Ransomware Insurance and Strategic Ransom Demand

Manos Perdikakis*

2nd March, 2026

Abstract

This paper studies ransomware insurance in a model where payments are determined through bargaining between firms and adversaries. Insurance shapes bargaining through two opposing effects: it improves firms' outside option (commitment) but also increases their ability to pay (liquidity). I show that which force dominates depends on whether adversaries observe insurance contracts. When contracts are observable, insurance reduces equilibrium ransom demands. When contracts are unobservable, this commitment channel weakens, and with sufficiently tight liquidity constraints, insurance can increase ransom payments and make firms worse off. Regulatory caps on ransom coverage can restore the commitment benefits of insurance.

1 Introduction

In this paper, I study markets for insurance against *ransomware* attacks. Ransomware is a type of malicious software which prevents a firm from accessing key functions and data, usually by encrypting files. After criminal groups gain access to a company's system and infect it with ransomware, they demand a *ransom* payment - typically in untraceable cryptocurrency - in exchange for the decryption keys. If companies refuse to pay, attackers threaten to intensify the disruption to their network and to also leak or sell the company's stolen data.[1][2] The 2025 ransomware attack on Marks and Spencer in the UK caused a 16% reduction in (short-term) stock value for the retailer, whose online shop was still not restored seven weeks after the attack [The Guardian, 2025], [Financial Times, 2025].

---

*The Brattle Group. This paper was initially written as part of my PhD thesis at the University of Oxford. I am grateful to my supervisor, Margaret Meyer, for her valuable feedback, and to Arina Nikandrova, Alexei Parakhonyak, and Greg Taylor for very helpful discussions.

[1] These "double extortion" attacks are the norm nowadays, and have contributed to a large increase in the amount of ransom payments, see van Rooyen [2024].

[2] The UK National Crime Agency views ransomware as the largest cyber security threat, [National Cyber Security Centre, 2024], and they expect it to remain so in the future as AI allows hackers to automate their operations and reduce the cost of performing ransomware attacks.

To mitigate such risks, firms increasingly purchase cyber insurance. According to Munich Re (a major reinsurance company), ransomware attacks were the leading driver of cyber insurance claims in 2024 [Munich Re, 2025].[3] Cyber insurance policies typically cover first- and third-party liability, as well as ransom payments (see, e.g., Marsh [2025], Acronis [2025]). Additionally, and very importantly, they also offer incident response services, to help companies better deal with an ongoing ransomware attack, as well as reputation management and regulatory compliance services, to help them deal with the aftermath.

However, there is evidence from cybersecurity reports that attackers seek information on victims' cyber insurance policies in order to determine optimal ransom demands (e.g., Coalition, Inc. [2025], Policyholder Pulse [2023]). This view is corroborated by the study of Cong et al. [2025] who claim that the more sophisticated ransomware groups scan a victim firm's network to locate possible cyber insurance policies before making ransom demands.[4]

This has raised concerns that the insurance sector is inadvertently contributing to higher ransom payments by insured victims, thus funding ransomware groups and motivating increased criminal activity. Following such concerns raised by French government officials, insurer AXA announced in 2021 that it would stop offering new policies that cover *ransom payments* in France (Euronews [2021]). More recently, the US Deputy National Security Adviser for Cyber and Emerging Technology advocated for banning insurance for ransom payments (Financial Times [2024]).

This paper investigates two main questions: first, whether the emergence of an insurance market for ransomware can reduce the welfare of firms, relative to the benchmark in which an insurance market does not exist. Second, whether and what regulatory interventions can raise the welfare of firms, and potentially decrease the expected revenue of hackers. In the model I develop to answer these questions, and motivated by the above discussion, I focus on the role of (1) hackers' ability to *observe* firms' insurance contracts and (2) firms' *liquidity constraints.*

There are several reasons why I study the game under different assumptions on contract observability, not least because the prospect of hackers observing insurance contracts is generating prolific commentary by market participants and policy makers. First, this allows us to identify what welfare effects are due to contract observability or lack thereof. Second, it is currently ambiguous whether contracts will be better modeled as observable once the insurance market has matured. Third, and related, studying the game under unobserved contracts allows us to

---

[3] According to the same source, global premium volume in the cyber insurance market is USD 15.3 billion in 2024 and expected to more than double by 2030.

