How Cyber-Insurance Influences the Ransomware Payment Decision: Theory and Evidence[⋆⋆⋆]

Anna Cartwright[1][0000-0003-1965-842X],
Edward Cartwright[2][0000-0003-0194-9368],
Jamie MacColl[3],
Gareth Mott[4][0000-0002-8788-769X],
Sarah Turner[5][0000-0003-1246-1528],
James Sullivan[3], and
Jason R.C. Nurse[5][0000-0003-4118-1680]

[1] Oxford Brookes Business School, Oxford Brookes University, Oxford, UK.
[2] Department of Accounting, Finance and Economics, De Montfort University, Leicester, UK. edward.cartwright@dmu.ac.uk
[3] Royal United Services Institute, Whitehall, London, UK.
[4] School of Politics and International Relations, University of Kent, Canterbury, UK.
[5] School of Computing, University of Kent, Canterbury, UK.

Abstract. In this paper we analyse how cyber-insurance influences the cost-benefit decision making process of a ransomware victim. Specifically, we ask whether organizations with cyber-insurance are more likely to pay a ransom than non-insureds. We propose a game-theoretic framework with which to categorize and distinguish different channels through which insurance may influence victim decision making. This allows us to identify ways in which insurance may incentivize or disincentivize payment of the ransom. Our framework is informed by data from semi-structured interviews with 65 professionals with expertise in cyber-insurance, cyber-security and/or ransomware, as well as data from the UK Cyber Security Breaches Survey. We find that perceptions are very divided on whether victims with insurance are more (or less) likely to pay a ransom. Our model can reconcile these views once we take into account context specifics, such as the severity of the attack as measured by business interruption and restoration and/or the exfiltration of sensitive data.

Keywords: Ransomware · Insurance · Cyber-security · Double extortion · Moral hazard · Negotiation

[⋆] On behalf of all authors, the corresponding author, Edward Cartwright, states that there is no conflict of interest.

[⋆⋆] This project was supported by the UK National Cyber Security Centre (NCSC) project ‘Ransomware: The Role of Cyber Insurance’.

1 Introduction

Crypto-ransomware is a form of malware that encrypts the files on a device so that the victim can no longer access their files or systems. The criminals then demand a ransom for the key to decrypt the files in order to restore access (11; 24; 33; 36). Loss of access to files poses a serious business interruption risk to organisations, even if they have a back-up (10). In recent years the ransomware threat has evolved further with a trend towards double extortion in which the criminals not only disrupt access to files but also threaten the leaking of sensitive data (22; 44). While it was relatively unknown 10 years ago, ransomware has quickly evolved to become one of the most common forms of malware (2; 25). It has proved lucrative for criminal gangs (32; 34) and highly disruptive for society and business. In particular, ransomware attacks have caused disruption to public services, including health care and education, and business, including critical infrastructure (e.g. 35; 36; 41; 42).

Against the backdrop of a worsening ransomware threat the role of cyber-insurance has come into the spotlight. In principle, cyber-insurance provides a market based solution to pool risk and limit organizations' exposure to the severe losses that can result from ransomware. The role of cyber-insurance in mitigating the ransomware threat has, though, been a topic of debate, with accusations that insurance companies have fuelled the payment of ransoms (e.g. 4; 12; 38; 46). In short, it is argued that paying the ransom becomes 'easier with insurance' and so organisations with insurance are more likely to pay a ransom than those without. The extent, however, to which this accusation stands up to close scrutiny is unclear. Moreover, even if there may have been some truth to the claim in the past, the cyber-insurance market is fast evolving and the balance may have changed as the insurance industry adapts (21). Furthermore, regulation may provide a viable way to address any socially inefficient incentives that result from insurance (27).

In this paper we provide a novel study of the influence cyber-insurance is having on the cost-benefit trade-off of paying a ransom. We first propose a game-theoretic framework with which to conceptualise the channels through which the presence of insurance can change the incentives of a victim organisation to pay a ransom. Consistent with a game-theoretic approach we distinguish four types of influence: (i) It changes the victim's payoff, or cost-benefit trade-off, from different actions, e.g. by providing financial cover for business interruption. (ii) Changing the set of actions available to the victim, e.g. access to skilled negotiators capable of brokering ransom payments. (iii) Changes to the belief updating of the victim, e.g. the insurer providing information on the likelihood of being able to decrypt files if a ransom is paid. (iv) Changes to the belief updating of the criminal, e.g. the criminal may become aware the victim has increased access to funds to pay a ransom. A key distinguishing feature of our framework is that it is dynamic. This allows us to capture ways in which insurance alters the timing of decisions. Indeed, a key benefit of insurance is that it can provide victims with more time to explore their options for system recovery.

We next interpret our framework based on semi-structured interviews with 65 professionals with expertise in cyber-insurance, cyber-security and/or ransomware. The interviews probed a range of issues around cyber-insurance and ransomware. Here we focus specifically on the interviewees views concerning the role of insurance in influencing the ransom pay decision. We identify and classify a number of potential channels through which insurance can have a positive or negative influence on the pay decision. We then look to evaluate the net effect of these influences. In doing so, we conjecture that in the case of a highly severe attack insureds may, on balance, be weakly more likely to pay the ransom. A highly severe attack one would be one with either large business interruption and/or the exfiltration of significant sensitive data. Insurance may incentivize payment in this case because it provides liquidity to pay the ransom and may provide expertise, e.g. in ransom negotiation, to facilitate payment. For a less severe attack we cannot identify a reason why insureds would be any more likely to pay a ransom than non-insureds.

In our main analysis we do not model the overall impact of insurance on ransomware. Our framework provides, however, novel insight on this overall impact. Specifically, cyber-insurers would argue that their clients are likely to have robust security practices in place. While the evidence for this claim is somewhat mixed there is no doubt that underwriting restrictions can play a positive role in selecting organisations with more cyber controls (47; 23; 37; 30; 29; 45; 43). This would mean that insureds will typically fall in impact scenarios where the ransom is less likely to be paid. To evaluate this argument we analysed data from the UK Government's Cyber Security Breaches Survey (2018-2021) giving us insight on 5878 businesses. Our analysis shows that businesses with cyber-insurance do indeed show significantly higher levels of cyber-security controls. We find no discernible change in insurance on the reported impact of a ransomware attack, but this could reflect the small number of businesses in the survey who experience a ransomware attack. If insureds have higher levels of cyber-security controls then insurance could have a net positive impact on reducing ransom payments.

In making the case that insurance may have a net positive impact we recognise that cyber-security practices are still lax in many organisations, including those with cyber-insurance. For instance, analysis of the Cyber Security Breaches Survey (see Section 5) shows that one in ten large businesses with cyber-insurance do not meet the UK's 'minimum level' Cyber Essentials controls. These organisations could face highly severe attacks where insurance may incentivize payment of the ransom. The overall picture is, therefore, complex and nuanced. This is reflected in our interviews where we found very divergent perceptions of whether insureds are more or less likely to pay a ransom than non-insureds. If, however, cyber-security practice is improving over time then it is reasonable to argue we are transitioning from a setting where insureds may have paid a ransom to a setting where insureds are less likely to pay than non-insureds. In a future where businesses take appropriate cyber-security actions our model suggests insurance could significantly lessen the incentives to pay a ransom.

We proceed as follows. In Section 2 we set out the methodological approach used in the paper. In Section 3 we propose a game-theoretic framework to conceptualise the ransomware pay decision. In Section 4 we report and discuss the interview data. In Section 5 we provide a unified interpretation of the impact of insurance on the ransom pay decision. In Section 6 we conclude.

2 Research methodology and data collection

The victim of a ransomware attack faces a basic economic cost-benefit trade-off as to whether to pay the ransom. In short, the organisation needs to decide whether the expected net benefit of paying the ransom exceeds the expected net cost. A complex range of factors will impact on this trade-off, ranging from tangible financial losses, such as the ransom payment, to intangible factors, such as the 'psychological cost' of rewarding or fuelling criminality. There is an emerging game-theoretic literature on the ransomware pay decision that identifies conditions under which a 'rational' victim would optimally choose to pay a ransom (e.g. 7; 9; 14; 17; 26; 28).

Relatively few game theoretic studies have considered the impact of insurance on the ransom pay decision. This, therefore, is an under-developed field of research. The few studies that have addressed the issue find that insurance is predicted to *increase* the likelihood of victims paying the ransom because, for example, it reduces liquidity constraints on victims (5), can lead to moral hazard (3; 49) or does not reduce the incentive to pay a ransom (17). These models, however, in order to provide a tractable, solvable game have necessarily focused on a handful of very specific channels through which insurance may impact on the ransom pay decision. Informative though these models are, they cannot capture the many varied channels through with insurance can influence the pay decision, and so may, for instance, fail to account for the positive benefits that insurance can provide post-breach. In Section 3 we take a complementary approach, and propose a general game-theoretic framework with which to conceptualise the ransom pay decision at a higher level of generality. While this approach does not provide a 'parameterized output' we show that it provides an informative framework with which to distinguish and characterise the influence of insurance. We subsequently look to refine our framework along three dimensions: (i) the potential channels through which insurance can impact on the benefit-cost calculation, (ii) the relative importance or weight of these channels, and (iii) the overall net impact of insurance on the likelihood of paying a ransom. Our approach is informed by two distinct sources of data which we detail in turn.

