Ransomware and Reputation

Anna Cartwright [1] **and Edward Cartwright** [2],*

[1] School of Economics, Finance and Accounting, University of Coventry, Coventry CV1 5FB, UK; ac8373@coventry.ac.uk

[2] Department of Strategic Management and Marketing, De Montfort University, Leicester LE1 9BH, UK

* Correspondence: edward.cartwright@dmu.ac.uk

Abstract: Ransomware is a particular form of cyber-attack in which a victim loses access to either his electronic device or files unless he pays a ransom to criminals. A criminal's ability to make money from ransomware critically depends on victims believing that the criminal will honour ransom payments. In this paper we explore the extent to which a criminal can build trust through reputation. We demonstrate that there are situations in which it is optimal for the criminal to always return the files and situations in which it is not. We argue that the ability to build reputation will depend on how victims distinguish between different ransomware strands. If ransomware is to survive as a long term revenue source for criminals then they need to find ways of building a good reputation.

Keywords: ransomware; repeated game; reputation; trust

1. Introduction

Broadly speaking, ransomware is a form of malware in which the victim of a cyber-attack is blackmailed. The term has recently become synonymous with crypto-ransomware wherein the files on a victims device are encrytped and a ransom is demanded for the key to decrypt those files [1–4]. While ransomware dates back to the AIDS Trojan in 1989, it has come to prominence over the last five years or so with an explosion in the number of ransomware strands and victims [5]. Critical to this rapid evolution has been the wider awareness of crypto-graphically sound techniques for encrypting files in a way that does not allow reverse engineering [6]. Such techniques mean a victim has no way of recovering encrypted files without a back-up or the private key held by the criminals.

Ransomware offers a viable, long term business model for criminals [7]. In particular, if a victim has no back-up and values the encrypted files then they may decide the ransom is worth paying. Survey evidence suggests that many businesses do indeed pay the ransom [8]. We also have evidence of large profits for ransomware criminals [9,10]. The example of South Korean web-hosting firm Nayana paying a $1 million ransom in 2017 amply demonstrates how lucrative ransomware can be for criminals. Given that the probability of being prosecuted is low, it is no surprise that we now see hundreds, if not thousands, of ransomware strands in the wild. Each strand may only last a few months before the authorities and security experts 'catch up' but new variants continually emerge.

While ransomware offers a long-term business model it is clear that this model relies on an element of trust between victim and criminal. Victims will only pay the ransom if they believe that paying the ransom gives them a good enough chance of getting their files back [11]. This, in turn, suggests that criminals would benefit from a reputation for honouring ransom payments. CryptoLocker and CryptoWall are two examples of ransomware strands that had a good reputation for returning files [12,13]. There are plenty of examples, like WannaCry, with a bad reputation. So, should the criminal return the files, take the money and run, or simply demand more money? Recent evidence suggests that victims who pay the ransom recover their data around 50% of the

time [14].[1] That may be enough to tempt victims to try their luck and pay up. So, have the criminals, consciously or not, hit upon a good strategy?

In this paper we analyze reputation formation in a simplified repeated game. In each period a new victim decides whether or not to pay the ransom and the criminal decides whether or not to return the files. We assume that victims learn from the past experience of other victims. We also allow that it is costly for the criminal to return access to files because of, say, the costs of a 'customer service' to guide victims on how to decrypt files. In a baseline case where victims are unresponsive to past experience we show that it is optimal for the criminal to not return files. In the other extreme where victims are highly responsive to past experience the criminal should always return the files. We then explore the middle ground to see where the tipping point lies above which it is optimal for the criminal to return files. We find that if a criminal's reputation will be based solely on their own actions then in all but the most extreme settings it is optimal for the criminal to return the files. If, however, the criminal's reputation is affected by what other criminals are doing then it may not be optimal to return the files.

We will discuss the implications of our findings for law enforcement more as we proceed. But we note here that the findings of this paper can help inform on the threat analysis of ransomware strands. If a particular strand has a poor reputation then it would be in the interests of law enforcement to advertise this, in order to undermine the business model of the criminals. Either way we would not expect the strand to last long. By contrast, a strand with a good reputation has a viable long term future, and so is of a higher threat. This may encourage copy-cats (who do not return files). Or it may be that criminals join behind a 'successful strand', as with Cerber [15]. Or that 'successful strands' are 'reborn' in the game of cat and mouse with law enforcement and security experts, as we saw with Locky [16]. All this suggests that reputation is a factor that should go into threat analysis, alongside more familiar factors such as the ease with which the strand can be detected by common anti-virus software.

Our work builds upon a number of papers using economics and game theory to analyze ransomware. Hernandez-Castro et al. (2017) show how the optimal ransom demand depends on the distribution of valuations in the population and the information criminals can discern about willingness to pay [7]. Caulfield et al. (2019) analyze how criminals can learn over time about the distribution of valuations in the population [17]. Caporusso et al. (2018) set out the basic rationale for using game theory to model the interaction between ransomware criminal and victim [18]. Laszka et al. (2017) consider the strategic role of back-ups [19] while Cartwright et al. (2019) look at the role of irrational aggression and deterrence [20]. We recognize that there is also a related literature looking at the game theory of hostage taking situations, typically framed in a terrorist context [20–23].

In none of the papers and models just mentioned is reputation explicitly modelled. Instead the focus is more on criminal and victim interaction in a one-shot context or how criminals can learn over time. The key contribution of our paper is, therefore, to move to a *repeated game* setting in which victims can collectively learn from past experience. This moves the focus somewhat away from the criminal onto how victims react to an attack. For instance, how do victims form beliefs and how does that influence their decision to pay the ransom [24]. We would argue that such questions should critically inform policy and advice on how to intervene in the ransomware business model. In particular, the criminals have already had time to learn from experience about how to make ransomware work and so can be expected to have a more sophisticated, forward looking strategy. To be ahead of the game we need to understand optimal behaviour in this setting.

We proceed as follows. In Section 2 we provide the model, in Section 3 we have our results, and in Section 4 we conclude.

---

[1]　Although it is not clear whether this is some criminals returning the files 100% of the time and some 0% of the time, or it is mixing by a particular criminal.