[4] The following quote by Jamie Hart, intelligence analyst at Digital Shadows, is revealing: "They've been in the [victim's] network, they've seen it, and they're going to argue that [the victim] has coverage and they can afford to pay.." (Gooding [2024]). The title of the piece is indicative of the active debate: "Does cyber insurance increase the risk of a ransomware attack?"

understand how equilibrium will be shaped if insurers and firms adapt by e.g. not saving digital forms of the insurance policies, or adopting other measures to ensure this information does not leak. Such measures are currently recommended by cyber-security specialists.

To answer these questions, I develop a theoretical model to study the strategic interaction between insurance providers, firms, and ransomware attackers. Firms face the risk of a ransomware attack and, following a successful breach, receive a ransom demand from an attacker who can alternatively monetize the stolen data if bargaining breaks down. In the baseline, a monopolist provides insurance against ransom payments and against the business interruption losses that arise if firms reject a ransom demand. The insurance contract therefore affects the firms' willingness to accept or reject a given ransom offer, and indirectly shapes attackers' optimal ransom demands.

In the insurance market, I abstract from moral hazard and private information between insurance buyers and sellers, to focus on the novel strategic aspect of ransomware insurance provision. A central focus of the analysis is whether attackers can observe breached firms' insurance contracts. When contracts are observed, insurance has commitment value for firms. In equilibrium, full insurance against business interruption losses limits attackers' ability to extract surplus through high ransom demands. Insurance **depresses ransom demands** to attackers' outside option, and as a result, attackers are worse off relative to a benchmark without insurance. In the presence of a fully extracting monopolist, firms are equally well off.

When insurance contracts are unobserved, insurers cannot directly influence attackers' ransom demands. Instead, attackers choose ransom demands based on their beliefs about firms' insurance coverage. The strategic complementarity between insurance coverage of ransom payments and ransom demands leads to multiple equilibria. Interestingly, and in contrast to the monopoly case of observed contracts, across equilibria, firms are made better off by the presence of the insurer. That is because *uninsured* firms benefit from a positive externality, and face lower ransom than if attackers knew they were facing an uninsured firm. Thus, the monopolist can extract less surplus in the insurance market and firms are better off.

I then study how firms' liquidity constraints affect the welfare consequences of ransomware insurance. Liquidity constraints limit firms' ability to pay ransom out of pocket in the absence of insurance. Under unobserved contracts, if firms are sufficiently liquidity constrained, the presence of insurance may increase attackers' expected revenue, even though it continues to insure firms against severe losses. In these cases, the presence of insurance markets makes firms *worse off*, and hackers better off. Figure 1 summarizes the welfare effects of ransomware insurance across the different regimes considered in the paper.

Welfare impact of ransomware insurance.

(a) No Liquidity Constraints

	Firms	Hackers
Observed	=	↓
Unobserved	↑	↓

(b) With Liquidity Constraints

	Firms	Hackers
Observed	=	↓
Unobserved	↑ / ↓	↑ / ↓

Figure 1: Upward (downward) arrow indicates weakly higher (lower) payoff relative to the counterfactual without a ransomware insurance market. Ambiguity reflects equilibrium multiplicity.

These results generate a role for regulation of ransomware insurance. By capping insurance coverage for ransom payments, regulators can preserve the commitment value of insurance while limiting attackers' ability to extract surplus from insured firms. Such regulatory interventions can make firms better off and attackers worse off relative to an unregulated insurance market.

In the next section, I review related literature. In Section 3, I present the model with observed contracts. I discuss the main modelling assumptions, derive the model's equilibrium and present extensions. In Section 5, I move on to the model with unobserved contracts and present in Section 6 the empirical predictions of the model in the absence of liquidity constraints. Finally, in Section 7, I extend the model to account for liquidity constraints, under both observed and unobserved contracts. I discuss the results' managerial and policy implications, and then conclude.