Semi-structured interviews were performed between September 2021 and February 2022 by the research team. Interviewees were chosen for their expertise in either the selling or purchase of cyber insurance products, the policy aspects around such insurance, or because of involvement in the process of ransomware and other cyber security or crisis management services. Subjects were recruited from existing contacts, recommendations from law enforcement and snowball selection. The interviews were conducted under a principle of full anonymity but we can say that the interviewees had wide and extensive experience of cyber insurance and/or ransomware in roles including head of incident response, chief claims officer, head of underwriting, chief information security officer (CISO), and chief executive officer (CEO). There was a split between those that specialise in working with large-sized corporates, and those working with small and medium enterprises (SMEs) in terms of under-writers, brokers and claims. The interviewees were geographically spread across the UK, continental Europe, Bermuda and the US.[6] Interviews typically took between 40 minutes and an hour and were conducted online. In total 65 people were interviewed across 52 interviews (some interviews contained multiple interviewees). Table 1 provides a breakdown of interviewees by category and role.

The interviews were designed to cover a wide range of issues around ransomware and cyber insurance, including the state of the cyber-insurance market, pre-incident support, policy exclusions, post-incident response and recommendations for policy intervention. The full set of interview questions is set out in the Appendix. In this paper we focus specifically on the question of whether insurance influences the decision to pay a ransom. Specifically, we focus on two lines of questioning used during the interviews: (a) Can you talk us through what happens during a ransomware attack? (b) Do you think insureds are more likely to pay a ransom than uninsureds? As we document below, not all interviewees expressed a firm opinion on these questions, but most did. Each interview was transcribed in written format. For the analysis reported in this paper two members of the research team independently coded the transcripts with a specific focus on the ransom pay decision. The coding aimed to: (1) identify all potential channels through which insurance may alter the costs or benefits of paying a ransom, and (2) identify the interviewees' views on whether the net effect of insurance is to make the victim more or less likely to pay a ransom. Codes were categorized to fit with our game-theoretic model.

The second distinct source of data we analysed is the UK Government's annual Cyber Security Breaches Survey from 2018-2021.[7] Individual level data from the years 2018-2021 was obtained from the UK Data Service and combined to give a dataset of 5,878 businesses.[8] The Survey covers a wide range of questions. Here we focus on a specific subset of relevant questions, including whether or not the business has cyber specific insurance, whether they were subject to a ransomware attack, the impact of that attack, and the cyber-security measures undertaken by the business as measured against the UK's National Cyber Secu-

[6] For instance, in terms of underwriters: Four specialise in the SME market, the rest in the large-sized corporate market in the US and UK. Brokers: One specialist in the UK SME market. One specialist in UK public sector and critical services. The rest specialised in working with mid and large sized corporates in the London market. Claims: Two specialise in the SME market in the UK; one specialises in large corporates in the US market.

[7] See https://www.gov.uk/government/statistics/cyber-security-breaches-survey-2021 for further details of the 2021 survey.

[8] The same business may appear over time but this is not identifiable in the data. We do not consider data for charities and education establishments.

Category	Subcategory/role	Count
Insurance Industry	Cyber insurance underwriter	10
	Cyber insurance broker	5
	Cyber insurance claims	3
	Cyber insurance executive	3
	Insurance industry association	3
	Cyber risk management services	2
	Cyber risk analytics	2
	Cyber reinsurance executive	1
	Cyber reinsurance underwriter	1
Cyber security	Digital Forensics and Incident Response	9
	Cyber security consultant	4
	Cyber threat intelligence manager	3
	Public policy	1
	Ransomware recovery	1
Purchasing organisations	Technology	2
	Local government	2
	Financial services	1
	Transport	1
	Defence	1
Government	Cyber policy	3
	Incident Management	1
Professional services	Breach counsel	2
	Insurance lawyer	1
Law enforcement	International law enforcement agency	1
	National law enforcement agency	1
Academia	Academic	1
Total		65

Table 1. Interview breakdown by category

rity Centre 10 Steps to Cyber Security.[9] Full information on the survey questions is also available from the UK Data Service. A summary breakdown of the number of observations by year and size of business is provided in Table 2 along with the proportion of respondents that had a cyber specific insurance policy. You can see that the Survey has a good coverage of micro, small, medium and large organisations. A growing number of papers are using the Cyber Security Breaches Survey data to investigate the cyber landscape in the UK (e.g 16; 20) but there is no study to our knowledge focused on ransomware or cyber-insurance.

3 Theoretical Framework to Model Ransom Payments

In our framework we focus on the decision making process of a victim of ransomware. We, thus, take as given that an organisation, called the *victim*, is

[9] Information on the 10 Steps can be found at https://www.ncsc.gov.uk/collection/10-steps

Size	2018		2019		2020		2021	
Large (250+ employees)	252	24.6%	214	34.1%	221	20.8%	190	20.5%
Medium (50-249 employees)	263	20.5%	301	35.9%	223	18.6%	199	18.6%
Small (10-49 employees)	349	12.9%	330	16.7%	286	15.4%	255	13.7%
Micro (0-9 employees)	655	6.8%	770	9.2%	644	2.6%	726	5.0%
Total	1519		1615		1374		1370	

Table 2. Summary of observations from the UK Cyber Security Breaches Survey. Number of businesses by size and year and proportion of businesses with a cyber specific insurance policy.

subject to a ‘penetrative’ ransomware attack by a malicious actor, called the *criminal*. The victim must ultimately decide whether to pay a ransom. Our main objective is to evaluate the impact of insurance on this decision. We can do this by comparing a victim without insurance to one with insurance. Naturally, the average organisation with insurance may look very different to the average organisation without insurance, because the choice to seek insurance, and the decision of the insurer to provide insurance, will depend on the characteristics of the organisation. We, thus, recognise that victims with insurance may differ from those without. An important limitation of existing game theoretic models of ransomware is that they are primarily static. The decision making process in a ransomware attack is, however, dynamic (see e.g. (48)) and one key benefit of insurance, as we shall discuss shortly, is that it can afford victims more time to make decisions. We, thus, propose a dynamic framework that can capture salient aspects of the decision making process over time. We assume that time runs in discrete periods $t = 0, 1, 2, 3, .., T$ beginning when the ransomware is first known to the victim and ending at some period T sufficiently far into the future.

We are naturally in a setting with high levels of imperfect and incomplete information. In particular, the victim may not know for sure the past actions of the criminal and vice-versa. For instance, the victim may not know if the criminal has already leaked data on the web. Similarly, the criminal may not know actions the victim has taken to recover from back-ups. The victim and criminal will also not be fully informed of the losses and gains of performing actions. For instance, the action to pay a ransom has an uncertain outcome, not only in terms of whether the criminals will make decryption keys available, but whether decryption will work. Moreover, the criminal may not know whether the victim has insurance, or may not know the terms of that insurance. The appropriate game theoretic framework to model a setting with incomplete information is a Bayesian game with Harsanyi types (15; 14). A Bayesian game can be briefly described as follows: Before time 0 ‘nature’ determines the type of the players. Each player is assumed to know their own type but not that of others. Players have initial beliefs about the type of others, which are then updated as time progresses, when they observe actions or receive additional information. To provide a specific example, there may be a type of criminal where the decryption key works and a type where the decryption does not work. Similarly, there may be a type of victim who is able to successfully decrypt and a type that is not. This can capture the inherent uncertainty that both criminal and victim will have on whether encrypted files can be successfully decrypted. Actions, e.g. trying to decrypt the files, allow the victim (and criminal) to update their beliefs.

In light of the preceding points we suggest that the interaction between the victim and criminal can be appropriately modelled as a dynamic Bayesian game. Informally, a dynamic Bayesian game is characterised by (a) a set of players, (b) the order of moves, (c) the players' payoffs as a function of the moves chosen, (d) the set of choices available to a player at each move, (e) the information each player has available when making his or her choice, and (f) the probability distribution over acts of nature (15). We define (a-f) in turn. In our setting the two key players are the victim and the criminal. Alongside this there may be a range of other players, such as incident response, insurance company, clients of the victim, and law enforcement. Given that we are interested on the decision of the victim to pay a ransom we primarily focus on how other players may influence the incentives of the victim.

We assume that in each period the victim and all other players have the opportunity to move. We denote by A_t the set of actions that can be performed in period t by the victim. Given that the potential actions a victim could take are vast and impossible to summarize, we are deliberately agnostic about what those actions are. They can include, for instance: paying the ransom, negotiating the ransom amount, performing data clean up, diagnostic checks, restoring from back-ups, migrating to the cloud, reporting to law enforcement, informing customers and suppliers etc. Let $a_t \subset A_t$ denote the actions performed in period t by the victim. Similarly, other players, including the criminal, can also perform actions. For instance, the criminal can change the ransom demand, change the deadline for ransom payment, leak or publish data on the web, launch a new attack against the victim etc. Let r_t denote the set of actions performed in period t by the other players.

At any period t there will be a history of past actions as given by $h_t = \{a_0, ..., a_t, r_0, ..., r_t\}$. The set of actions A_t available to the victim, and to other players, will be highly path dependent. For instance, the action to decrypt files would naturally follow the act of paying a ransom. Similarly, a criminal's success in launching a new attack will depend on how well the victim has cleaned up their system and instigated new security measures. The set of available actions A_t will, thus, depend on history h_{t-1}. At any time t the victim (and other players) can be assumed to have beliefs about: (i) the type of the other players, (ii) the past actions of the other players. Let $\beta(h_t)$ denote the beliefs of the victim.