2. Model

There is a continuum of potential victims. Each victim i is characterized by the amount she values her files v_i. The population distribution of valuations is given by function q, where $q(v)$, for any $v \geq 0$, is the proportion of victims who value their files more than v. We shall as an example consider the linear distribution $q(v) = a - bv$, where a and b are parameters.[2] In stage 1 of the game a victim is randomly drawn from the population and her electronic device is attacked by a criminal. The victim knows her valuation v_i. The only information the criminal has is that v_i is drawn from distribution q. He does not, therefore, know v_i.

In stage 2 of the game the criminal chooses a take-it-or-leave-it ransom demand r. In stage 3 of the game the victim decides whether or not to pay the ransom. Let p denote the victim's choice, where $p = 1$ indicates pay and $p = 0$ indicates not pay. In stage 4 of the game the criminal decides whether or not to release the files back to the victim. Let g denote the criminal's choice, where $g = 1$ indicates return and $g = 0$ indicates destroy. The payoff of the victim is given by

$$u_i(v_i, r, p, g) = v_i g - rp. \tag{1}$$

For instance, the victims payoff is v_i if she gets her files back without paying the ransom (which could be equivalent to not being attacked) while her payoff is $-r$ if she pays the ransom and does not get her files back. The payoff of the criminal is given by

$$\pi(v_i, r, p, g) = rp - gc \tag{2}$$

where $c \geq 0$ is the cost of returning files. For instance, the criminal's profit is r if the victim pays the ransom and the criminal does not return the files.

Let us briefly comment on the interpretation of the cost parameter c. If the criminal can effortlessly return a victim access to her files, $c = 0$, then it is a weakly dominant strategy for him to do so. This makes it is 'easy' for the criminal to build a reputation. In reality, however, we can expect that there are costs to returning access to files. These costs may include: checking the payment matches a particular victim, returning the key to that victim, and, perhaps most importantly, guiding the victim on how to decrypt their files and dealing with queries about files that fail to decrypt [13]. This latter point is crucial in terms of reputation because giving back the private key is not enough—the files need to be successfully recovered by the victim in order that she would feel the ransom payment was honoured. Hence, c is likely to be positive. This provides an incentive for the criminal to not return files.

In the following we assume that the above game is repeated indefinitely in periods $t = 1, 2, \ldots$. Let v_t, g_t, r_t and p_t denote the respective valuation and choices in period t. Given that victims are chosen randomly a victim will only ever be a victim in one period. Victims are assumed to be self-focused and so they ignore the externality effect that paying the ransom may have on future victims. Specifically, the payoff of the victim in period t is

$$u_t(v_t, r_t, p_t, g_t) = v_t g_t - r_t p_t. \tag{3}$$

The criminal is forward-looking and discounts future payoffs with discount factor δ. His aggregate payoff is thus

$$\Pi = \sum_t \delta^t \pi(v_t, r_t, p_t, g_t). \tag{4}$$

We assume that the objective of the victim in period t is to maximize u_t and the objective of the criminal is to maximize Π. We also assume that both victims and criminal are risk neutral.

---

[2] We restrict $q(v)$ to lie in the interval $[0, 1]$ as appropriate.

Whether the victim is willing to pay the ransom will depend on his beliefs about the likelihood of the criminal returning the files. Let $\beta_t \in [0,1]$ denote the perceived probability of the criminal returning the files in period t if the ransom is paid. For simplicity we assume that the victim believes the probability of retrieving the files if the ransom is not paid is 0. Let *history* at time t be given by the vector $h_t = \{r_1, p_1, g_1, \ldots, r_{t-1}, p_{t-1}, g_{t-1}\}$ of past ransom demands, whether the ransom was paid, and whether files were returned. We will assume that beliefs β_t are a function of h_t and r_t.[3] Clearly, in reality, a victim would not have access to the complete history. Instead they would need to pick up snippets of information from personal and social networks, forums, search engines etc.

To look at how beliefs may be shaped it is useful to relate our model to that on reputation in repeated games. The basic model in this literature involves a long-lived player interacting with a sequence of short lived players [25–27].[4] For instance, a monopolist interacting with a sequence of potential entrants [30,31]. The key question addressed in the literature is whether the long-lived player has an incentive to generate a reputation. This could be a monopolist generating a reputation for being tough against entrants or, in our case, a criminal generating a reputation for honouring ransom payments.

Models of reputation are driven by some level of informational asymmetry between the long-lived and short-lived players. One variant on the theme is a game of imperfect information and perfect recall in which the short-lived players do not know the payoff function of the long-lived player but can observe all actions. In our setting this can be equated with victims not knowing c and/or δ [30]. Or it could be the short-run players put some positive probability the criminal is a 'commitment type' who will always honour ransom payments [31,32]. In either case beliefs β_t are shaped by history h_t because the past actions of the criminal reveal information about his type.

Throughout the following we take a partial equilibrium approach in which the criminal maximizes expected payoff taking as given the beliefs of victims. This allows us to abstract away from formally modelling incomplete information or equilibrium belief formation and focus on the incentives of the criminal. It also means we do not restrict to 'rational' victims. In an Appendix A, however, we briefly consider equilibrium belief formation and Bayes Nash equilibrium.

3. Results

As a trivial benchmark case let us begin with the case of *independent beliefs* in which β_t is independent of h_t for all t. An assumption of independent beliefs is appropriate in a setting where there are many strands of ransomware which are indistinguishable, or treated as indistinguishable by victims. Hence, victims lump all ransomware attacks together and beliefs are based on the overall probability of getting files back. In this case the criminal that we are modelling in our game would have no (or very limited effect) on the beliefs of victims.[5] It is simple to show that the criminal's optimal strategy is to never return files.

Proposition 1. *In the case of independent beliefs, if $c > 0$ it is optimal for the criminal to set $g_t = 0$ for all t.*

Proof. For any period t we can see that the victim will pay the ransom if and only if $r_t < v_t \beta_t$. Consider a specific period τ in which the ransom is paid. Given that β_t for all $t > \tau$ is independent of g_τ, returning the files costs the criminal c and has no benefit. □

---

[3] The victim observes r_t before making her decision and so beliefs may also be conditioned on this. In our results this will not be an issue because the ransom is constant over time. So in the text we use h_t as shorthand for h_t and r_t.

[4] A variation on this theme, less relevant to us, is a patient player interacting with a less patient player [28] or two long-lived players interacting a finite number of times [29].

[5] To formally capture this one would need a game with multiple criminals in which beliefs are shaped by the collective behavior of independent criminals.