2 Related Literature and Contribution

The work closest to mine is Balasubramanian [2021], who also studies insurance with strategic ransomware attackers. In his model, ransomware insurance offers risk-sharing value but by raising firms' liquidity always leads to higher equilibrium ransom demands and more active hackers. However, in contrast to my model, insurance coverage does not discriminate between payments for business interruption and ransom, rather a uniform coverage amount is specified. This implies that insurance increases the expected payoff to hackers by relaxing firms' liquidity constraints.

The recent contribution of Cartwright et al. [2023] discusses how cyber insurance may affect the level of ransom paid, and recognizes that insurance for *incident response* and ransom-rejection costs should lower the incentive to pay ransom. Meurs et al. [2023] provide a first attempt at *empirically* understanding whether firms that hold cyber insurance policies are attacked more frequently and pay higher ransom, using data from self-reported ransomware attacks in the Netherlands. The paper by Cong et al. [2025] studies negotiation transcripts and document operational details of major ransomware groups, including how attackers may seek information about victims' financial resources and insurance arrangements.

More broadly, in the economics of cyber-risk and ransomware, Ahnert et al. [2022] and August et al. [2025] study environments in which attackers choose between monetizing stolen data and deploying ransomware, a choice that parallels the outside option in my model. Ahnert et al. [2022] study an environment with unobserved investments in security. The emergence of ransomware shifts the incidence of cyber-attack losses from consumers to firms and can induce higher levels of security relative to an environment with only traditional data breaches. In August et al. [2025], the authors ask how social welfare changes once the availability of cryptocurrency makes ransomware attacks feasible, relative to a world with only traditional breaches. The option to pay ransom can potentially benefit firms if equilibrium ransom is small relative to the harm they suffer from a traditional data breach.

Laszka et al. [2017] offer a model of strategic investment in backup and ransom negotiations between firms and adversaries. Cartwright and Cartwright [2019] deal with the interesting question of reputation formation by long-lived ransomware gangs, who after every successful attack choose whether to return the firm's data or not. The main finding is that equilibria with reputational incentives are more likely to arise when there is a small number of well-known adversaries, rather a mass of potential entrants. August et al. [2022] study a setting in which consumers of a software can choose whether to implement security patches or not, and the probability of a successful ransomware attack depends on the *mass* of the unpatched consumers. They study how ransomware affects optimal software pricing and welfare. My work is also related to the papers by Fainmesser et al. [2023] and de Cornière and Taylor [2024], who study cyber-security provision by firms in settings with strategic adversaries. Both papers highlight a negative externality that customers of the same service exert on each other, since firms that service more customers and collect more data become more valuable data breach targets.

In studying the role of contract observability, my work relates to the seminal contributions of Katz [1991] and Bolton and Scharfstein [1990], which ask whether contracts between a principal and an agent can serve as *pre-commitments* for the agent in subsequent games. In Bolton and Scharfstein [1990], under publicly observed contracts, a lender's optimal financing contract is chosen to protect the borrower against predation. However, offering that same contract is not credible under unobserved contracts, hence predation occurs in equilibrium. In the setting I study, contract observability determines the extent to which insurance can act as a commitment device that shapes attackers' ransom demands.

3 Baseline Model

The game takes place over a single period comprising several stages. There is a single insurance provider, a unit-mass of identical *firms*, which operate in an industry subject to ransomware attacks and a unit mass of identical *adversaries*,[5] , i.e., hackers who engage in ransomware attacks. The firms maximize expected utility and are risk averse with concave Bernoulli utility function over end-of-period wealth levels that satisfies $u' > 0, u'' < 0$. The adversaries and insurer are risk neutral profit maximizers. The insurer plays first by choosing an insurance contract with terms $(p, M_r, M_b) \in R^3_+$. Aside from the premium, p, the insurance contract is comprised of (weakly positive) terms M_r and M_b, whose function I explain below. Then, each firm has the option of paying the premium to purchase the insurance contract. After a firm has made an insurance-purchase decision, it operates in the market and is "matched" with an adversary who attempts to breach its network. Each adversary is successful with *exogenous* probability $q \in (0, 1)$. I assume that successes occur independently of whether insurance has been purchased or not. In addition, the probability q is known to insurer and firm at the time the insurance contract is signed.