Over time the victim will incur losses as a result of the attack. Denote by $l_t(h_t)$ the losses incurred in period t and $L_t(h_t) = \sum_{\tau=0}^{t} l_\tau(h_\tau)$ be the cumulative loss up to period t. Typically we would expect $l_t > 0$, meaning a net loss, but we do not rule out the possibility $l_t < 0$, meaning a net gain because, for example, new ways of operating have increased productivity. Actions will have costs for the victim. For instance, performing diagnostic tests, servicing payment of a ransom, or migrating to the cloud will incur capital and labour costs for the victim. We denote by $c(h_t)$ the costs incurred by the victim in period t, and $C(h_t)$ the cumulative cost over time.[10]

The victim faces a complex, forward looking problem in which their objective is to choose a set of actions in period t so as to minimize the net loss from the attack given their beliefs, $\min E[f(L + C|\beta)]$, where f is an objective function that takes into account attitudes to risk (i.e. risk averse or risk seeking) and time preference (i.e. discount rate). The criminal, and other players such as insurers, incident response and law enforcement, will similarly be trying to maximize their gain net of costs. There are various models to help organizations consider the various cost trade-offs they face (6). The issue we focus on in this paper is whether the presence of insurance influences the cost-benefit trade-off of paying the ransom. Before we directly tackle this question we briefly reflect on the decision making process.

3.1 Pay or delay

Previous game-theoretic models of the ransom decision has focused on a binary pay or not pay decision (e.g. 3; 9). A dynamic setting makes clear, however, that the decision to be made by the victim is not whether or not to pay but whether *to pay now or delay a decision until later*. This crucially allows the victim to pursue multiple, simultaneous actions that collectively allow an updating of information and beliefs, β, and, thereby, feed into an overall strategy to minimize losses. A particular example is to pursue negotiations with the criminal while simultaneously trying to recover from back up and, more generally, assess the extent of the damage done by the attack (1; 28). The potential to pursue ransom negotiations while simultaneously exploring alternatives was a theme in the interviews. For instance, an underwriter said:

> We'll be negotiating in parallel whilst we try to recover from back up and if it's just not possible to recover from back up or in a timely way that saves the business, then obviously we'll pursue the negotiation more aggressively.

Similarly, a risk officer at a purchasing organisation said:

> We'd do everything, we back all of the horses, and we don't have any precedent about whether we pay or don't pay, everything is dealt with on a case by case, but, certainly, we would, we would seek to recover a service, and data, as well as pursue a negotiation if that was sort of the right business choice to make. And we would assume that neither one of those things, neither thing, would succeed, so we'd want to have, sort of plan Bs and Cs.

[10] The criminal and other players, such as insurers and law enforcement, will also incur financial gains and losses, e.g. through the sale of breached data. We abstract from that here to focus on the victim.

The optimal strategy in the ransom response involves, therefore, pursuing actions that will inform and update beliefs. This has not been studied in the prior theoretical literature but, as we shall see, is potentially crucial to the benefits of insurance.

Another important point to highlight, as recognised in our model, is the number of actors that may inform the victim's decision making process. In a game theoretic model it is standard to consider *a victim*, but the victim's decision will typically be determined, as time progresses, by a very wide range of actors, including the Board or Trustees, IT providers, incident response, legal teams, law enforcement, ransom negotiators, public relations teams, customers or clients, and the insurers (1). This process may be highly complex, particularly if actors have non-aligned objectives (48). Ultimately, though, it is for the victim organisation, whether that be the CEO, board of directors etc., to decide to pay a ransom. Other actors play a role of informing the victim organisation. For instance, a head of cyber insurance said:

> we always are respectful of the boundary lines, and one big boundary line is we can put you in touch with the best experts, we can tell you what we typically see, we can tell you what we've seen from our experience. [If] you're going to pay a ransom, that is your decision as a business. We are not going to say pay, we are not going to say don't pay, that's not our call, it's your business it's your call.

One key theme running through the interviews was the notion that paying the ransom 'is the last resort'. Clearly no organisation wants to pay a ransom and so the notion that paying should be a last resort is very natural. It is crucial, though, to pick apart what that means in practice. Our dynamic, Bayesian framework allows a characterization of the process. Simultaneous actions can be taken to update beliefs and inform on the actions that will minimize the expected net loss from the attack. Only if all other options (e.g. recover from back up, prove relatively costly, and the ransom is possible, e.g. sanctions checks are passed, and the ransom is expected to minimize loss, e.g. the ransom demand has been negotiated down and criminals will offer decryption keys) is the ransom paid. This is shown schematically in Figure 1. We outline in a highly stylized way one possible path on the game tree of an equilibrium strategy. The strategy involves a combination of actions to mitigate loss, as well as updating beliefs. Actions, including whether to pay the ransom, are based on updated beliefs. The ransom is only paid if there comes a point in time where it seems payment of the the ransom will minimize expected net loss.

3.2 Modelling the influence of insurance

We turn now to the role of insurance in influencing the ransom pay decision. In principle, insurers can influence all aspects of the dynamic decision making process illustrated in Figure 1. To capture this we distinguish four game theoretic channels through which insurance can influence the cost-benefit decision of the

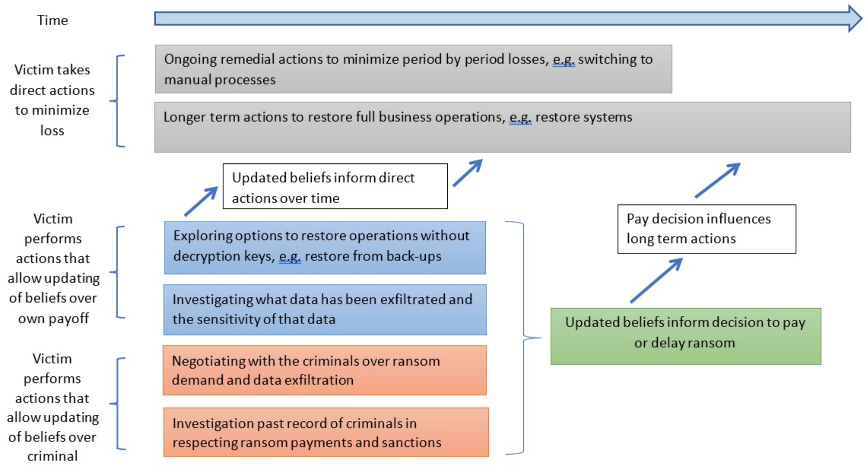
Time
Victim takes direct actions to minimize loss
Ongoing remedial actions to minimize period by period losses, e.g. switching to manual processes
Longer term actions to restore full business operations, e.g. restore systems
Updated beliefs inform direct actions over time
Pay decision influences long term actions
Victim performs actions that allow updating of beliefs over own payoff
Exploring options to restore operations without decryption keys, e.g. restore from back-ups
Investigating what data has been exfiltrated and the sensitivity of that data
Updated beliefs inform decision to pay or delay ransom
Victim performs actions that allow updating of beliefs over criminal
Negotiating with the criminals over ransom demand and data exfiltration
Investigation past record of criminals in respecting ransom payments and sanctions

Fig. 1. Schematic representation of the dynamic Bayesian game and ransom pay decision

victim. These are: (i) Impact on the victims payoff, or cost-benefit trade-off as captured by loss function l_t, and cost function, c_t. (ii) Revise the action set A_t available to the victim in some period t. (iii) Refine the way the victim updates beliefs β about the type of the criminal and moves of nature. (iv) Refine the way the criminal (and potentially other players) update beliefs about the type of the victim (or other players). One could also consider changes to the payoff or action set of the criminal (and other players) but we suggest these are of lesser importance.

Based on the existing literature and the data from the interviews we identified and characterised a main set of potential influences insurance can have on the pay decision. These are summarized in Table 3. For each influence we indicate whether the likely effect of insurance, for most victims, is to lessen the incentives to pay, increase the incentives to pay, or have an ambiguous influence. In interpretation, we are looking at how insurance may influence available actions, A_t, payoff, l_t, c_t, or updated beliefs, β, of the victim in a way that changes the cost-benefit trade-off of paying the ransom. We highlight that our objective at this stage is to identify a set of *potential* influences. We will use the interview data to identify the main influences (and, where relevant, identify the direction of influence). We briefly summarize these influences in turn.

A basic objective of insurance is clearly to provide financial support to victims. (37) identified the top 10 most common covered losses in cyber-insurance. This top 10 covers financial recompense for business interruption, reputation loss, extortion etc. Such coverage has the effect of lowering the effective loss $l(t)$ over time, because financial losses are transferred on to the insurer. Also, the insurer will cover (or partly cover) costs, such as incident response or payment of the ransom, and thus lower the costs $c(t)$ incurred by the victim. This fundamentally changes the payoff function of the insured, but can change incentives in different ways. For instance, access to funds to pay a ransom creates a form of moral hazard and can, thus, increase the incentive to pay (5). By contrast, finance that cover losses and access to incident response should, in general, lower the incentive to pay a ransom because it means the victim has less to gain from engaging with the criminals. These effects can be dynamic in changing the flow of losses and costs. For example, paying the ransom has an up-front cost, involves subsequent costs to perform the decryption, and will only realistically lower the losses from business interruption after a delay while decryption takes place.

In terms of the action set available to the victim, insurance means a range of options and services become available (48). The value of these services are context specific. For instance, a large organisation may have ample in-house expertise in public relations and legal counsel to manage an attack, and ample financial reserves to, say, pay a ransom. A small business, by contrast, may have very limited ability to deal with an attack in the absence of insurance. Access to post-breach services that support a victim's recovery naturally lowers the incentive to pay a ransom. For instance, they may allow recovery without need to pay the ransom. Access to services that facilitate payment of the ransom, however, naturally increase the incentive to pay a ransom. For instance, the greater access to funds and liquidity may mean payment is feasible where it would not have been otherwise. The influence of access to legal advice is ambiguous because it will depend on whether that legal advice is to not pay (and hence an insured is less likely to pay a ransom) or that payment can be made (and hence an insured is more likely to pay). One important element of insurance is the extra time it may afford victims. This comes in, at least, three forms: First, the extra resource and expertise that insurers can provide 'buys time' in the sense that more can be done within a fixed period of time to address a breach. For instance, to more quickly identifying how to restore systems or recover data. Second, insurers bring additional expertise in negotiating with criminals and pushing back deadlines. Third, that insurers cover business interruption losses again, 'buys time' for the victim to work out alternative options. Increased time should lower the incentive to pay a ransom because of additional options it provides.