Proposition 1 encapsulates the basic intuition that if reputation is irrelevant then criminals have no incentive to return files. In this setting the only rational belief for victims is $\beta_t = 0$. But, then victims are not willing to pay and so ransomware is not profitable. We can already see in Proposition 1 that reputation is vital if ransomware is to be a long term business model [11]. In particular, the ransomware business model is dependent on beliefs, β_t, being somehow tied with history, h_t.

To appreciate the potential power of reputation let us consider another extreme. Suppose that we have *grim-trigger beliefs* in which (1) $\beta_t = 0$ if $p_{t-1} = 1$ and $g_{t-1} = 0$, (2) $\beta_t = \beta_{t-1}$ otherwise. In this case, if the criminal once fails to return the files when a victim pays then no other victim will ever trust him in the future. If the criminal always return the files then initial beliefs are retained.[6] The following result shows that it can be optimal for the criminal to return files.

Proposition 2. *In the case of grim-trigger beliefs it is optimal for the criminal to set $g_t = 1$ when $p_t = 1$ for all t if δ is sufficiently large.*

Proof. Suppose that the criminal (a) charges a ransom $r_t^\star$, where $r_t^\star > c$ and $q(r_t^\star \beta_1) > 0$, and (b) sets $g_t = 1$ if $p_t = 1$, in all periods t. Consider period τ and suppose that the victim pays the ransom in this period. Should the criminal return the files? The expected future payoff of the criminal if he returns the files is

$$\Pi_\tau = r^\star - c + q(r^\star \beta_1)(r^\star - c)\left(\frac{\delta}{1-\delta}\right). \tag{5}$$

If the criminal does not return the files his future payoff is $\Pi_\tau = r^\star$. It is, therefore, optimal to return the files if

$$\delta > \bar{\delta} = \frac{c}{q(r^\star \beta_1)(r^\star - c) + c}. \tag{6}$$

It is clear that $\bar{\delta} \in (0, 1)$. □

To illustrate Propositions 1 and 2 consider the linear distribution $q(v) = a - bv$. In the case of independent beliefs suppose beliefs are fixed $\beta_t = \beta_1$ for some initial belief β_1. The expected profit of the criminal if he sets ransom $r^\star$ in each period t is

$$\Pi = \frac{q(r^\star)r^\star}{1-\delta} = \frac{(a - br^\star)r^\star}{1-\delta}. \tag{7}$$

Hence we get optimal ransom demand

$$r^\star = \frac{a}{2b}. \tag{8}$$

Consider next grim-trigger beliefs with the same initial belief β_1. Suppose the criminal sets ransom $r^{\star\star}$ and always returns the files if the ransom is paid. His expected payoff is

$$\Pi = \frac{q(r^{\star\star})(r^{\star\star} - c)}{1-\delta} = \frac{(a - br^{\star\star})(r^{\star\star} - c)}{1-\delta}. \tag{9}$$

Here we get optimal ransom demand

$$r^{\star\star} = \frac{a + cb}{2b}. \tag{10}$$

Note that the optimal ransom is higher in the case of grim-trigger beliefs because the criminal will pay the cost of returning the files. Returning the files is optimal (see the proof of Proposition 2) if

---

[6]　Grim-trigger beliefs are consistent with reputational models where a failure to return the files serves as a signal the criminal is not a commitment type [25].

$$\delta > \bar{\delta} = \frac{c}{\frac{a - cb}{2b} + c} = \frac{2bc}{a + cb}. \tag{11}$$

For plausible parameter values the value of $\bar{\delta}$ is likely to be low. For instance, if $a = b = 1$ and $c = 0.01$ we have $\bar{\delta} = 0.02$.

3.1. Sampling Recent Victims

Propositions 1 and 2 show that whether it is optimal for the criminal to return files will primarily depend on belief formation. For the remainder of the paper we consider a particular model of belief formation based on empirical frequencies. In motivating this approach we begin by recognizing that there is an extensive literature on learning in games [33,34]. It is widely recognized that beliefs typically do not obey Bayes rule (see Selten (1991) for an early, light-hearted take on the matter, with reference to models of reputation [35]). Here we focus on a belief-based model of learning in which the victim's beliefs are based on the past record of the criminal in returning files. Our approach is similar to that of Young (1993) [34,36] and consistent with beliefs being biased by event frequencies [37].

Our model of belief formation can be explained as follows. Given period τ and history h_τ let $P(h(\tau)) = \{t : p_t = 1\}$ denote the subset of periods in which the ransom is paid. Let n be an exogenous parameter that measures *sample size*. If $|P(h(\tau))| < n$ then we assume $\beta_\tau = \beta_1$ for some initial beliefs β_1. If $|P(h(\tau))| \geq n$ then denote by $P_n(h(\tau))$ the n largest values in $P(h(\tau))$, i.e., $P_n(h(\tau)) \subset P(h(\tau))$ are the most recent n periods in which the ransom was paid. We assume that

$$\beta_\tau = \frac{1}{n} \sum_{t \in P_n(h(\tau))} g_t. \tag{12}$$

In interpretation, this means that beliefs are based on the proportion of times the criminal returned files the last n times the ransom was paid. So, we can think of victims as sampling the last n victims who paid the ransom.[7]

The following result is analogous to Proposition 2 in showing that it can be optimal for the criminal to return the files. This must be the case if sample size is small. The logic here is that it is 'easy' for the criminal to get a 'bad reputation' when the sample size is small, and once he has that reputation nobody ever pays the ransom. So, the better long term strategy is to pay the ransom.

Proposition 3. *If the sample size is $n = 1$ and δ is sufficiently large then it is optimal for the criminal to set $g_t = 1$ when $p_t = 1$ for all t.*

Proof. In period t we can see that the victim will pay the ransom if and only if $r_t < v_t \beta_t$. Suppose that the victim pays the ransom in period τ. If the criminal returns the files and continues with a strategy of charging r_τ and returning files, then his expected payoff in subsequent periods is

$$\Pi = r_\tau - c + q(r_\tau)(r_\tau - c)\left(\frac{\delta}{1 - \delta}\right). \tag{13}$$

Note that this equation takes account of the fact $\beta_t = 1$ for all $t > \tau$. If the criminal does not return the files his expected payoff is $\Pi = r_\tau$. This takes account of the fact that $\beta_t = 0$ for all $t > \tau$. It is, therefore, optimal to return the file if $\delta > \bar{\delta}$, where $\bar{\delta}$ is as defined in the proof of Proposition 2. □

---

[7]　In comparing our model to other models of belief-based learning we make the following remarks. Young (1993) allows that the individual does not necessarily sample the last n events. This adds a further stochastic element. Other models allow the past to gradually be forgotten and so recent events are given higher weight [38].