If an attack is unsuccessful, which occurs with probability $(1 - q)$, the game ends. If the attack is successful, the firm and corresponding adversary enter the bargaining stage of the game.

Bargaining Subgame. Given the information available to him at the beginning of the bargaining subgame, the adversary makes a take-it-or-leave-it (TIOLI) ransom offer r to the victim firm. If bargaining is successful and the firm accepts to pay ransom r, the adversary's payoff is r, regardless of whether he is bargaining with an insured or uninsured firm. On the other hand, if there is bargaining breakdown, the adversary earns his outside option of s, the payoff from monetizing information stolen by the firm.

In the event of successful bargaining firms suffer $r + b^{low}$, which is the ransom plus business interruption (BI) damage caused to the firm upon the occurrence of a breach. Without loss of generality, I assume $b^{low} = 0$. This is the part of BI costs that is *sunk* at the time of bargaining, hence a positive value would not qualitatively impact any of my results.[6] In the event of bargaining breakdown, the firm suffers harm b^{high}, and since I focus on $b^{low} = 0$, I will omit the superscript and refer to b^{high} as b. If the firm has purchased an insurance contract with terms M_r, M_b, it receives M_r compensation if it pays the ransom, and M_b compensation

---

[5] For most of the following analysis, we can also consider a single firm and a single adversary.

[6] This does not mean that I assume b^{low} is small in absolute size or relative to any of the other parameters in the model. In fact, as August et al. [2025] discuss, even when ransom is paid, there are significant additional costs from a ransomware attack.

if it rejects the offer.

The analysis proceeds by first characterizing equilibrium outcomes when insurance contracts are publicly observed by the attacker, and then studying how these outcomes change when contracts are unobserved. Before that, I briefly discuss the modelling of payoffs presented above.

3.1 Discussion of modelling and assumptions

Interpretation of b. In reality, the harm that adversaries can cause to the firm following rejection of the ransom offer has various components: (1) business interruption induced by hackers not releasing data and/or network resources that are necessary for the firm's operations, (2) exposure to regulatory punishment (e.g., GDPR fines) if data leaked (see this recent example, Woodruff Sawyer [2024]), (3) reputational harm if consumers become aware of data sale.[7] The recently observed occurrence of "double extortion" means precisely that hackers can threaten to cause both business interruption and data theft in the event of bargaining breakdown (see Cartwright et al. [2023], Cong et al. [2025], Böhme and Schwartz [2010]). All components of the harm caused to the firm following bargaining breakdown are increasing in the "size" of the breach, which in this model is represented by the scalar b.

Interpretation of M_b. Varying degrees of insurance can be provided against all the aforementioned components of the threat: the insurer can provide (1) technical remedies and support via specialized IT staff to mitigate business interruption, (2) legal support to reduce regulatory penalty, (3) crisis management support and credit monitoring for consumers. These are indeed some of the key components of modern cyber insurance policies (see also Ransomware Task Force, Acronis [2025], Marsh [2025]). For example, the Financial Times reported after the recent data breach of $M\&S$ that their cyber insurance policy "..would cover both first-party losses, such as lost sales and incident response costs, as well as third-party losses, such as legal liabilities related to the data breach..".[8] I abstract away from specifying the means of compensation to firms that suffer business interruption and only refer to payment M_b, the insurance payment in the event of ransom rejection.

Symmetric information between adversary and victim. As Cong et al. [2025] find, one of the first pieces of information that ransomware attackers look for in a firm's network is the financial statement, which gives them detailed knowledge of the firm's assets and profitability. The hackers use estimates of revenue lost as a result of encrypted resources to infer the value

---

[7] Current data-protection regulation requires consumers to be notified only if their data is leaked, but investors may have to be informed even in the absence of personal data leakage.

[8] `https://www.ft.com/content/723b6195-1ce7-4b5f-94f5-729e9152c578`, accessed July 24, 2025.

of b.