The influence of insurance on the information set of the victim can operate on many levels and, as noted in Table 3 has highly ambiguous effects. To illustrate, consider the fact that insureds will likely have access to more reliable information on the criminal's past record. If the criminals have a strong reputation of providing decryption keys and not publishing exfiltrated data when a ransom is paid, then such information would incentivize payment of the ransom because it means the ransom payment could be seen as 'good value for money' (8). If, by contrast, the criminals have a poor reputation of providing decryption keys or honouring promises then there is little incentive to pay the ransom. Similarly, access to forensics and incident response can create differing incentives depending on whether the news, from the perspective of the victim, is good or bad.

Channel	Possible influence of insurance	Incentive
Payoff function	Financial recompense for business interruption	Lower
	Financial recompense for reputation loss and/or damages	Lower
	Resource for incident response and forensic investigation	Lower
	Recompense for ransom	Increase
Action set of victim	Access to public relations services	Lower
	Access to data/system restoration and forensics	Lower
	Increased financial liquidity to spend resource on incident response	Lower
	Access to expertise on limiting the publication of sensitive data	Lower
	Increased time to make a decision	Lower
	Access to ransom negotiation services	Increase
	Increased liquidity to spend on the ransom	Increase
	Access to legal advice, including around the sanctions list and payment of ransoms	Ambiguous
Information set of victim	Information on how to conduct ransomware negotiations	Increase
	Information on the criminals past record of ransomware negotiation	Ambiguous
	Information on the likelihood of decryption	Ambiguous
	Information on the likelihood of sensitive data being published	Ambiguous
	Expert advice on post-breach recovery including options and timescales	Ambiguous
	Expert advice on the likely financial losses from different recovery paths	Ambiguous
Information set of criminal	Criminal aware of increased liquidity of victim to pay a ransom, or sub-limits to payment	Increase
	Criminal aware the victim has access to post-breach services	Increase
	Criminal familiar with the insurers stance on paying ransoms	Ambiguous

Table 3. Potential channels of influence for insurance on the ransomware pay decision and whether it lowers, increases incentives for insureds to pay or the effect is ambiguous

A victim who can quickly recover has less incentive to pay a ransom while a victim who learns it would be a long road to recovery has more incentive to pay a ransom. Given the ambiguous possible influence of information we shall pay particular attention to this in analysing the interview data.

In modelling the influences of insurance it is critical to recognise that it may also influence the actions of the criminal which then impact on the victim. In particular, if the criminal knows the victim is insured then that signals information about the victim. Indeed, criminals may want to actively target businesses with insurance. For example, if the criminal knows the victim is insured, and knows the range of payments that the insurance company has permitted in the past then they may be able to set the ransom demand at a level which the victim is more willing to pay. In short, 'familiarity' between the criminals and insurers not only allows a flow of information to the victim but also allow greater 'coordination' between the victim and criminal. This will ultimately act to increase the incentive to pay the ransom (potentially because that ransom is set at a lower level). There are some interesting parallels here with regard to kidnap and ransom (e.g. 39; 40). Potentially the victim may be in a better position if the criminal does not know the victim has insurance (18). This quote from a ransom negotiation, reported in (18), illustrates the issue

> Look, we know about your cyber insurance. Let's save a lot of time together? You will now offer 3M, and we will agree.

Our theoretical framework, as illustrated in Figure 1 and Table 3, provides a general approach to model the many channels through which insurance can impact on the ransomware decision. Specifically, rather than consider a very specific parameterized model that directly captures one particular aspect of the ransomware decision, we have operated at a level of generality that allows us to understand the complex web of influences that insurance may have. To refine our understanding and evaluate which channels may be of most importance we turn to our empirical data.

4 The impact of insurance on the pay decision

The channels of influence discussed above, and summarized in Table 3, show that insurance can, in theory, lower or increase the incentive to pay a ransom. In this section we summarize the data from our expert interviews with an objective to identify the perceived main channels of influence of insurance and the net impact it has on the pay decision.

4.1 Cheaper to pay than rebuild

One recurring theme in the interviews, with views on both side of the argument, was the perception that paying the ransom is the cheapest way for a business and, therefore, the insurer to minimise the losses incurred with business interruption. For instance, an incident response manager said:

> If it's going to cost the insurer less to just pay the ransom than have to support an organisation through rebuilding through all the business impacts then the insurer is incentivised to just pay it.

Similarly, an insurance broker said:

> As an insurer, if your client chooses not to pay that ransom, quite often it can cost us a lot more money because we don't get the solution.

A claims officer at a cyber insurer even commented that:

> Sometimes the decryption key is a much faster means of recovery even if you had backups.

Other interviewees, however, pushed back on the notion that paying the ransom is a 'less costly solution'. For instance, a lawyer said

> You get a decryption key and it can still take 3 weeks to get your systems up and running and iron out all the bugs – or indeed much longer. So paying the ransom isn't a panacea.

Similarly, a cyber-security executive said:

> There is no truth to the paying and decrypting your data as fast... or faster than recovering from backups. People think that its this magical light switch, you know, as soon as you pay, everything is just fine again... it couldn't be further from the truth. It's a very time consuming, laborious process that involves a lot of trouble shooting. And that's just to decrypt the data. Then you have to remediate the machines, you have to do the data migration, you have to test the data. When you restore from backup, you're getting clean everything.

The CISO of a large organisation reflected on the incentives to pay:

> [Organisations may pay] to get their business back up and running and seeming like an easy fix ... the idea that it is as simple as just a decryption key, and if I could just get hold of that I could crack on as if nothing has happened. In reality that's so far from the truth but I do think that some people do think that that's the case.

In terms of our model these split perceptions point to a high degree of uncertainty on the payoff consequences of paying a ransom. Some suggest that paying the ransom is a quick and effective way of reducing losses, $l(t)$, while others suggest it will do little to alleviate losses. The important question for us is how the presence of insurance influences the perceived benefit (or loss reduction) of paying the ransom. We identify three channels through which insurance may have influence. One is the extra information that insureds can receive on the relative returns from paying the ransom, such as the probability of obtaining working decryption keys versus the costs of a rebuild. As we summarized in Table 3, information can have an ambiguous effect depending on what that information is. Given that information will be covered in detail in a subsequent sub-section we postpone discussion of this issue.

In explaining the other two channels of influence we recognise a potential misalignment of the incentives of the victim and insurer regarding system recovery and rebuild. In particular, the insurers role is to return the victim to their pre-breach status, even if that means returning the victim to a relatively cyber-insecure position. The victim, by contrast, may see the breach as an opportunity to reappraise their overall cyber-security strategy. Given that the organisation naturally has a longer term horizon for the organisation than the insurer (on a fixed term contract) they have an incentive to invest more in system restoration and rebuild than the insurer. This was set out by a cyber-security consultant

> we can't go and rebuild the client's environment and all that stuff, there's very strict rules or unwritten rules about what goes into that and so there's a tension .... we help the client respond, but we don't leave the client in a better position by installing software controls etc. than they were before. Because that's not the insurers responsibility.

This tension creates competing incentives. On the one hand insurance lowers the cost of post-breach restoration and so, in that sense, decreases incentives to pay the ransom. On the other hand insurance could restrict the action set of a victim in a way that rebuild seems less appealing. For instance, an incident response manager commented:

> if the insurance coverage wasn't there at all, they might've said, well it's going to hurt us a little bit more to rebuild but we'll do it.

The net effect of these various channels of influence is unclear and will likely be context dependent. It is, therefore, far from clear that insureds are more likely to pay than non-insureds because paying the ransom is 'cheaper' than rebuild. One factor, however, that is crucial is the quality of post-breach support. This naturally feeds into the information that insureds receive and the quality of system restoration that insurers can offer. A cyber-security executive expressed concern about this eco-system:

> what has happened over the last two years is we've had this race to the bottom on the forensics side, and a lot of those companies are also doing the data restoration; they make a ton of money off the data restoration..... Most carriers don't realise they are getting absolutely robbed right now, because they're getting the wrong vendors on the wrong incidents.

This view is at odds with the findings of (48) and so may not be representative. If insureds are getting a 'poor service' then this may influence the decision to pay the ransom.

4.2 Liquidity and moral hazard

Another recurring theme in the interviews is how insurance legitimizes and facilitates ransom payments. As we discussed in Section 3 this can be considered a form of moral hazard, in which the victim uses ransom negotiators and potentially pays a ransom, because the cost is covered by the insurer. It is also a consequence of increased liquidity, because a victim with insurance has more access to money to pay the ransom. For instance, an insurance broker said:

> there's no doubt that cyber insurance has led to more ransom payments because it's a different decision for the insured, right? In that, if I have insurance, I don't bear the full cost of it obviously, so I'm now, I have a lot more resources to be able to consider making a payment.

A couple of interviewees highlighted the role that ransom negotiators play in facilitating payment. For instance, an underwriter said of ransomware negotiators:

> they help facilitate the payment as well, right, because paying in cryptocurrency is not something that most businesses are set up to do

A CISO of large business, when asked about the influence of insurance, reflected on how insurers may be able to lower ransom demands:

> From my point of view it probably will probably affect the amount of ransom. If your insurers are providing you with some decent negotiators, you might reduce the amount you're paying.

This could incentivize payment, albeit with a lower amount of money going to the criminals.

