \delta V(1), & \text{if } s_2\text{=P}. \end{cases}$$

[3] According to Neuhauser [8], most state and local governments have passed legislation on data-breach disclosure, thereby requiring private companies to inform affected customers about a breach and also report the same to state authorities.

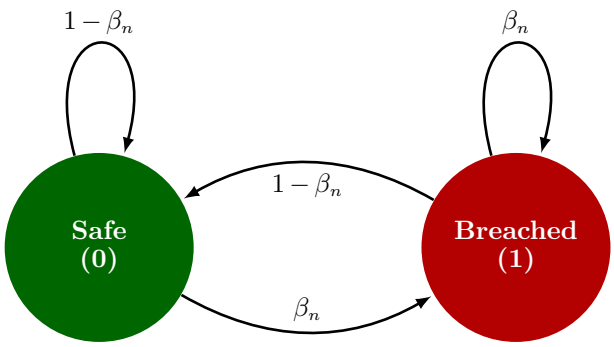

Figure 1. State Transition Diagram for a Firm

In contrast, if firm 1 is breached in the current period, it will not only incur the future expected cost, but it will also have to incur either r or c in the current period depending on whether it adopts N or P. Furthermore, firms' strategy, s, should be reflected in the breach probability, β_n. Therefore, we also get the following two relationships:

$$v_{10}(s) = \begin{cases} c+\delta V(0), & \text{if } s_1\text{=N} \\ r+\delta V(1), & \text{if } s_1\text{=P}, \end{cases} \text{ and}$$

$$v_{11}(s) = \begin{cases} c+\delta V(0), & \text{if } s\text{=(N,N)} \\ c+\delta V(1), & \text{if } s\text{=(N,P)}, \\ r+\delta V(1), & \text{if } s\text{=(P,N)}, \\ r+\delta V(2), & \text{if } s\text{=(P,P)}. \end{cases}$$

Therefore, for a given $s \in \mathcal{S}$, we have four simultaneous equations which can be solved for the four unknowns. The results are summarized in Table 1.

3. Equilibrium

Based on the incurred costs specified in Table 1, we can now find out firm 1's optimal strategy given firm 2's and vice versa. For ease of exposition, the expected costs are arranged in the form of a cost (payoff) matrix shown in Table 2; firm 1's cost is the first entry in each cell of this matrix. Now, it can be shown from Table 1 that, regardless of the state x, $v_x(\text{N,N}) \le v_x(\text{P,N})$ if and only if $c \le \frac{r}{1-\alpha\beta\delta}$, or equivalently, $\gamma \le \frac{\alpha\beta\delta}{1-\alpha\beta\delta} \triangleq \gamma_1$. Furthermore, again regardless of the state x, $v_x(\text{P,P}) \le v_x(\text{N,P})$ if and only if $c \ge \frac{r(1-\alpha\beta\delta)}{1-2\alpha\beta\delta}$, that is, iff $\gamma \ge \frac{\alpha\beta\delta}{1-2\alpha\beta\delta} \triangleq \gamma_2$. Clearly, γ_2 is always greater than γ_1, so the following result is immediate from Table 2:

PROPOSITION 1. *Not paying a ransom is the dominant strategy if $\gamma \le \gamma_1$. Paying a ransom is the dominant strategy if $\gamma \ge \gamma_2$. In all other situations, it is optimal to pay a ransom only if the other firm does not.*

Proposition 1 delineates the optimal strategy of each firm and, hence, also the equilibrium. Specifically, if γ is sufficiently large ($\gamma \ge \gamma_2$), no firm wants to pay c, and they both prefer paying r instead. As a result, the equilibrium strategy is to play P in every period. Exactly the opposite is true when γ is sufficiently low ($\gamma \le \gamma_1$). In between γ_1 and γ_2, however, we have a situation where the payoff matrix resembles that of the so-called "game of chicken," where one firm playing P and the other playing N in each period is an equilibrium. Formally:

PROPOSITION 2. *If $\gamma \le \gamma_1$, both firms playing N in every period is the equilibrium. If $\gamma \ge \gamma_2$, both firms choosing P is the equilibrium. If $\gamma_1 < \gamma < \gamma_2$, one firm playing P and the other choosing N is the equilibrium.*

Proposition 2 can be illustrated in Figure 2, where the two thresholds, γ_1 and γ_2, are plotted as a function of the externality effect parameter, α, thereby partitioning the entire (α,γ)-space into three regions, two symmetric and an asymmetric one nestled in between the two. Figure 1 clearly shows that the (N,N) region expands with an increasing α, as does the asymmetric region. In essence, when the externality effect increases, so do the resistance against paying ransom and the firms' willingness to tolerate a higher loss ratio.