Our next result shows that at the opposite limit of a large n it is not optimal to return the files to every victim. In this case the large sample size means that the criminal can take the money and run some of the time without it fundamentally influencing his reputation.

Proposition 4. *For any $\delta < 1$ there exists $\bar{n}$ such that if $n > \bar{n}$ it is not optimal for the criminal to set $g_t = 1$ when $p_t = 1$ for all t.*

Proof. We proceed by contradiction. Suppose that a victim in period τ pays the ransom. Moreover, assume the criminal has a perfect record of returning files and so $\beta_\tau = 1$. If the criminal returns the files then his expected payoff in future periods is at most

$$\Pi_o = r^\star - c + q(r^\star)(r^\star - c)\left(\frac{\delta}{1-\delta}\right) \tag{14}$$

where $r^\star$ is the optimal ransom for $\beta_t = 1$. Fix a particular realization of values $(v_{\tau+1}, v_{\tau+2}, \dots)$. Given this we can determine $(p_{\tau+1}, p_{\tau+2}, \dots)$ and $(g_{\tau+1}, g_{\tau+2}, \dots)$. Suppose that there exists some period $\bar{t}$ at which point n subsequent victims have paid the ransom. That is $\sum_{t=\tau+1}^{\bar{t}} g_t = n$.

Now suppose that the criminal does not return the files in period τ but does continue to return the files to those who pay in all subsequent periods $t > \tau$. There will be at least n periods in which we have beliefs $\beta_t = (n-1)/n$. In this case the victim will only pay if $r_t < v_t(n-1)/n$. We set $r_t = r^\star(n-1)/n$. Consider again the realization of values $(v_{\tau+1}, v_{\tau+2}, \dots, v_{\bar{t}})$. We have fixed things so that $(p_{\tau+1}, p_{\tau+2}, \dots, p_{\bar{t}})$ and $(g_{\tau+1}, g_{\tau+2}, \dots, g_{\bar{t}})$ will remain unchanged. The only difference is the lower ransom. The 'drop' in payoff compared to Π_o is, therefore,

$$\Delta = \sum_{t=\tau+1}^{\bar{t}} \delta^{(t-\tau)} g_t r^\star \left(1 - \frac{n-1}{n}\right) = \sum_{t=\tau+1}^{\bar{t}} \frac{\delta^{(t-\tau)} g_t r^\star}{n}. \tag{15}$$

This can be bounded

$$\Delta \leq \sum_{t=\tau+1}^{\tau+n} \frac{\delta^{(t-\tau)} r^\star}{n} = \frac{r^\star \delta (1 - \delta^n)}{n(1-\delta)}. \tag{16}$$

The 'gain' in payoff compared to Π_o from not returning the files in period τ is c. For sufficiently large n we have $\Delta < c$. Hence the criminal can increase his payoff.

In the proceeding we assumed that there existed some period $\bar{t}$ at which point n subsequent victims have paid the ransom. Potentially one can have a realization of values where this is not the case. Here, however, the drop in payoff from not returning the files is strictly less than Δ (for $\bar{t} = \tau + n$). We observe, therefore, that irrespective of the sequences of valuations the criminal earns a payoff above Π_o if he does not return the files in period τ. □

3.2. Simulation Results on Sample Size

Propositions 3 and 4 formally demonstrate that it may or may not be optimal for the criminal to return files depending on how sensitive victim beliefs are to the criminal's actions. We considered, however, the relative extreme cases of $n = 1$ and n large. To explore what happens in the middle ground we report the results of simulations in which we see how the profit of the criminal is influenced by n. In the simulations we consider the linear distribution of valuations $q(v) = 1 - v$.[8] We track the profits of the criminal over 1000 victims. Throughout we assume that victims start with the 'optimistic' initial belief $\beta_1 = 1$ and then update beliefs according to Equation (12). Each data point is the average of 1000 simulations.

---

[8]　Parameters $a = b = 1$ can be set without loss of generality.

We consider two variants on criminal strategy. Let us first explain our baseline *random model*. This is characterized by an exogenous parameter $w \in [0,1]$. For any period t, if the victim in period t pays the ransom then the criminal returns the files with probability w. Moreover, the criminal sets ransom demand

$$r_t = \frac{\beta_t(1+c)}{2} \tag{17}$$

which corresponds to the optimal ransom demand for beliefs β_t. So, in our baseline model the criminal will, over time, return the files $100w\%$ of the time.

Our baseline model is somewhat biased by the effect which drives Proposition 3. To appreciate why, suppose that there are n instances in a row where the criminal does not return the files. Then beliefs will be zero and no victim will pay the ransom. This means the 'bad reputation' sits forever. To control for this we consider a *lower-bound model* in which the criminal always guarantees that beliefs are at least w. Specifically, for any period t, if the victim in period t pays the ransom then (1) the criminal returns the files with probability w, unless (2) failure to return the files would result in beliefs falling below w, i.e., it would be the case that $\beta_{t+1} < w$, in which case he returns the files, to make sure $\beta_{t+1} > w$.