An additional factor that came out clearly in some of the interviews is how insurers may change perceptions of paying the ransom. In our model this would be operating through the information channel in that it changes the beliefs of the victim. For instance, a cyber-security executive said:

> they legitimise the ransomware market by pushing into a default setting on payment .... having on staff negotiators seems to kind of legitimise the market, you're basically saying that this is such a normal thing to do that we've hired a guy.

Overall, therefore, there seemed a clear perception from many of the interviewees that moral hazard, liquidity and the normalisation of ransom negotiation are factors that incentivize insureds to pay. Although, this is not to say that other channels of influence may not counter-balance this effect.

4.3 Insurance gives advice and information

We turn our attention next to the role that insurance can provide in informing the victim. Many of the interviewees argued that insurance can lessen the incentive for the victim to pay by offering them more options, expertise and time. For instance, an underwriter pointed to the additional actions and information afforded to insureds:

> If they were uninsured, they just wouldn't have the options that we're giving them, potentially they wouldn't have access to the same level of expertise to try to recover their data, to access publicly available decryption keys if that is the case and to work through all of that process, to get to the point where it is the last option to pay the ransom, so with our help I think it's a lot more likely that the ransom wouldn't be paid.

To gather information takes time and insurance allows this. A representative of an insurer pointed to the role insurance plays in changing the payoff function to afford more time to insureds:

> a lot of our clients are going well, okay, that means we can have some down time, we can have some business interruption, we can do these things and we haven't got to pay these people and that's actually not a bad thing.

More time can also have behavioural implications, as reflected by an underwriter:

> I think one of the big benefits we've seen with folks who have policies, is it gives them the chance to take a step back and evaluate what's really going on and not to rush themselves into a decision to pay a ransom quickly. If you're not insured and you have a ransom of a say a quarter of a million dollars, people panic, we know when people panic, they make poor decisions.

As remarked by an interviewee from another insurer this behavioural implication can be particularly important for SMEs where owners have built up a business over many years and may panic when faced with a ransomware incident.

More ambiguous is the effect of information that insureds can provide about the criminals. An insurer broker remarked:

> we've got statistics now ... and depending on who we believe is actually using it on us, we know what their average ransom demand is, we know how often they pay, or how often they actually give you the key. ... so we can have a good idea where we're gonna end up very quickly if we go down the payment route.

Payment of the ransom may become more likely with insureds if the criminal's behind it have a 'good reputation' of honouring ransom payments. The net effect of this will depend on the proportion of ransom attacks by criminal's with a good reputation. But, as some of the interviewees pointed out, the existence of a negotiation eco-system gives the impression that insurers are willing to 'trust' some criminal gangs. Sophos' 2021 annual survey suggested that around 65% of data was restored following a ransom payment (41). A two-thirds chance of restoring data may seem like a reasonable risk if the ransom demand is proportionate.

One particular strand of information that is important in informing the ransom decision concerns sanctions. An underwriter explains:

> The insurers do not want to pay ransoms, end of, because the actual checks that you have to do to make sure the money's not going to a terrorist entity or organisation, you can never be 100% sure, particularly in the cyber space

There was also a view that, in a highly dynamic environment, the insurers or reinsurers may refuse to reimburse a payment because they are not allowed to pay. The net effect of information, therefore, given the various channels it operates through would seem to lessen the incentives to pay the ransom. This provides a counter-weight to the effects of moral hazard and liquidity.

4.4 Double extortion is important

A final theme we highlight is that of double extortion. Many interviewees expressed a view that double extortion is making victims more willing to pay a ransom. The question for us to explore is whether insurance influences the cost-benefit trade-off of data leakage. Several interviewees expressed a view that double extortion is complicating the cost-benefit trade-off. For instance, the manager at a software company said:

> much rather deal with a business interruption event than a data breach, 10 out of 10 times. I mean, business interruption, it's a lot more easy to control, easy to identify, and from an insurance perspective, make a claim for, compared to data breach, where you have a lot of reputational damages that aren't insured

A claims officer at an insurer remarked

> lot of ransomware victims in 2021 are more concerned about the potential publication or disclosure on the dark web of their information than they even are about recovering their systems. And I'm fond of saying, that's an entirely different underwriting calculus than can these organisations recover, how quickly can they get their business back up and running, service their clients, build their widgets, whatever it might be. Because now we're underwriting whether they feel shame or embarrassment

In terms of whether an insureds are more or less likely to pay there were several interviewees suggested insurance incentivizes payment. For instance, a consultant at a cyber-security firm said:

> I think from my experience it can be damage limitation somewhat from the insurers because ... the sooner potentially you pay that, the less damage going down the line, maybe the less data's exposed.

There was also a clear view from some of the interviewees that paying a ransom to avoid data leak is not a good cost-benefit trade-off. For instance, a cyber-security executive remarked:

> with data exfiltration, the only thing you may avoid, is a little bit of PR, when like, the data hits one of these leak sites, which frankly is a flash in the pan.... Paying doesn't absolve any company from its legal obligations to notify... we do see companies getting extorted again, after they've paid the first time. And there's no way to prevent that, right, because they're just keeping a copy

The information flow to victims on data leakage may, therefore, be slightly misaligned and result in insureds being incentivized to pay the ransom.

4.5 Net effect on insureds

We have seen some factors that suggest insureds may be more likely to pay a ransom, e.g. moral hazard and liquidity, and some that suggest insureds may be less likely to pay a ransom, e.g. more information and options. It is natural, therefore, to consider the net effect. In Table 4 we provides a summary of the perceptions of our interviewees. Recall that there were 52 interviews in total. In 32 of the interviews we identified a stated opinion on whether victims with insurance were more likely to pay than victims without insurance. We have categorized these as a clear yes, maybe yes, neutral and clear no. Note that neutral implies an opinion that insurance has no effect (rather than the absence of an opinion). We also categorized the main reason given for the view where we categorise according to our model.

We highlight that the data in Table 4 reflects the perceptions of our interviewees and so should not be seen as definitive evidence of the reality on the ground. One interesting thing to observe is a a clear difference in opinion between those working in cyber insurance (middle column) and those in the wider cybersecurity community (right column). In particular, the majority view of those working in cyber insurance is that insureds are less likely to pay. Interviewees, in particular, highlighted how insurance expands the action set of the victim. Those working outside cyber insurance expressed, however, an almost universal opinion that insureds are more likely to pay. A range of reasons were given but liquidity and moral hazard were a recurring theme. These differing opinions between those working in the insurance industry and outside the industry could be interpreted or explained in different ways as we shall now explore.

5 Evaluating the net impact of insurance on the pay decision

In this section we apply our model to try and reconcile the very divergent views we heard in the expert interviews on the impact of insurance on the pay decision. In doing so we we recognise that the losses from the attack are critical to the pay decision. Connolly et al. (2020) (10) proposed an impact assessment instrument for ransomware attacks that allows a measure of severity. They distinguished five categories of negative outcome: business interruption, recovery time, affected

Opinion	Insurance	Non-insurance
Insureds more likely to pay		
Yes	2	8
Maybe yes	1	8
Neutral	5	2
No	6	0
Why yes		
Moral hazard	1	7
Cheaper than rebuild	1	4
Liquidity	1	4
Avoid data publication	0	2
Facilitate payment	0	2
Why no		
Gives more options	6	0
Legal blocks	2	0
Total	14	18

Table 4. Summary of perceptions on whether ransomware victims with insurance more likely to pay ransom

devices, encrypted information critical to business and information loss. One limitation of this framework is that it does not cover loss of sensitive data and so that could be seen as sixth category. For each category they proposed a measure of severity from low to high. Adopting this approach we can talk of a low severity or high severity ransomware attack. For instance, a high severity attack could be one with high levels of business interruption, high recovery time and high loss of sensitive data.

The severity of a ransomware attack for an organisation will naturally depend on a range of factors specific to the organisation (31). Some factors will reflect the sector in which the organisation works; for instance, the extent to which operations are digitized and the amount of sensitive information stored. Crucially, additional factors will depend on the pre and post-breach cyber-security actions of the organisation (13; 10). This is where, we conjecture, insurance can prove crucial. To illustrate we will contrast a high severity and low severity attack.

Consider an attack that is high severity, in the sense that it is highly severe on several of the six potential impacts. This would seemingly fit the 'last resort' category where insurers would advise paying a ransom, and factors like moral hazard and liquidity facilitate that payment. In this scenario non-insureds may also be inclined to pay a ransom. For instance, on the subject of table-top ransomware exercise run in their company, a CISO commented:

> [the conversation] went quite quickly from a "well we definitely shouldn't pay money to bad people", to "we should definitely pay money". It doesn't take very long to get through that 180. People don't want to do it, but then, what we actually need, we need every hope of recovering as quickly as possible.

While all victims in this ‘last resort’ scenario may consider payment, a business without insurance may, as we discuss in Section 4.2, be unable to afford the ransom or lack the capabilities to work through the options and negotiate with the criminals. For instance, a lawyer, when asked about insurers giving pre-approval for ransom payments, remarked:

> there’s two reasons why you might want to pay a ransomware threat actor. Number one is because your business is down and the only way to get it back up is to decrypt, and the only way to get the decryptor is through the extortion, that’s number one. And the second reason is because the threat actor stole your data and threatens to release it to the world .... In those two instances, then we do recommend, make recommendations about paying ransomware threat actors and we do help facilitate those transactions.