It can be shown that the simplest approach of sticking to the same strategy—either P or N—is sufficient in this game. Further, since action history of the competitor is not observable, it is not possible to devise a response that is contingent on actions taken by the other firm in prior periods (such as a "tit-for-tat" strategy). Overall, as c increases relative to r, we eventually transition from (N,N) to (P,P) every period, but an asymmetric outcome is eminently possible even in our symmetric setting at moderate values of γ ($\gamma_1 < \gamma < \gamma_2$).

Table 1. Net Present Total Cost to Firm 1 in Different States

$s \rightarrow$	(N,N)	(N,P)	(P,N)	(P,P)
$v_{00}(s)$	$\frac{c\beta\delta}{1-\delta}$	$\frac{c\beta\delta}{(1-\delta)(1-\alpha\beta\delta)}$	$\frac{r\beta\delta}{(1-\delta)(1-\alpha\beta\delta)}$	$\frac{r\beta\delta}{(1-\delta)(1-2\alpha\beta\delta)}$
$v_{01}(s)$	$\frac{c\beta\delta}{1-\delta}$	$\frac{c\beta\delta(1+\alpha(1-\delta))}{(1-\delta)(1-\alpha\beta\delta)}$	$\frac{r\beta\delta}{(1-\delta)(1-\alpha\beta\delta)}$	$\frac{r\beta\delta(1+\alpha(1-\delta))}{(1-\delta)(1-2\alpha\beta\delta)}$
$v_{10}(s)$	$\frac{c(1-\delta(1-\beta))}{1-\delta}$	$\frac{c(1-\delta(1-\beta+\alpha\beta(1-\delta)))}{(1-\delta)(1-\alpha\beta\delta)}$	$\frac{r(1-\delta(1-\beta))}{(1-\delta)(1-\alpha\beta\delta)}$	$\frac{r(1-\delta(1-\beta+\alpha\beta(1-\delta)))}{(1-\delta)(1-2\alpha\beta\delta)}$
$v_{11}(s)$	$\frac{c(1-\delta(1-\beta))}{1-\delta}$	$\frac{c(1-\delta(1-\beta))}{(1-\delta)(1-\alpha\beta\delta)}$	$\frac{r(1-\delta(1-\beta))}{(1-\delta)(1-\alpha\beta\delta)}$	$\frac{r(1-\delta(1-\beta))}{(1-\delta)(1-2\alpha\beta\delta)}$

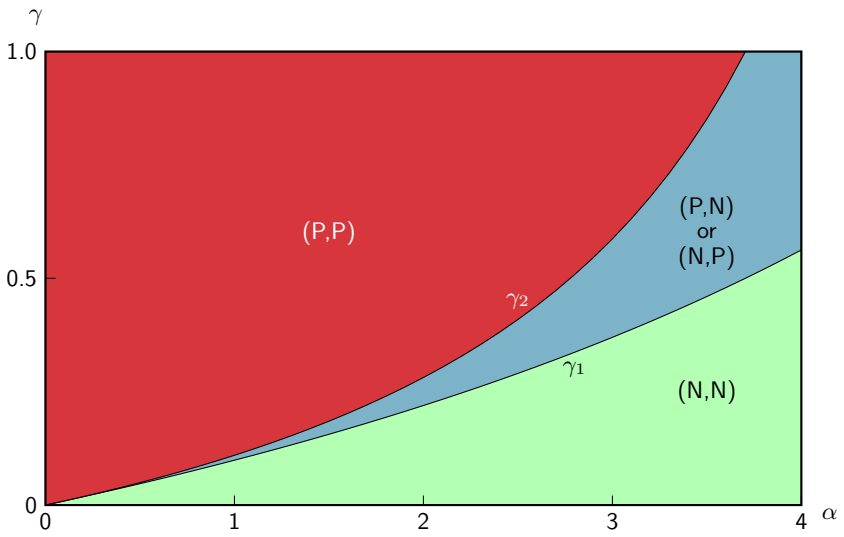

Figure 2. Equilibrium Regions for $\beta = 0.1$ and $\delta = 0.9$

Table 2. Cost Matrix For a Given State $x \in$ {00,01,10,11}

		Firm 2: N	Firm 2: P
Firm 1	N	v_x(N,N), v_x(N,N)	v_x(N,P), v_x(P,N)
	P	v_x(P,N), v_x(N,P)	v_x(P,P), v_x(P,P)