Consider first the case where the cost of returning files is $c = 0.01$ and the discount factor is $\delta = 1$. Figure 1 plots profit as a function of w for four different values of n. It is readily apparent that profit is increasing in w. In other words the criminals profit is increasing in the likelihood of returning the files and the optimal strategy is to always return files to those that pay the ransom.
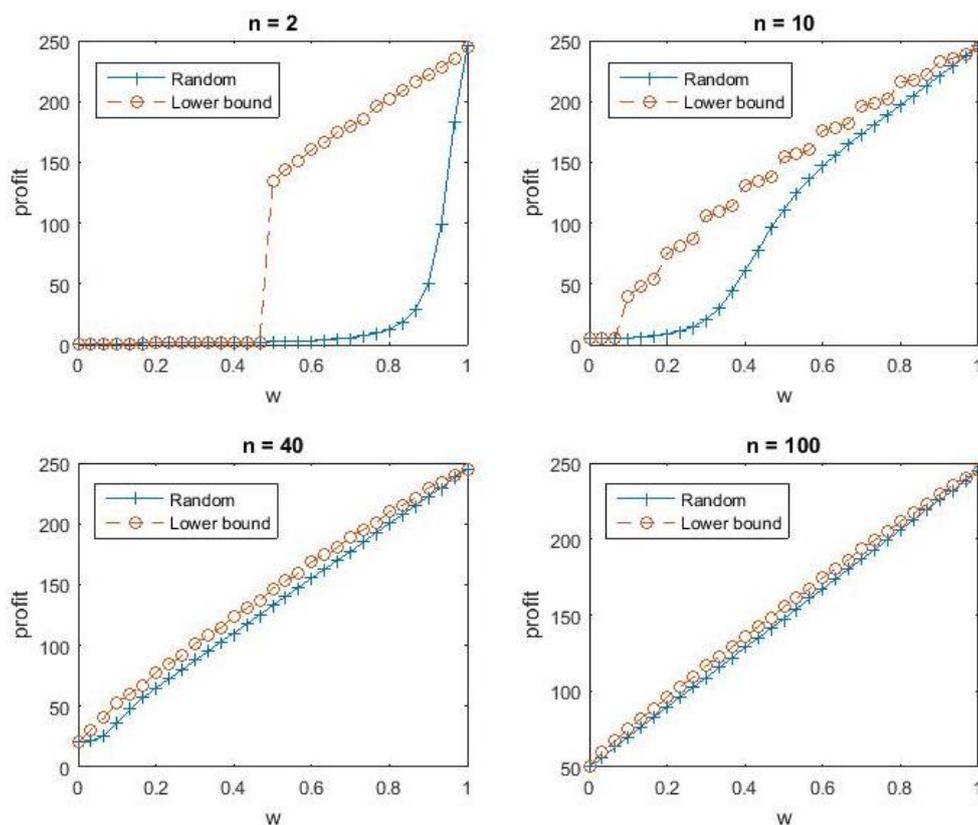
Figure 1. Criminal's profit as a function of w for four different values of n when $c = 0.01$ and $\delta = 1$.

Figure 2 shows what happens if we increase the cost of returning files to $c = 0.1$ and lower the discount factor to $\delta = 0.9$. Here we see that, with the exception of $n = 2$, profits are highest when $w = 0$ and so the criminal's optimal strategy is to not return files. Note, however, that this finding is critically dependent on two things: (1) a discount factor of 0.9 means the criminal is very much focused on the short term, and (2) it takes time for victims to update beliefs. So, the criminal essentially exploits

victims during the time it takes them to collectively learn about his strategy. Clearly, this strategy is not consistent with a long term business model.
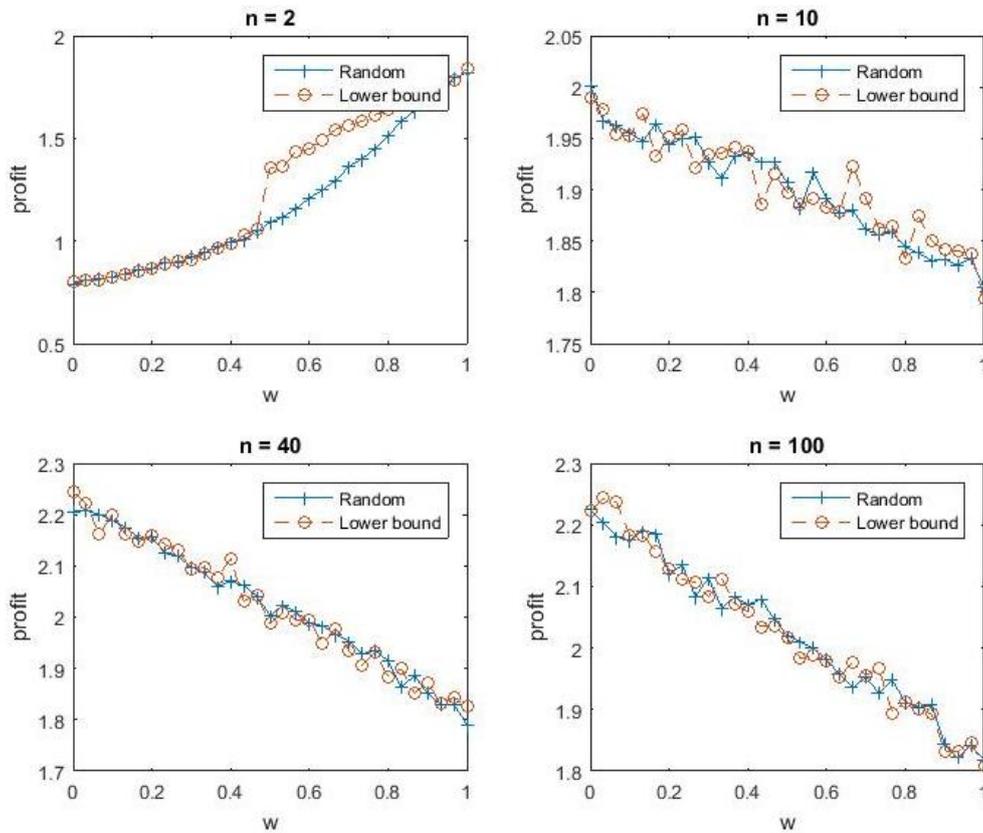
Figure 2. Criminal's profit as a function of w for four different values of n when $c = 0.1$ and $\delta = 0.9$.

To illustrate the criminal's trade-off in choosing between a high and low probability of returning files consider Figure 3. This plots profit over time for three different values of w when $n = 100$ and $c = 0.1$. We can see that the strategy of never returning files $w = 0$ succeeds in the short run because it takes time for victims to update beliefs. But, over the long run the strategy of always returning files $w = 1$ is a clear winner.
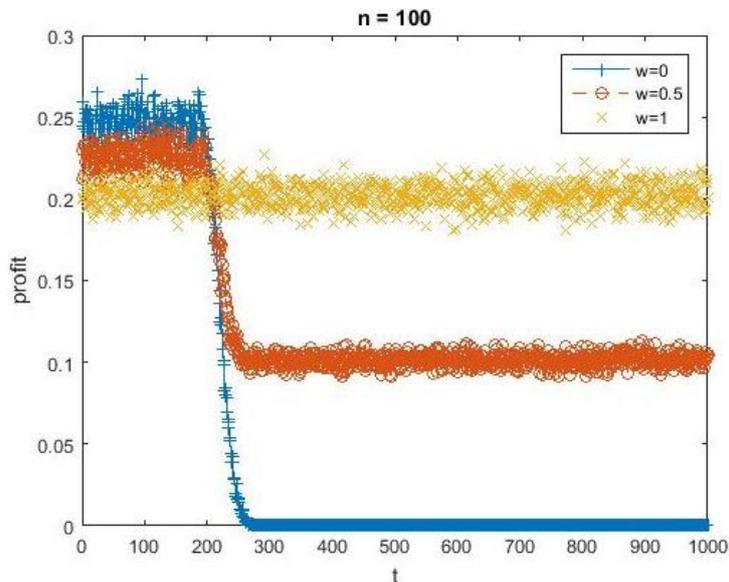
Figure 3. Criminal's profit over time for three different values of w when $c = 0.1$.