In terms of our model, and the channels identified in Table 3, insurance influences payoffs (recompense of the ransom), action set (access to ransom negotiation services, legal advice and increased liquidity), and information (on negotiation, decryption and data leakage). The influence of these channels we depend on the characteristics of the business. For instance, an SME is more likely to be liquidity constrained and lacking in negotiation expertise than a large business. Overall, however, we argue the balance in this ‘last resort’, high severity case is towards insureds being more likely to pay a ransom than non-insureds. As a CEO remarked:

> I think the times where companies pay the ransom because they have insurance, that’s more likely when, if they’re going bankrupt or something, or if they don’t have strong cashflows, but they’ve got a million dollar limit or something, those scenarios, yeah, they’re gonna make the payment because they have insurance, otherwise they wouldn’t have had the ability to.

Consider next an attack with low severity. Then we conjecture that, on balance, the victim may be less likely to pay if they are insured than not insured. This is where the information and options that insurers can provide, discussed in Section 4.3, come to the fore to avoid a ransom payment. As an underwriter confirmed

> if there are good backups in place, then the ransom almost always just becomes off the table.

This was a view echoed by other interviewees from the insurance sector. Again, in this setting a victim without insurance is also unlikely to pay. Crucially, however, the time and information that insurance provides may be critical to the victim having an awareness they can recover without paying the ransom. In terms of our model, and the channels identified in Table 3, insurance would influence the action set (access to public relations, data/systems restoration and forensics and access to expertise on data exfiltration) as well as information (expert advice on post-breach recovery). Again, these factors could be particularly relevant for SMEs who lack the in-house expertise to respond to an attack. On balance, therefore, we conjecture that insureds who suffer a low severity attack are no more likely to pay, and may be less likely to pay if the expertise of the insurers allows alternative options.

If the influence of insurance depends on the severity of attack then the net impact of insurance will depend on the relative frequency with which insureds and non-insureds suffer a severe attack. In particular, the preceding discussion focuses on the question: (1) Fixing the severity of attack, are insureds more likely to pay the ransom than non-insureds? We argued the answer may depend on the severity of attack we are analysing. The complimentary question, we now consider is: (2) Are insureds more or less likely to suffer attacks of a particular severity. It can be argued that businesses with insurance are likely to have more robust security practices in place because of underwriting criteria and pre-breach advice (47; 23; 37; 30; 29; 45; 43). If so, this would mean insureds will typically fall in impact scenarios with less severe attacks. We look to the Cyber Security Breaches Survey (CSBS) 2018-2021 for data to inform on this issue, recognising that this means a focus on UK business.

One proxy for whether victims will suffer a severe impact from attack is the implementation of cyber-security controls within the organisation (36). In the CSBS participants are asked a range of questions about the controls they implement; these are then scored against the UK Government's Cyber Essentials controls and the NCSC 10 Steps to Cyber Security. Informally, Cyber Essentials can be seen as a minimum acceptable standard of controls while the 10 Steps represents a higher level of aspiration. In Figure 2 we plot the proportion of businesses in the CSBS satisfying the Cyber Essentials controls as well as the average number of steps (out of 10) a business satisfies. We also plot the proportion of businesses claiming they have a secure back-up and those with an incident management process. Across all four measures we can see that insureds show a high level of controls. This difference is particularly pronounced for micro and small businesses. Regression analysis, reported in Table 5, controlling for business size, sector, and use of digital technology, shows that the coefficient on insurance is highly statistically significant. This is consistent with the notion that businesses with insurance are more cyber secure.

Given that insureds have, on average, higher controls we may expect a lower severity impact from a ransomware attack. The CSBS allows a proxy measure of this impact. Specifically, we have data on whether the business experienced a ransomware attack in the previous 12 months and we have data on the impact cyber-attacks (overall if there was more than one attack) had on the business over the last 12 months. In Figure 3 we plot the proportion of ransomware victims who experienced the various impacts measured. You can see that there is no discernible difference between insureds and non-insureds in terms of impact. We also analysed data on time to recover and again found no discernible difference. We do not, therefore, find any evidence that insureds suffer a less severe impact from ransomware than non-insureds. This may be because the CSBS data is

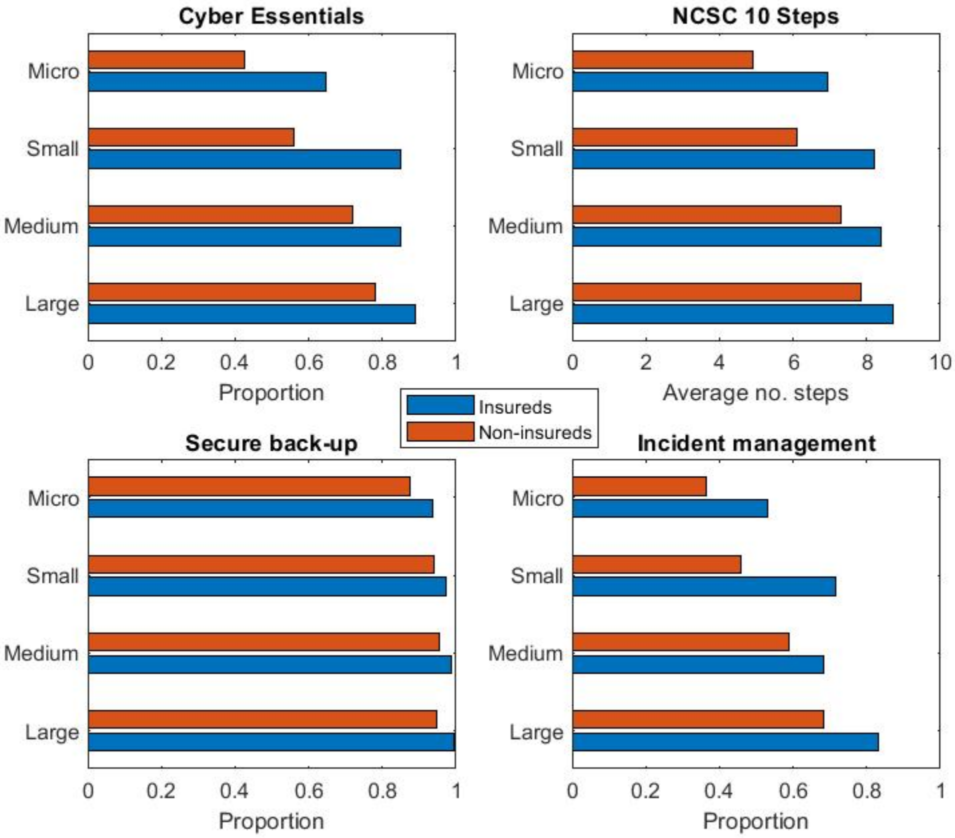

Fig. 2. Proportion of businesses satisfying 5 security controls of Cyber Essentials, the average number of NCSC 10 Steps, the proportion with secure back-up and the proportion with incident management

not fine-grained enough to pick up a difference. In particular, there were only 336 businesses in the sample that suffered a ransomware attack and so this is a relatively small number of observations. It may also reflect that businesses with insurance are better aware of disruption because of the support provided by the insurer.

Even though the evidence from the CSBS is mixed there is, at least, some strong evidence that insureds have a higher level of cyber-security controls. That could lend support to the notion that insureds are, on average, likely to suffer lower severity attacks. Although, this will depend on whether the controls are the 'right ones' to stop and mitigate ransomware. The evolution of underwriting standards is likely critical to this net impact. The following quote from an underwriter illustrates:

	10 Steps	Essentials	Back-up	Incident
Insurance	1.06 (0.08)***	0.41 (0.06)***	0.05 (0.02)**	0.17 (0.015)***
Business size				
Micro	-2.09 (0.09)***	-0.80 (0.06)***	-0.09 (0.03)	-0.25(0.02)***
Small	-1.21 (0.10)***	-0.53 (0.07)***	0.006 (0.03)	-0.16(0.02)***
Medium	-0.34 (0.10)***	-0.19 (0.07)***	0.02 (0.03)	-0.07 (0.02)***
Business sector				
Administration	0.35 (0.26)	0.15 (0.17)	0.23 (0.07)***	0.05 (0.05)
Construction	-0.03 (0.26)	0.11 (0.17)	0.28 (0.08)***	0.00 (0.05)
Education	0.65 (0.27)	0.23 (0.19)	0.22 (0.08)***	0.05 (0.05)
Entertainment	0.05 (0.27)	0.00 (0.18)	0.15 (0.08)*	0.06 (0.05)
Finance	1.30 (0.27)***	0.42 (0.18)**	0.17 (0.08)**	0.19 (0.05)***
Food	-0.57 (0.26)	-0.16 (0.17)	0.019 (0.08)	-0.01 (0.05)
Health	0.84 (0.27)***	0.24 (0.18)	0.19 (0.08)***	0.11 (0.05)**
Information	0.86 (0.27)***	0.38 (0.17)**	0.27 (0.08)***	0.12 (0.05) **
Professional	0.68 (0.26)***	0.31 (0.17)*	0.26 (0.07)*	0.08 (0.05)
Transport	0.12 (0.27)	0.17 (0.18)	0.19 (0.07)***	0.08 (0.05)
Utilities	-0.01 (0.26)	0.14 (0.17)	0.19 (0.03)***	0.00 (0.05)
Digital use				
Social media	0.64 (0.06)***	0.22 (0.04)***	0.08 (0.02)***	0.06 (0.01)***
Online booking	0.02 (0.06)	0.04 (0.04)	0.01 (0.02)	0.01 (0.01)
Online banking	0.12 (0.13)***	0.11 (0.04)***	0.11 (0.02)***	-0.00 (0.01)
Industry control	0.77 (0.13)***	0.20 (0.08)**	0.14 (0.04)***	0.09 (0.02)***
Number observations	5878	5878	5878	5878
R squared	0.40	0.18	0.09	0.41

Table 5. Regression result of security measures on insurance and control variables. Additional controls (not reported) for year, social enterprise and cyber priority. *** $p < 0.01$, ** $p < 0.5$, * $p < 0.1$.