Insurance does not affect attack probability. Regarding insurance, I assume that adversaries (1) get matched to insured and uninsured firms with the same probability, i.e. there is no ex-ante targeting of either type of firm and (2) and insurance does not offer better security, i.e. the probability of a successful attack is q for both insured and uninsured firms alike. It is ex-ante ambiguous how insurance could affect the total probability of a successful ransomware attack; on one hand, if hackers anticipate larger payments by firms with insurance (holding other firm characteristics fixed), this would induce them to find out which firms are insured and target them. On the other hand, pooling cyber-security expertise means cyber insurers can provide advice to their customers on how to avoid ransomware attacks, which is a common function of insurance in many other contexts. I abstract away from these considerations by assuming q remains independent of the decision to get insured.

Credibility of adversaries. Finally, it is important that the firm does not suffer harm b if it pays the ransom offer. The underlying assumption is that the ransomware group credibly commits not to further harm the firm if the ransom is paid, which is plausible due to reputational concerns of the large ransomware groups that dominate this space. The analysis of ransom negotiation transcripts by Cong et al. [2025] offers strong evidence of such reputational concerns.

3.2 Equilibrium with observed contracts

Bargaining subgame

First, I derive a firm's best response to a ransom offer r. An uninsured firm will accept a ransom offer r if and only if $r \leq b$. An insured firm with insurance terms M_r, M_b will accept a ransom offer r if and only if $r - M_r \leq b - M_b$, where $b - M_b$ is the net harm from accepting the ransom demand. Thus, when facing a firm with insurance contract (M_r, M_b), a best-response[9] for the adversary is to make a TIOLI offer:

$$BR^A(M_r, M_b) = r(M_r, M_b) := \max\{b + M_r - M_b, s\} \tag{1}$$

This remains true for the case of an uninsured firm, for which $M_r = M_b = 0$. If the firm's maximum willingness to pay is below the hacker's outside option s, the hacker demands a ransom high enough to induce rejection.

Proposition 1. *Given an insurance contract* (M_r, M_b)*: the following is true in any subgame*

---

[9] For $r(M_r, M_b) > s$, the best response is unique. In general, the best-response is to ask for any ransom $r \geq r(M_r, M_b)$ when $r(M_r, M_b) = s$.

equilibrium[10] *of the bargaining game that begins after the successful breach of a firm with insurance contract* (M_r, M_b).

1. *If* $r(M_r, M_b) > b - M_b + M_r \iff s > b - M_b + M_r$, *the firm finds it strictly optimal to reject. The adversary earns* s *and the firm suffers harm net of insurance payments* $(b - M_b)$.

2. *If* $r(M_r, M_b) = b - M_b + M_r \iff s \leq b - M_b + M_r$, *the firm is indifferent between accepting and rejecting the offer. Regardless of decision, the harm net of insurance payments is equal to* $b - M_b$.

Thus, for any insurance contract, the insured firm's equilibrium monetary loss when suffering a breach is $(b - M_b)$. *For an uninsured firm, it is* b.

By the above result, the expected utility of an insured firm does not directly depend on M_r, since the adversary fully adjusts the ransom demand to extract M_r:

$$U^I(p, M_r, M_b) = (1 - q)\, u(w - p) + q\, u(w - p - (b - M_b)) \tag{2}$$

The expected utility of an uninsured firm is:

$$U^N = (1 - q)\, u(w) + q\, u(w - b) \tag{3}$$

Taking into account how the choice of insurance contract shapes equilibrium outcomes in the bargaining subgame, the monopolist insurer selects the profit-maximizing contract to offer.