You might think the short term advantage of setting $w = 0$ stems from our assumption that victims are initially optimistic (with beliefs $\beta_1 = 1$). This, however, is not the key factor. To illustrate, Figure 4 shows what happens when we set $\beta_1 = 0.5$. As you can see we get similar trade-offs across w. The key factor driving the initial advantage of $w = 0$ is the time it takes for reputation to take hold. The smaller is n the sooner reputations are formed and so the sooner profits are higher with $w = 1$.
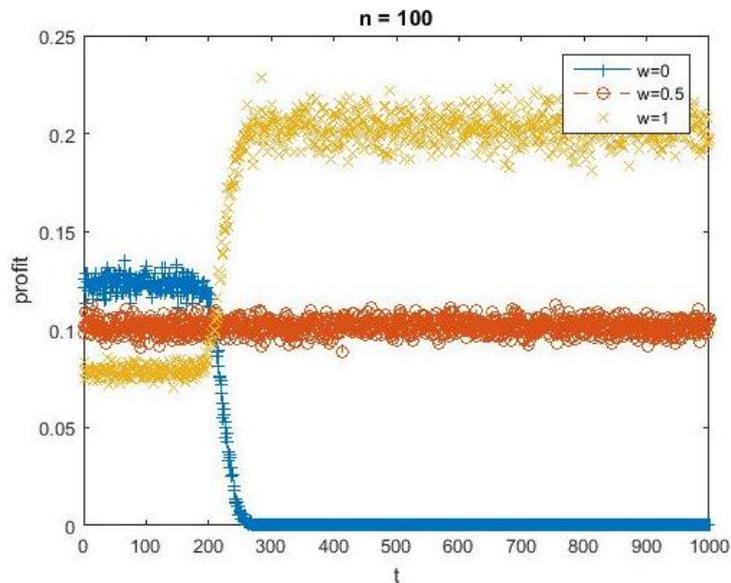
Figure 4. Criminal's profit over time for three different values of w when $c = 0.1$ and $\beta_1 = 0.5$.

In summary, our simulations suggest that a criminal interested in maximizing overall profit will take a long term view and always return files to victims that pay. We observed no gains from a strategy that entails sometimes not returning the files. This suggests that Proposition 4 captures a very specific case. Only the criminal interested in short term gains (low discount factor) would optimally choose to not return files.

4. Conclusions

In this paper we have looked at the incentives for ransomware criminals to return files to victims. We have done so by studying a repeated game in which a criminal interacts with a new victim each period. A victim is assumed to form beliefs based on the past actions of the criminal. In simple settings we were able to obtain explicit results on when it is optimal for the criminal to always return the files and when it is not. We then used simulations to further our understanding around this issue.

To summarize and interpret our results first imagine a world in which there was only one ransomware criminal. In this setting the beliefs of victims will likely be shaped by the actions of the criminal. If so, the criminal will do best to always return the files. This results from two different factors (1) returning the files gives the criminal a 'good reputation', and (2) the criminal can benefit in the long term from that good reputation. Only the most short term of criminals would take the money and run. So, in this stylized setting incentives are clearly to return files.

In reality we know that there are hundreds, if not thousands, of competing ransomware strands. So, next consider a world in which there is a large number of ransomware strands and no distinction between them. Then it is difficult for the criminal to create a personalized reputation. Put another way, the reputation of the criminal is likely to be influenced by factors outside of his control. In this setting the optimal strategy is to not return the files. This, though, cannot succeed in the long run and ultimately would mean that ransomware has no viable future [11]. In the short run, some victims will no doubt keep on paying the ransom in the 'hope' of getting their files back. Ultimately, however, the more news spreads that files are not returned the less trust there will be.

As we discussed in the introduction, evidence suggests that victims recover files around 50% of the time [14]. This would suggest we are somewhere in between the two polar cases discussed above. Some ransomware criminals, we would suggest, are trying to create a good reputation for their 'product' to create a long term viable business model. Other criminals, by contrast, simply want to cash in on short term gains and have no incentive to honour ransom payments. Indeed, we observe fake ransomware that does not encrypt files at all [6]. This melting pot of competing motives could influence the long-term evolution in ransomware in different ways.

If ransomware is to survive in the long term then the criminals behind a particular strand of ransomware, call it Locky, would want to separate their reputation from that of other strands of ransomware. In practical terms this means having a 'brand' that is recognized on search engines or forums that victims may consult after attack. It also means the criminals behind Locky would have an incentive to eliminate any competitors that may muddle their reputation. On the flip side, competitors would have an incentive to imitate Locky so as to free-ride on their good reputation. So, we could have a game of cat and mouse between competing criminal gangs.

From a law enforcement perspective things are somewhat complex. On the one hand if they interfere in the reputation of Locky then this would undermine the business model of Locky. For instance, law enforcement could frame things in a way that lowers victim's beliefs about getting their files back (without, necessarily, explicit deception). This would, hopefully, have the long run effect of removing the criminal's desire to continue with Locky. But, on the flip side, if Locky typically does return files then there is a short term welfare loss of discouraging victims from paying the ransom and getting their files back. So, there is a sense in which some victims have to sacrifice their own interests to undermine the long-term business model.

It is interesting that current advice tends to skim over the likelihood of getting files back. This is possibly because of a recognition that victims who are 'desperate' are going to be tempted to pay the ransom. Other reasons for not paying the ransom are that the victim may be targeted more in the future or may only get some of their files back before being asked for a larger ransom [4]. In our model we have abstracted away from such issues but the basic insights from our paper still hold true. In particular, victims are likely to be influenced by beliefs, whether that be the probability of getting files back or the probability of the ransom increasing. The long term viability of ransomware relies on a good reputation.