> My perspective is ... paying the ransom is usually the very last option on the table. ... what we will... advise is, to try and recover without paying the ransom first, and when we're underwriting businesses, the key part is trying to work out in the recovery piece, whether they do have that alternative solution, so they're not stuck with the only option being paying the ransom.

Over time it would be expected that underwriters will obtain a better understanding of cyber risk and, thus, be able to identify the controls that proxy for a lower risk of severe impact. If so, our analysis suggests the net impact of insurance can be positive.

We finish by noting that there is an interesting dynamic effect at play here as well. In the past underwriting controls may have been relatively lax meaning that insureds were suffering high severity impact attacks. We conjectured that in such attacks insureds may be, on average, more likely to pay the ransom. This may be reflected in the perceptions of some the experts we interviewed, given their experience. Over time, as controls are increasing, the net impact of insurance will, according to our model, become more positive. This more forward

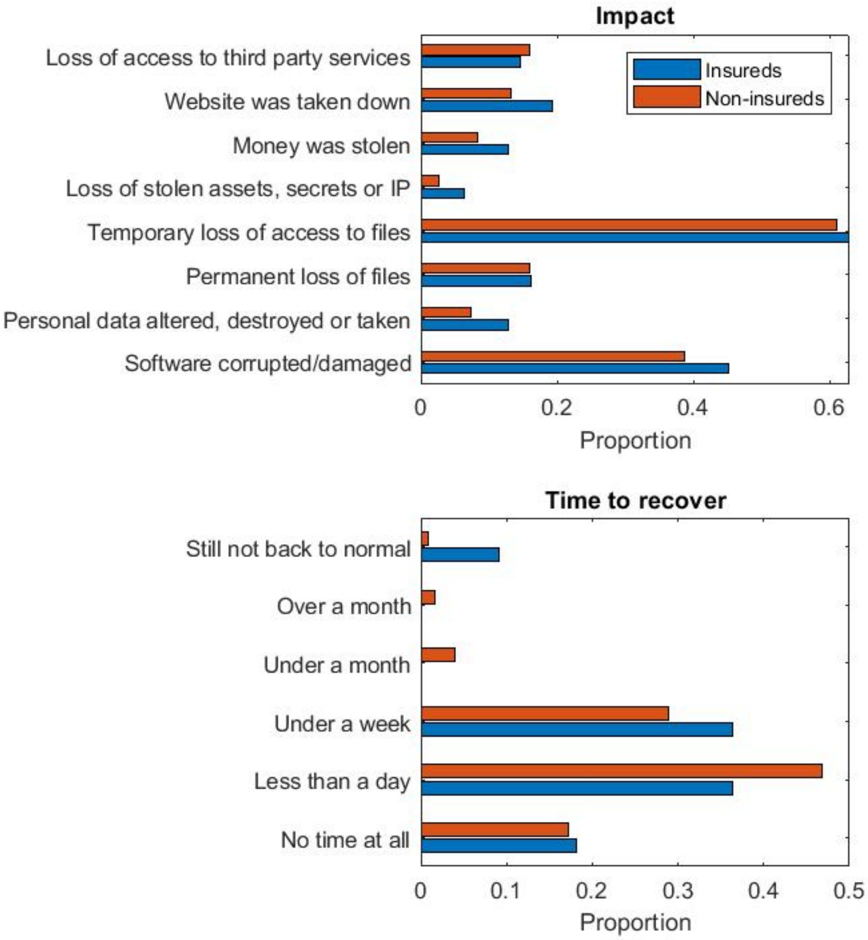

Fig. 3. Disruption to ransomware victims in terms of impact and time to recover

looking positive view of the role of insurance may be reflected in the views of those working in the cyber-insurance sector. As a cyber-security lead remarked:

> at the end of last year we saw a lot of cases where the default was let's just pay this thing and get on with our lives. There wasn't always the mature understanding of what's been lost, what's been impacted, what can we recover. It was almost like a knee jerk reaction that we've got the insurance and lets just claim. That may have led to some of the big claims that we have today and why insurers are being a little bit

> more, I'd say guarded with what coverage they're going to provide going forward in this area.

6 Conclusion

In this paper we have analysed the impact that insurance is having on the decision of ransomware victims to pay a ransom. Ransomware is one of the foremost cyber-threats to organisations and, given that it provides a viable business model to criminals, is likely to remain a threat for many years to come (19). It is, thus, vital to see what role insurance can play in pooling risk and mitigating the losses that results from ransomware. Some have claimed that the insurance industry has fuelled ransomware by increasing the incentives for victims to pay a ransom (e.g. 4; 12; 38; 46). Existing game theoretic work also supports that broad premise (e.g. 5; 49). We, however, come to a different conclusion and argue that insurance need not be incentivizing ransom payment.

Our theoretical framework was designed to capture two key aspects of the ransom decision making process: (i) it is dynamic in being played out of a period of time, which may be months rather than days, and (ii) there are large amounts of imperfect and incomplete information on the victims side across a wide range of issues, including their ability to recover from back-up, the extent of data loss, and the viability of paying the ransom. In a dynamic Bayesian game setting, such as this, information acquisition and interpretation is critical to the optimal strategy (15). This, as many of the experts we interviewed remarked, is where insurance can play a crucial role; insurance can afford the victim time, because business interruption is compensated and expert advice is available, and it can provide the information needed to take actions that minimize losses.

We demonstrated, though, that the impact of information can be ambiguous. If the victim faces a low severity impact attack and has the means to restore from back-up etc. then information allows the victim a solution that does not involve paying the ransom. By contrast, if the victim faces a high severity attack and the criminals have a reputation of 'honouring' ransom payments then information points towards paying the ransom. And many of our expert interviewees expressed a view that insurance incentivizes and facilitates payments in cases where all other options have run out. In this case insurance provides liquidity to victims and normalises the notion of paying the ransom. The net effect of insurance, thus, depends to a large extent on the ability of organisations to have viable options to recover other than paying the ransom.

It is widely acknowledged that cyber-security practices in many organisations are still below the level one would hope. This view was expressed by cyber-security experts in our study and reflected in the UK's Cyber Security Breaches Survey. For instance, as you can see in Figure 2, around 1 in 5 large businesses in the survey are not even meeting the minimum Cyber Essentials standard. Lax controls inevitably mean that ransomware victims may end up paying the ransom as a last resort. We provided evidence that insureds have a higher level of cyber-security controls, suggesting they are less vulnerable to a severe impact.

If cyber-security standards increase over time and the insurance industry proactively supports better controls then we enter a landscape where insureds should be less likely to pay a ransom than non-insureds. Whether or not we have reached that point yet is not clear. The experts we interviewed were very split on whether the net effect of insurance is to increase or lessen the incentive to pay the ransom. The direction of travel, though, could be towards the insurance industry being part of the solution.

There are a wide range of issues that it would be valuable to study in future work. Ideally, additional data can help further our understanding. For instance, we could identify only 336 businesses in the CSBS that had suffered a ransomware attack (out of 5878 businesses surveyed). The inferences that one can draw from such a relatively small number of victims are naturally limited. Ideally we would have data on a large number of victims, including whether they paid the ransom, the severity of attack, and the presence of insurance. This would allow us to directly test the impact of insurance on ransomware payments. Accessing ransomware victims is, however, difficult (10) and so it is unlikely such data will be available anytime soon. In the absence of such additional quantitative evidence it remains important to continue collecting qualitative evidence from experts in the field and victims of ransomware attacks, particularly in the fast paced environment of ransomware and cyber insurance.

Appendix

Theme	Question
Overview	What is your role and industry? How many years of experience do you have in your industry? Have you always worked in cyber? If not, what other sector(s) did you work in before? What types of organisations do you provide services or cover for? Any specific sectors?
State of the market	What is the state of the cyber insurance market and how has this changed over the last few years Has ransomware played a role in the changes in the cyber insurance industry? If so, how? Is ransomware coverage included as standard as part of your policies? Are all organisations that request cyber insurance coverage currently able to attain it (i.e. to find a company willing to underwrite the policy)? How is the insurance industry dealing with the challenges posed by ransomware?
Potential positive impacts of ransomware	What is the role of cyber insurance in helping organisations mitigate the threat of ransomware, particularly its impacts? Do insurers require specific security controls to underwrite an insurance policy? If so, which ones? How do you assess an organisation's risk? Have underwriting practices changed since ransomware became more of a threat? (mention adverse selection here) Do you provide any pre-incident services that specifically aim to mitigate the threat of ransomware? Are they in-house services or provided by an external vendor? Do policies include any kind of warranties or security obligations that need to be fulfilled for a ransomware claim to be paid? For instance, patching vulnerabilities etc.
Claims/Incident Response	What sort of post-incident services do you/cyber insurers provide (DFIR, breach counsel etc.)? Are these in-house capabilities or thirdparty? What are the roles of the various post-incident services (including your organisation) after a ransomware attack on an insured organisation? Can you talk us through what happens when a policyholder comes to you during a ransomware attack? Do you have an in-house breach response/claims handling team? How much influence do you have on policymaker's decisionmaking during a ransomware attack? What are your priorities during a ransomware incident? Who is the most influential external actor (DFIR, insurer, breach counsel etc.) during a ransomware incident? What of information do you get from DFIR reports?
Payments	From an economic perspective, there is an argument that it is often cheaper to pay the ransom than to try to recover from backups. What are your thoughts on this? Are organisation that are insured more likely to pay ransom demands? Are there any other incentives that push victims to pay the ransoms? Under what conditions are insurers willing to pay the ransom? How does an insurance company decide what ransom amount they are willing to pay? Do you think ransom payments will be included in cyber insurance coverage in the future? Has cyber insurance incentivised more ransomware attacks?
Recommendations	What are the policy and business implications to banning ransomware payments? How can we disrupt the ransomware business model? Which stage (access, cash-out etc.) do we have the best chance of disrupting? What kind of intelligence do insurers have to help law enforcement/government? What recommendations would you make for government, insurers, cyber vendors etc. to help mitigate the threat posed by ransomware? Would you support mandatory reporting for ransomware incidents? How would mandatory reporting affect insurers?