3.3 Monopolist insurer

I assume there is symmetric information between firms and insurer and that the insurer makes a TIOLI offer of an insurance *contract* to firms. The risk neutral monopolist insurer chooses a contract comprising of a premium p and terms M_r, M_b to maximize revenue net of insurance payments, subject to the firm being indifferent between buying insurance and not doing so.

$$\max_{p, M_r, M_b} \mathbb{E}\Pi(M_r, M_b, p) = \, p - \mathbb{E}\Pi(\text{insurance payment})$$

$$\text{subject to: } IC^{\text{hacker}}, IC^{\text{firm}}, IR^{\text{firm}}$$

---

[10] There exists a multiplicity of subgame equilibria that only differ in the hacker's ransom demand in the case $r(M_r, M_b) < s$ in which the hacker wants the demand to be rejected. Across every one of these subgame equilibria, payoffs of both parties are constant.

The IC constraints of firm and hacker refer to the strategies of those players in the *bargaining* subgame, whereas the IR constraint of the firm refers to its incentive to buy insurance relative to the outside option of staying uninsured. By Proposition 1 and the firm's risk aversion, the profit-maximizing contract must offer $M_b^* = b$, which provides the firm with full insurance, for any value of M_r. By offering full insurance, the insurer simultaneously reduces the firm's willingness to pay ransom to $b - M_b + M_r = M_r$. Thus, the business interruption component of insurance offers both *commitment* and *risk-sharing* value to insured firms.

In turn, the comparison between M_r and the adversary's value for monetizing the breach outside of ransom payments, s, determines whether the equilibrium features acceptance or rejection of the ransom offer, and thus, whether the expected insurance payment of the monopolist is $q\, M_r$ or $q\, M_b$.

The monopolist thus chooses between two candidate profit-maximizing contracts: one that induces ransom acceptance and one that induces rejection. It is clear that if the monopolist wants to induce ransom acceptance, the optimal contract offers $M_r^* = s$, and induces equilibrium ransom of $r(b, s) = s$. This guarantees that the ransom is as low as possible, subject to the adversary not exploiting the outside option. An even lower value of M_r would induce the adversary to take advantage of their outside option, so the monopolist is effectively "bribing" the adversary just enough to not cause harm b to the firm. A higher value of M_r would increase the insurer's payout without raising the premium it can extract from the firm.

To compare profits between the two candidate optimal contracts, first note that between the two cases, p^* will be the same, given by the solution to the binding IR constraint, $u(w - p^*) = U^N$. This is because both candidate contracts offer full insurance to the insured firm and the value of the outside option of the firm, U^N, is the same across the two scenarios. The comparison of profits then simply boils down to a comparison of expected insurance payouts, $q\, s$ and $q\, b$.

Proposition 2. *In the game with a monopolist insurer and observed contracts:*

- *If $b > s$, in the unique SPNE, the insurer offers $M_b = b$, $M_r = s$, p^*. In equilibrium all firms purchase insurance. The adversary makes offer $r(s, b) = s$ to an insured firm, and firms accept the ransom offer. Off-path, adversaries make ransom offer $r(0, 0) = b$ to an uninsured firm, and that offer is accepted.*

- *If $b < s$, there is a continuum of payoff-equivalent SPNE, which on the equilibrium path*[11] *differ only in the value of $M_r^* \in [0, s]$. In every equilibrium, all firms purchase insurance. The adversary makes offer $r(M_r, b) = s$ to an insured firm, and firms reject the ransom*

---

[11] If $b < s$, equilibria can also differ in the ransom demand, as explained in the footnote of Proposition 1, but the payoffs are constant across equilibria.

offer. Off-path, adversaries make ransom offer $r(0,0) = s$ to an uninsured firm, and that offer is rejected. Across every equilibrium, premium p^* is the same.

In any equilibrium, p^* is set to make the firm indifferent between purchasing insurance and staying uninsured. In any equilibrium, adversaries earn expected payoff of qs.

Benchmark: No insurance market. Consider the benchmark equilibrium in which there is no active insurance provider, and all active firms are successfully breached with probability q. I denote the expected utility of firms in the benchmark by U^0 and the corresponding hacker payoff by π^0. Maintaining the TIOLI assumption for hacker ransom demands yields the following:

Lemma 1 (No insurance market)**.** *Without an active insurance market there is a unique equilibrium. Hackers earn expected payoff $\pi^0 = q \max\{s, b\}$ and firms' expected utility is $U^0 = q\, u(w-b) + (1-q)\, u(w)$.*

At this point we can ask how equilibrium changes relative to the case in which the insurance market does not exist, in particular, how the welfare of firms and adversaries change relative to π^0, U^0. Answering that question will be key to addressing policy concerns about the existence of markets for ransomware insurance.