4. Policy and Welfare

We now examine whether or not the equilibrium outcomes described in the preceding section are also the first best, that is, if they are also optimal from a policymaker's point of view. Specifically, can a ban on ransom payments improve welfare and, if yes, under what conditions? And, if a ban is not possible or desirable, how else can a policymaker intervene to counter this threat?

To be able to answer these questions, let us compare the total cost in equilibrium with what the firms would have incurred together if ransom payments were prohibited. Interestingly, the result of this comparison depends on where we start the game. For the rest of the discussion, we will assume that, in the beginning, neither firm experienced a breach, that is, the game started with both being in the not-breached state in the first period. It is important to note that the following analysis as well as any insights obtained from it are robust qualitatively. Our conclusions would be similar even if we were to assume a different starting point—either 10, 01, or 11—instead of 00 assumed here.

To calculate the welfare, we only consider the costs borne by the two firms; we exclude from the welfare calculus the attacker and its gains. This is done for two reasons. First, these attacks often originate in other countries well beyond the jurisdiction of the nation to which the victim belongs. It is unlikely that a policymaker would be interested in counting any gains by a foreign entity. Second, it is not clear that the attacker would plow the extracted ransom back into the economy in a productive way. For example, the ransom payments may actually provide the

attacker with more resources to breach security in the future, which can be detrimental to welfare.

When ransom payments are not banned, the total cost incurred by the two firms in equilibrium, denoted v_{T} henceforth, can be written as:

$$v_{\mathrm{T}} = \begin{cases} 2v_{00}(\mathrm{N,N}), & \text{if } \gamma \leq \gamma_1, \\ v_{00}(\mathrm{N,P}) + v_{00}(\mathrm{P,N}), & \text{if } \gamma_1 < \gamma < \gamma_2, \\ 2v_{00}(\mathrm{P,P}), & \text{otherwise.} \end{cases} \quad (1)$$

In contrast, when ransom payments are completely banned, the total cost would simply be $2v_{00}(\mathrm{N,N})$. Comparing this with v_{T} leads to the following result:

Proposition 3. *Let* $\gamma_3 = \frac{2\alpha\beta\delta}{1-2\alpha\beta\delta}$. *Then,* $\gamma_3 > \gamma_2 > \gamma_1$. *Further, if* $\gamma_1 < \gamma < \gamma_3$, *prohibiting ransom payments improves welfare, i.e.,* $2v_{00}(\mathrm{N,N}) < v_{\mathrm{T}}$.

The result in Proposition 3 is best viewed in Figure 3, which partitions the (α, γ)-space into four regions. In two of these regions, the (N,N) region in its entirety and the portion of the (P,P) region shaded gray, the market outcome is the same as the first-best outcome, implying that the free market achieves the social optimum, and government intervention is unwise. A complete ban on paying ransom can, however, be the social optimum in the other two regions: (i) the entire asymmetric region, (P,N) or (N,P), where $\gamma_1 < \gamma < \gamma_2$, and (ii) the portion of the (P,P) region shaded in red, $\gamma_2 \leq \gamma < \gamma_3$. Notably, the ban works in slightly different ways in the two regions. In region (ii), the ban benefits both the firms, but in region (i), it hurts the ransom-payer even as it improves the total welfare.

Recall that, when $\gamma \geq \gamma_2$, both firms play P, i.e., pay a ransom of r in the event of a breach. They do so because c is significantly high compared to r. What Proposition 3 tell us is that, when $\gamma \geq \gamma_3$, this is indeed the most beneficial strategy for both firms. However, when $\gamma < \gamma_3$, it is unfortunately not the best for them to pay a ransom. Why do they then pay a ransom when $\gamma_2 \leq \gamma < \gamma_3$? The answer is simply that they face the classic "prisoner's dilemma"—even though (N,N) is the ideal spot for them to be in, absent any coordination, implicit or explicit, they are unable to reach that point, and they end up settling for an equilibrium that is actually not in their interest. Note that the fact that they cannot observe each other's actions makes any coordination through a tit-for-tat strategy practically impossible. Therefore, they remain locked in a prisoner's dilemma, which, although an equilibrium, is certainly not in their interest. If ransoms are outlawed, they will both gain in such a situation. The lessons for a policymaker is quite apparent. Banning ransoms can be futile if γ is very large ($\gamma \geq \gamma_3$), but such a policy should definitely be a consideration when firms remain locked in a prisoner's dilemma.