Author Contributions: Conceptualization, A.C. and E.C.; methodology, A.C. and E.C.; formal analysis, A.C. and E.C.; writing—review and editing, A.C. and E.C.

Funding: This research was funded by the Engineering and Physical Sciences Research Council (EPSRC) for project EP/P011772/1 on the EconoMical, PsycHologicAl and Societal Impact of RanSomware (EMPHASIS).

Acknowledgments: We thank three reviewers of an earlier version of the paper, as well as the Editor, for their constructive and useful suggestions. We also thank members of the EMPHASIS project for interesting discussions around ransomware.

Conflicts of Interest: The authors have no conflict of interest.

Appendix A

In this appendix we look at how independent and grim-trigger beliefs can be consistent with Bayes Nash equilibrium in a model with reputation.

Consider, first, grim-trigger beliefs. Suppose that there are two types of criminal: (1) A low-cost criminal with $c_L = 0$ and (2) A high-cost criminal with $c_H > 0$. For the low type of criminal there is a weakly dominant strategy to return the files. This could be interpreted as a commitment type who will return the files [25]. A high-cost would prefer to not return the files if this could be done without influencing the beliefs of future victims. Let z denote the prior probability the criminal is low-cost (and $1 - z$ the probability high-cost).

Informally, a Bayes Nash equilibrium in this setting consists of: (a) an optimal strategy for the criminal given his type and the belief formation of victims, (b) an optimal strategy for victims given their beliefs, and (c) victims form beliefs consistent with Bayes rule and the strategy of the criminal.

We introduce some preliminary notation. Given history h_τ in period τ, let L_τ be an indicator variable that takes value 1 if there is period $t < \tau$, where $p_t = 1$ and $g_t = 0$ and takes value 0 otherwise. Let M_τ be an indicator variable that takes value 1 if there is a period $t < \tau$, where $p_t = 1$ and takes value 0 otherwise. In interpretation L_τ records if a victim paid and did not get her files back, and M_τ records if any victim has paid. Also, let $r^\star(\beta)$ denote the optimal ransom to set in a one-shot game with beliefs equal to β.

It is simple to show that for δ sufficiently high the following is a Bayes Nash equilibrium with grim-trigger beliefs:[9] (a.i) Irrespective of type the criminal chooses ransom demand $r^\star(1)$ in all periods. (a.ii) Irrespective of type the criminal honours all ransomware payments, i.e., sets $g_t = 1$ if $p_t = 1$. (b) The victim in period t chooses $p_t = 1$ if $v_t \beta_t > r_t$ and chooses $p_t = 0$ otherwise. (c) The victim in period t has beliefs $\beta_t = 1$ if $L_t = 0$ or $\beta_t = 0$ if $L_t = 1$. Note that on the equilibrium path the criminal always honours ransom payments and so victims expect him to honour ransom payments. Hence the high type criminal maintains a good reputation.

It is simple to show that for δ sufficiently low the following is a Bayes Nash equilibrium with grim-trigger beliefs: (a.i) Irrespective of type the criminal chooses ransom demand $r^\star(\beta_t)$ in all periods. (a.ii) The low type criminal honours all ransomware payments, i.e., sets $g_t = 1$ if $p_t = 1$. The high type criminal never honours ransom payments, i.e., sets $g_t = 0$ if $p_t = 1$. (b) The victim in period t chooses $p_t = 1$ if $v_t \beta_t > r_t$ and chooses $p_t = 0$ otherwise. (c) The victim in period t has beliefs $\beta_t = z$ if $M_t = 0$, beliefs $\beta_t = 0$ if $L_t = 1$ and beliefs $\beta_t = 1$ if $M_t = 1$ and $L_t = 0$ otherwise. Here we have a separating equilibrium in which the low-cost criminal returns the files and a high-cost criminal does not. The first time a victim pays the type of criminal is, therefore, revealed.

Finally, we provide an example with independent beliefs. In the previous example the low-cost type had a weakly dominant strategy to return the files. Here assume that c is small but positive $c > 0$ for the low-cost criminal. Then both would prefer to not return the files if this would not change beliefs. The following is a Bayesian Nash equilibrium: (a.i) Irrespective of type the criminal chooses ransom demand $r^\star$ in all periods (where $r^\star$ can be any positive number). (a.ii) Irrespective of type the criminal does not honour ransomware payments, i.e., sets $g_t = 0$ if $p_t = 1$. (b) The victim in period t chooses $p_t = 1$ if $v_t \beta_t > r_t$ and chooses $p_t = 0$ otherwise. (c) The victim in period t has beliefs $\beta_t = 0$. This equilibrium is characterized by 'pessimistic' beliefs, which make it impossible to generate a positive reputation.

References

1. Hull, G.; John, H.; Arief, B. Ransomware Deployment Methods and Analysis: Views from a Predictive Model and Human Responses. *Crime Sci.* **2019**, *8*, 2. [CrossRef]
2. Kharraz, A.; Robertson, W.; Balzarotti, D.; Bilge, L.; Kirda, E. Cutting the gordian knot: A look under the hood of ransomware attacks. In Proceedings of the International Conference on Detection of Intrusions and Malware, and Vulnerability Assessment, Milan, Italy, 9–10 July 2015; Springer: Cham, Switzerland, 2015; pp. 3–24.
3. Mansfield-Devine, S. Ransomware: Taking businesses hostage. *Netw. Secur.* **2016**, *2016*, 8–17. [CrossRef]
4. Richardson, R.; North, M. Ransomware: Evolution, mitigation and prevention. *Int. Manag. Rev.* **2017**, *13*, 10–21.
5. F-Secure. F-Secure State of Cyber-Security Report 2017. 2017. Available online: https://www.f-secure.com/documents/996508/1030743/cyber-security-report-2017 (accessed on 14 May 2019).

---

9   For a precise definition of δ high see Equation (6).

6. Aurangzeb, S.; Aleem, M.; Iqbal, M.A.; Islam, M.A. Ransomware: A Survey and Trends. *J. Inf. Assur. Secur.* **2017**, *6*, 48–58.

7. Hernandez-Castro, J.; Cartwright, E.; Stepanova, A. Economic Analysis of Ransomware. *arXiv* **2017**, arXiv:1703.06660v1.

8. Trend-Micro. Ransomware: The Truth Behind the Headlines. 2016. Available online: https://www.trendmicro.co.uk/media/misc/ransomware-the-truth-behind-the-headlines.pdf (accessed on 14 May 2019).