Table 6. Interview questions

Bibliography

[1] Ransomware: to pay or not to pay? EY (2020), https://www.ey.com/en$_u k/consulting/ransomware-to-pay-or-not-to-pay$

[2] Agrafiotis, I., Nurse, J.R.C., Goldsmith, M., Creese, S., Upton, D.: A taxonomy of cyber-harms: Defining the impacts of cyber-attacks and understanding how they propagate. Journal of Cybersecurity **4**(1), tyy006 (2018)

[3] Ahnert, T., Brolley, M., Cimon, D.A., Riordan, R.: Do you know where your data sleeps at night? cyber security and ransomware in financial markets. Cyber Security and Ransomware in Financial Markets (March 14, 2022) (2022)

[4] Bajak, F.: Cyber insurance industry in crosshairs of ransomware criminals. Insurance Journal (2021), https://www.insurancejournal.com/news/national/2021/07/07/621416.htm

[5] Balasubramanian, A.: Insurance against ransomware. Available at SSRN 3846111 (2021)

[6] Beck, C., Fleisher, B.: Does it ever make sense for firms to pay ransomware criminals? Insurance Journal (2021), https://www.insurancejournal.com/news/international/2021/07/08/620508.htm

[7] Caporusso, N., Chea, S., Abukhaled, R.: A game-theoretical model of ransomware. In: International Conference on Applied Human Factors and Ergonomics. pp. 69–78. Springer (2018)

[8] Cartwright, A., Cartwright, E.: Ransomware and reputation. Games **10**(2), 26 (2019)

[9] Cartwright, E., Hernandez Castro, J., Cartwright, A.: To pay or not: game theoretic models of ransomware. Journal of Cybersecurity **5**(1), tyz009 (2019)

[10] Connolly, L., Wall, D.S., Lang, M., Oddson, B.: An empirical study of ransomware attacks on organizations: an assessment of severity and salient factors affecting vulnerability. Journal of Cybersecurity **6**(1), tyaa023 (2020)

[11] Connolly, L.Y., Wall, D.S.: The rise of crypto-ransomware in a changing cybercrime landscape: Taxonomising countermeasures. Computers & Security **87**, 101568 (2019)

[12] Dudley, R.: The extortion economy: How insurance companies are fueling a rise in ransomware attacks. Pro Publica (2019)

[13] Fagioli, A.: Zero-day recovery: the key to mitigating the ransomware threat. Computer Fraud & Security **2019**(1), 6–9 (2019)

[14] Fang, R., Xu, M., Zhao, P.: Determination of ransomware payment based on bayesian game models. Computers & Security p. 102685 (2022)

[15] Fudenberg, D., Tirole, J.: Game theory. MIT press (1991)

[16] Furnell, S., Heyburn, H., Whitehead, A., Shah, J.N.: Understanding the full cost of cyber security breaches. Computer Fraud & Security **2020**(12), 6–12 (2020)

[17] Galinkin, E.: Winning the ransomware lottery: A game-theoretic model for mitigating ransomware attacks. arXiv preprint arXiv:2107.14578 (2021)

[18] Hack, P., Wu, Z.Y.: "we wait, because we know you." inside the ransomware negotiation economics (2021)

[19] Hernandez-Castro, J., Cartwright, A., Cartwright, E.: An economic analysis of ransomware and its welfare consequences. Royal Society open science **7**(3), 190023 (2020)

[20] Kemp, S., Buil-Gil, D., Miró-Llinares, F., Lord, N.: When do businesses report cybercrime? findings from a uk study. Criminology & Criminal Justice p. 17488958211062359 (2021)

[21] Kenneally, E.: Ransomware: a darwinian opportunity for cyber insurance. In: Connecticut Insurance Law Journal Fall Symposium Edition. vol. 28 (2021)

[22] Kerns, Q., Payne, B., Abegaz, T.: Double-extortion ransomware: A technical analysis of maze ransomware. In: Proceedings of the Future Technologies Conference. pp. 82–94. Springer (2021)

[23] Khalili, M.M., Liu, M., Romanosky, S.: Embracing and controlling risk dependency in cyber-insurance policy underwriting. Journal of Cybersecurity **5**(1), tyz010 (2019)

[24] Kharraz, A., Robertson, W., Balzarotti, D., Bilge, L., Kirda, E.: Cutting the gordian knot: A look under the hood of ransomware attacks. In: International conference on detection of intrusions and malware, and vulnerability assessment. pp. 3–24. Springer (2015)

[25] Lallie, H.S., Shepherd, L.A., Nurse, J.R.C., Erola, A., Epiphaniou, G., Maple, C., Bellekens, X.: Cyber security in the age of covid-19: A timeline and analysis of cyber-crime and cyber-attacks during the pandemic. Computers & Security **105**, 102248 (2021)

[26] Laszka, A., Farhang, S., Grossklags, J.: On the economics of ransomware. In: International Conference on Decision and Game Theory for Security. pp. 397–417. Springer (2017)

[27] Lemnitzer, J.M.: Why cybersecurity insurance should be regulated and compulsory. Journal of Cyber Policy **6**(2), 118–136 (2021)

[28] Li, Z., Liao, Q.: Game theory of data-selling ransomware. Journal of Cyber Security and Mobility pp. 65–96 (2021)

[29] MacColl, J., Nurse, J.R.C., Sullivan, J.: Cyber insurance and the cyber security challenge. RUSI Occasional Paper (2021)

[30] Nurse, J.R.C., Axon, L., Erola, A., Agrafiotis, I., Goldsmith, M., Creese, S.: The data that drives cyber insurance: A study into the underwriting and claims processes. In: 2020 International Conference on Cyber Situational Awareness, Data Analytics and Assessment (CyberSA). pp. 1–8. IEEE (2020)

[31] Ochoa, R., Ticse, D., Herrera, E., Vargas, J.: Ransomware scenario oriented financial quantification model for the financial sector. In: 2021 IEEE Sciences and Humanities International Research Conference (SHIRCON). pp. 1–4. IEEE (2021)

[32] Oerlemans, J.J.: Laundering the profits of ransomware. European J. Crimec Criminal Law Criminal Justice **28**, 121–152 (2020)

[33] Oz, H., Aris, A., Levi, A., Uluagac, A.S.: A survey on ransomware: Evolution, taxonomy, and defense solutions. arXiv preprint arXiv:2102.06249 (2021)

[34] Paquet-Clouston, M., Haslhofer, B., Dupont, B.: Ransomware payments in the bitcoin ecosystem. Journal of Cybersecurity **5**(1), tyz003 (2019)

[35] Rege, A., Bleiman, R.: Ransomware attacks against critical infrastructure. In: ECCWS 2020 20th European Conference on Cyber Warfare and Security. p. 324. Academic Conferences and publishing limited (2020)

[36] Reshmi, T.: Information security breaches due to ransomware attacks-a systematic literature review. International Journal of Information Management Data Insights **1**(2), 100013 (2021)

[37] Romanosky, S., Ablon, L., Kuehn, A., Jones, T.: Content analysis of cyber insurance policies: How do carriers price cyber risk? Journal of Cybersecurity **5**(1), tyz002 (2019)

[38] Sabbagh, D.: Insurers 'funding organised crime' by paying ransomware claims. Guardian (2021), https://www.theguardian.com/technology/2021/jan/24/insurers-funding-organised-by-paying-ransomware-claims

[39] Shortland, A.: Governing criminal markets: The role of private insurers in kidnap for ransom. Governance **31**(2), 341–358 (2018)

[40] Shortland, A.: Kidnap: Inside the ransom business. Oxford University Press (2019)

[41] Sophos: The state of ransomware 2021 (2021), https://www.sophos.com/en-us/content/state-of-ransomware

[42] Spence, N., Niharika Bhardwaj, M., Paul III, D.P.: Ransomware in healthcare facilities: a harbinger of the future? Perspectives in Health Information Management pp. 1–22 (2018)

[43] Sullivan, J., Nurse, J.R.: Cyber security incentives and the role of cyber insurance. RUSI Emerging Insights Paper (2021)

[44] Tuttle, H.: Ransomware attackers turn to double extortion. Risk Management **68**(2), 8–9 (2021)

[45] Uuganbayar, G., Yautsiukhin, A., Martinelli, F., Massacci, F.: Optimisation of cyber insurance coverage with selection of cost effective security controls. Computers & Security **101**, 102121 (2021)

[46] Wolff, J.: As ransomware demands boom, insurance companies keep paying out. Wired (2021), https://www.wired.com/story/ransomware-insurance-payments/

[47] Woods, D., Agrafiotis, I., Nurse, J.R.C., Creese, S.: Mapping the coverage of security controls in cyber insurance proposal forms. Journal of Internet Services and Applications **8**(1), 1–13 (2017)

[48] Woods, D.W., Böhme, R.: How cyber insurance shapes incident response: A mixed methods study. In: Workshop on the Economics of Information Security (2021)

[49] Yin, T., Sarabi, A., Liu, M.: Deterrence, backup, or insurance: A game-theoretic analysis of ransomware. In: The Annual Workshop on the Economics of Information Security (WEIS) (2021)