There is another lesson, which is that asymmetric outcomes are never socially beneficial. Banning ransom payments is always a good idea in such cases as well. However, a policy maker needs to be mindful that, in this case, the ban does not uniformly impact both firms. One firm, the one that pays ransom, will be hurt, while the other one that does not pay will gain. Collectively, though, the welfare increases.

Finally, when $\gamma \leq \gamma_1$, the situation from a welfare perspective is actually similar to what happens when $\gamma \geq \gamma_3$—in both cases, the equilibrium strategy turns out to be the most beneficial. Therefore, it is only when γ is between γ_1 and γ_3 that the equilibrium diverges from what is socially optimal.

5. Discussion

Current US law criminalizes extortion—receiving, possessing, or disposing of money that at any time has been delivered as ransom—but there is no generally applicable law prohibiting individuals or organizations from making such payments, except when such payments are being made to sanctioned entities such as terrorists. The question is thus simple. Should we enact new laws banning ransom payments in this era of grave cybersecurity threats? Although it may appear from our results that banning ransom may appear socially desirable in certain cases—specifically, when $\gamma_1 < \gamma < \gamma_3$—simply outlawing ransom payment could be a heavy-handed and often an undesirable way for the policymaker to intervene. Here is why.

First, in a significant portion of the region—specifically, when $\gamma \leq \gamma_1$—such an intervention is completely unnecessary. This is because firms are unlikely to pay the ransom in this region, irrespective of whether or not such a ban exists. The fact that this region expands as α, the externality parameter, increases tells us that, to an extent, the market can self-regulate itself and government intervention is essentially of no value. Furthermore, in the region where $\gamma \geq \gamma_3$, banning ransom would result in an outcome that is socially suboptimal, so a ban is detrimental in that portion of the parameter space as well.

Therefore, we are left with only a narrow strip of the parameter space ($\gamma_1 < \gamma < \gamma_3$) where a ban may potentially be of some value. However, even then, a ban could be an excessive measure. The fact remains that a firm's c, and hence its γ, is private information unknown to the policymaker. To complicate

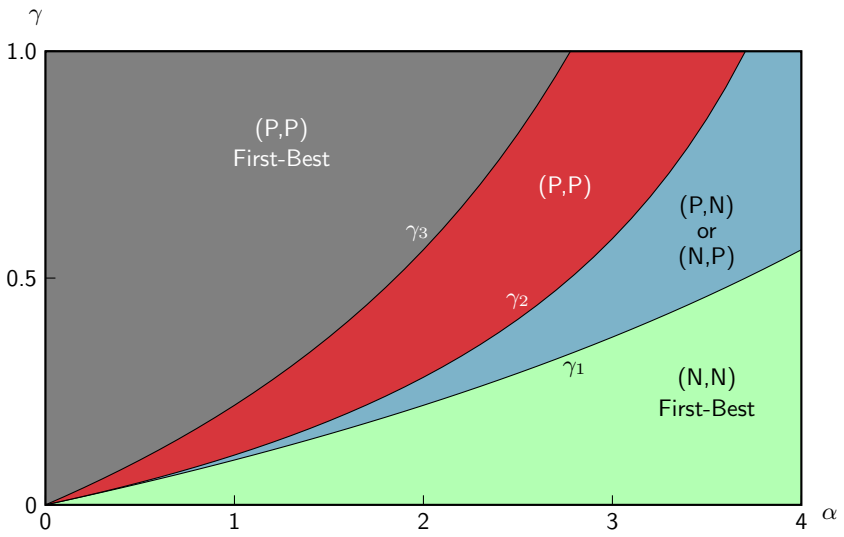

Figure 3. Equilibrium and First-Best Regions for $\beta = 0.1$ and $\delta = 0.9$