9. Huang, D.Y.; Aliapoulios, M.M.; Li, V.G.; Invernizzi, L.; Bursztein, E.; McRoberts, K.; Levin, J.; Levchenko, K.; Snoeren, A.C.; McCoy, D. Tracking ransomware end-to-end. In Proceedings of the 2018 IEEE Symposium on Security and Privacy (SP), San Francisco, CA, USA, 20–24 May 2018; pp. 618–631.

10. Paquet-Clouston, M.; Haslhofer, B.; Dupont, B. Ransomware Payments in the Bitcoin Ecosystem. *arXiv* **2018**, arXiv:1804.04080.

11. Cusack, B.; Ward, G. Points of Failure in the Ransomware Electronic Business Model. In Proceedings of the Twenty-Fourth Americas Conference on Information Systems, New Orleans, LA, USA, 16–18 August 2018.

12. Rashid, F. 4 Reasons Not to Pay Up in a Ransomware Attack. InfoWorld. 2016. Available online: https://www.infoworld.com/article/3043197/4-reasons-not-to-pay-up-in-a-ransomware-attack.html (accessed on 14 May 2019).

13. VioletBlue. Customer Service Matters When It Comes to Ransomware, Engadget. 2016. Available online: https://www.engadget.com/2016/09/09/customer-service-matters-when-it-comes-to-ransomware/ (accessed on 16 May 2019).

14. CyberEdge. Fifth-Annual Cyberthreat Defense Report. 2018. Available online: https://cyber-edge.com/cdr/#about-this-report (accessed on 16 May 2019).

15. Bursztein, E. Unmasking the Ransomware Kingpins. EliE. 2017. Available online: https://elie.net/blog/security/unmasking-the-ransomware-kingpins/ (accessed on 16 May 2019).

16. Palmer, D. The Godfather of Ransomware Returns: Locky Is Back and Sneakier than Ever. ZD Net. 2017. Available online: https://www.zdnet.com/article/the-godfather-of-ransomware-returns-locky-is-back-and-sneakier-than-ever/ (accessed on 16 May 2019).

17. Caulfield, T.; Ioannidis, C.; Pym, D. Dynamic Pricing for Ransomware. 2019. Available online: http://www0.cs.ucl.ac.uk/staff/D.Pym/ransomware-dynamic.pdf (accessed on 14 May 2019).

18. Caporusso, N.; Chea, S.; Abukhaled, R. A Game-Theoretical Model of Ransomware. In Proceedings of the International Conference on Applied Human Factors and Ergonomics, Orlando, FL, USA, 27–31 July 2018; Springer: Cham, Switzerland, 2018; pp. 69–78.

19. Laszka, A.; Farhang, S.; Grossklags, J. On the Economics of Ransomware. In Proceedings of the International Conference on Decision and Game Theory for Security, Vienna, Austria, 23–25 October 2017; Springer: Cham, Switzerland, 2017; pp. 397–417.

20. Cartwright, A.; Cartwright, E.; Hernandez-Castro, H. To pay or not: Game theoretic models of ransomware. *J. Cybersecur.* **2019**, forthcoming.

21. Brandt, P.T.; George, J.; Sandler, T. Why concessions should not be made to terrorist kidnappers. *Eur. J. Political Econ.* **2016**, *44*, 41–52. [CrossRef]

22. Lapan, H.E.; Sandler, T. To bargain or not to bargain: That is the question. *Am. Econ. Rev.* **1988**, *78*, 16–21.

23. Selten, R. A simple game model of kidnapping. In *Mathematical Economics and Game Theory*; Springer: Berlin/Heidelberg, Germany, 1977; pp. 139–155.

24. Zarifis, A.; Cheng, X. The Impact of Extended Global Ransomware Attacks on Trust: How the Attacker's Competence and Institutional Trust Influence the Decision to Pay. In Proceedings of the Twenty-Fourth Americas Conference on Information Systems, New Orleans, LA, USA, 16–18 August 2018.

25. Cripps, M.W.; Mailath, G.J.; Samuelson, L. Imperfect monitoring and impermanent reputations. *Econometrica* **2004**, *72*, 407–432. [CrossRef]

26. Fudenberg, D.; Levine, D.K. Maintaining a Reputation when Strategies are Imperfectly. *Rev. Econ. Stud.* **1992**, *59*, 561–579. [CrossRef]

27. Fudenberg, D.; Kreps, D.M.; Maskin, E.S. Repeated games with long-run and short-run players. *Rev. Econ. Stud.* **1990**, *57*, 555–573. [CrossRef]

28. Celetani, M.; Fudenberg, D.; Levine, D.K.; Pesendorfer, W. Maintaining a reputation against a long-lived opponent. *Econometrica* **1996**, *64*, 691–704. [CrossRef]

29. Kreps, D.M.; Milgrom, P.; Roberts, J.; Wilson, R. Rational cooperation in the finitely repeated prisoners' dilemma. *J. Econ. Theory* **1982**, *27*, 245–252. [CrossRef]

30. Kreps, D.M.; Wilson, R. Reputation and imperfect information. *J. Econ. Theory* **1982**, *27*, 253–279. [CrossRef]

31. Milgrom, P.; Roberts, J. Predation, reputation, and entry deterrence. *J. Econ. Theory* **1982**, *27*, 280–312. [CrossRef]

32. Weinstein, J.; Yildiz, M. Reputation without commitment in finitely repeated games. *Theor. Econ.* **2016**, *11*, 157–185. [CrossRef]

33. Fudenberg, D.; Levine, D.K. *The Theory of Learning in Games*; MIT Press: Cambridge, MA, USA, 1998.

34. Young, H.P. *Individual Strategy and Social Structure: An Evolutionary Theory of Institutions*; Princeton University Press: Princeton, NJ, USA, 2001.

35. Selten, R. Evolution, learning, and economic behavior. *Games Econ. Behav.* **1991**, *3*, 3–24. [CrossRef]

36. Young, H.P. The evolution of conventions. *Econometrica* **1993**, *61*, 57–84. [CrossRef]

37. D'Acremont, M.; Schultz, W.; Bossaerts, P. The human brain encodes event frequencies while forming subjective beliefs. *J. Neurosci.* **2013**, *33*, 10887–10897. [CrossRef]

38. Feltovich, N. Reinforcement-based vs. Belief-based Learning Models in Experimental Asymmetric-information Games. *Econometrica* **2000**, *68*, 605–641. [CrossRef]