things further, based on investment in technology, a firm's γ could also change over time. Although such changes are outside the scope of the model, it is difficult to accept that a government can truly ascertain whether the current context indeed falls within this externality strip of $\gamma_1 < \gamma < \gamma_3$. Therefore, if the policymaker indeed takes the draconian step of banning all ransom payments for good, an unintended consequence of such a step could be an overall loss in social welfare. Finally, we must not forget that it is much easier to pass a law against ransom payments than to actually enforce it. For, it is often difficult to track these payments—large ransom payments are often broken into several smaller ones, are made using cryptocurrencies, and are routed through different accounts, with the money making multiple hops before finding its way to the attacker [13]. Therefore, if payments are outlawed, a firm facing a dire consequence may resort to illegal transactions, which could be difficult for a government to track and prosecute, making the underlying law ineffective in the process. In summary, a law prohibiting all ransom payments could be an overkill and should perhaps be avoided.

Now, even if a complete ban is not desirable, is it still possible for the policymaker to intervene? And, can it act in a way that eventually incentivizes firms so that they refrain from complying with ransom demands? The short answer is yes; it can. This the policymaker can do by reducing γ and bringing the context to the green (N,N) region in Figure 3.

How can the policymaker reduce γ? Recall that $\gamma = \frac{c-r}{r}$ can be decreased either by decreasing c or by increasing r. Therefore, the policymaker can reduce γ at once by partially bailing out firms who refuse to pay the ransom. If necessary, the policymaker can tie such bailouts to a precondition that a firm can receive a bailout only once and the firm receiving the governmental aid must invest a certain portion of the aid towards more sophisticated recovery systems and better security education of its employees. This way, the firms can be incentivized to not accede to ransom demands and to reduce the externality effect it imposes on other firms.

Corporate bailout is not necessarily the only strategy the policymaker can adopt. It can also tax the ransom payments themselves, thereby effectively increasing the r that a firm must pay. Of course, implementing a tax can still be a challenge. As mentioned earlier, ransom payments are typically done in a manner that is difficult for the government to track [13].

Interestingly, these two approaches are not mutually exclusive, and the policymaker can implement both, that is, it can tax ransom payments and, at the same time, provide aids to firms for refusing to comply with ransom demands. This way, by a judicious mix of tax and subsidy, the policymaker can not only stop ransom payments effectively, but it can also do so in a somewhat revenue-neutral manner.

6. Conclusion

Ransomware attacks are becoming more and more frequent these days, and, at the same time, there has also been a steady increase in the amount demanded as ransom. Since refusing to pay can become very

costly, the victim often ends up paying the ransom. However, such an action by an individual firm has an externality. By paying the ransom, the firm essentially strengthens the hand of the attacker, which in turn increases the intensity of such attacks in the future. Within this context, we wanted to study whether the market can adequately address this externality and, in case the market fails to fully internalize it, how a policymaker should intervene.

We found that, although a complete ban on ransom payments could be an overkill, the policymaker can indeed push victims towards non-payment by adopting suitable subsidies and/or taxes. Of course, our results should not be taken to mean that we are suggesting that policymakers all over should immediately try to intervene. Our implications, if any, are that a policymaker must pause and consider the unintended consequences carefully before intervening. And, if there must be an intervention, it should not be draconian.

In fact, policymakers cannot be faulted if they decide to not intervene, leaving the market to itself even within the externality region ($\gamma_2 < \gamma < \gamma_3$). By not intervening, a policymaker could see, to the detriment of public welfare, a growth in such attacks in the short term, but the long-term consequences could actually be quite desirable. This is because, as the saying goes, once bitten twice shy! When the threat of such attacks grows in the short run, individual firms would have all the incentives necessary for them to invest in sophisticated backup and recovery technology that can restore its critical business functions quickly and cheaply, making c go down drastically. As more and more firms adopt proper recovery technology, γ is likely to go down, pushing the firms towards the green (N,N) region in Figure 3. Therefore, the market may be able to address this externality in the long run, even though there could be some short-term blues.

Our work has a few limitations. In order to keep the analysis simple, we only consider a static policy of P or N, throughout the time horizon. A more general case would be where firms can dynamically change their ransom-payment behavior from one period to the next depending on the state that they are actually in. Our preliminary analyses indicate that there is no material impact on our results, insights, and conclusions when a dynamic policy is adopted. However, a formal treatment is necessary before we can truly generalize the results. Also, our setup considers only two firms; for the results to be useful, they must extend to any number of firms. Finally, we assume that the breach probability, β, is static and exogenous. However, this probability may change as firms invest more in IT security. We are working on these issues to get a more complete picture of the economic incentives that underlie this important context.

References

[1]T. August, D. Dao, and M. Niculescu, "Economics of Ransomware Attacks," Social Science Research Network, March 12, 2019. URL https://ssrn.com/abstract=3351416.

[2]B. Barth, "Dread Zeppelin: Ransomware targets health care and IT sectors in U.S. and Europe," SC Media, December 13, 2019. URL https://www.scmagazine.com/home/security-news/ransomware/dread-zeppelin-ransomware-targets-health-care-and-it-sectors-in-u-s-europe/. Accessed June 11, 2020.

[3]S. Brown, "How does ransomware work and what technologies best prevent it?," Rutter Networking Blog, March 2020. URL https://www.rutter-net.com/blog/how-does-ransomware-work-and-what-technologies-best-prevent-it. Accessed June 11, 2020.

[4]H. Cavusoglu, S. Raghunathan, and W. T. Yue, "Decision-Theoretic and Game-Theoretic Approaches to IT Security Investment," Journal of Management Information Systems, 2008 (25:2), pp. 281–304.

[5]A. Ghoshal, A. Lahiri, and D. Dey, "Drawing a Line in the Sand: Commitment Problem in Ending Software Support," MIS Quarterly, 2017 (41:4), pp. 1227–1447.

[6]J. Hernandez-Castro, A. Cartwright, and E. Cartwright, "An economic analysis of ransomware and its welfare consequences," Royal Society Open Science, 2020 (7:3), pp. 1–14.

[7]A. Laszka, S. Farhang, and J. Grossklags, "On the Economics of Ransomware,", in GameSec 2017, the 8th Conference on Decision and Game Theory for Security, Vienna, Austria, October 2017, p. 5.C.

[8]A. Newhauser, "Can the Law Stop Ransomware?," US News and World Report, April 2018. URL https://www.usnews.com/news/national-news/articles/2018-04-13/can-the-law-stop-ransomware. Accessed June 17, 2020.

[9]D. Olenick, "Ransomware attack forces DCH Health Systems to turn away patients," SC Media, October 2, 2019. URL https://www.scmagazine.com/home/security-news/ransomware/ransomware-attack-forces-dch-health-systems-to-turn-away-patients/. Accessed June 11, 2020.

[10]———, "Ransomware: To pay or not to pay," SC Media, October 1, 2019. URL https://www.scmagazine.com/home/security-news/ransomware/ransomware-to-pay-or-not-to-pay/. Accessed June 11, 2020.

[11]———, "New York considers bills banning ransom payments," SC Media, January 27, 2020. URL https://www.scmagazine.com/home/security-news/government-and-defense/new-york-considers-bills-banning-ransom-payments/. Accessed June 14, 2020.

[12]———, "Travelex paid $2.3 million ransom, report," SC Media, April 10, 2020. URL https://www.scmagazine.com/home/security-news/ransomware/travelex-paid-2-3-million-ransom-report/. Accessed June 3, 2020.

[13]T. Seals, "Ransomware 'Remediation' Firm Exposed: Researchers Weigh in on Paying," Threatpost, May 17 2019. URL https://threatpost.com/ransomware-pay-or-not/144833/. Accessed June 17, 2020.

[14]F. Shi, "Ransomware Attacks: Why It Should Be Illegal to Pay the Ransom," Dark Reading, February 4, 2020. URL https://www.darkreading.com/risk/ransomware-attacks-why-it-should-be-illegal-to-pay-the-ransom/a/d-id/1336905?_mc=rss_x_drr_edt_aud_dr_x_x-rss-simple. Accessed may 1, 2020.

[15]B. Sussman, "As Ransomware Payments Double, Some Want Them Banned," Secureworld, January 27, 2020. URL https://www.secureworldexpo.com/industry-news/ransomware-payments-double-some-want-ransoms-payment-ban. Accessed May 3, 2020.

[16]H. R. Varian, "Managing Online Security risks," The New York Times, June 2000. URL http://www.nytimes.com/library/financial/columns/060100econ-scene.html. Accessed June 28, 2020.