The equilibrium expected utility of firms in the absence of an insurance market, U^0 is the same as the equilibrium expected utility of uninsured firms in the game with a monopolist insurer, U^N. This is a direct consequence of the observed contracts assumption.[12] In equilibrium of Proposition 2, firms are indifferent between buying insurance and not doing so, $U^* = U^I = U^N$. Thus, firms' welfare is unaffected by the entry of an insurance monopolist who fully extracts their equilibrium surplus. On the other hand, adversaries earn revenue $\pi^{\text{obs}} = qs \leq \pi^0 = q \max\{s, b\}$ and are thus (weakly) worse off in equilibrium relative to the absence of an insurance market.

Corollary 1. *Under the presence of the monopolistic insurer, the adversary's expected equilibrium payoff is $\pi^{obs} = qs$, in any equilibrium. The presence of the monopolistic insurer strictly reduces the hacker's payoff relative to the case in which insurance provision is unavailable if $b > s$ and leaves it unchanged otherwise. Firms' welfare is unaffected by the presence of the monopolist, $U^{obs} = U^0$.*

4 Extensions

In this section, I verify the robustness of the baseline results by relaxing key assumptions regarding the bargaining process, the structure of the insurance market, and the entry of adversaries.

---

[12] And as we shall see later, of the assumption that there is no endogenous margin of hacker participation.

For brevity, I only consider the case of $b > s$.

4.1 Ransom determined via Nash Bargaining

I first show that the results extend to the case of Nash Bargaining between hackers and victims over the ransom payment, with bargaining power of the adversary given by $\beta \leq 1$. If ransom is paid in equilibrium, the value is given by the Nash Bargaining Solution $r^{NB}(M_r, M_b) = (1-\beta)s + \beta(M_r - M_b + b)$. The baseline corresponds to $\beta = 1$. As in the baseline, it remains true for any β that ransom is paid in the bargaining stage if and only if it is the efficient outcome (given the insurance contract), i.e., if $M_r - M_b + b \geq s \iff r^{NB}(M_r, M_b) \geq s$.

Proposition 3. *Fix $b > s$: for any $\beta \in (0,1)$, all firms purchase insurance in the unique equilibrium and the equilibrium contract involves $M_r^* = s, M_b^* = b$ and p^* such that the insurer extracts all surplus of firms. For any β, the adversary's equilibrium payoff in the bargaining stage is equal to their outside option, s. The equilibrium insurance premium is increasing in β.*

The equilibrium premium is increasing in β, because the ransom paid by uninsured firms off the equilibrium path is increasing in hackers' bargaining power.

4.2 Competitive insurance market

Replacing the monopolist with a perfectly competitive market changes the distribution of surplus between firms and insurers but leaves the equilibrium M_r, M_b terms unchanged. Assume that a continuum of identical insurance providers operate in the market, and they simultaneously post their contracts at the beginning of the game. The insurers are not capacity constrained, hence all firms purchase insurance from the seller whose contract offer yields the highest expected utility. In equilibrium, each insurer will offer the unique contract that maximizes insured firms' expected utility, subject to the break-even constraint.

In Proposition 2, the monopolist insurer offers the terms (M_r, M_b) that maximizes the customer's expected utility, by minimizing the payment to the adversary and also offering full insurance.[13] Thus, the same (M_r, M_b) is also offered in competitive equilibrium with homogeneous insurers. The difference is that the premia charged will equal expected insurance payments, so that insurers just break even.

Proposition 4. *If $b > s$, the game with competition between identical insurance firms has a unique subgame perfect equilibrium in symmetric strategies. All insurers set $M_r^* = s, M_b^* =$*

---

[13] Note that the (M_r, M_b) pair that maximizes the expected utility of insured firms does not depend on the insurance premium charged.

$b, p^* = qs$. All firms buy insurance and accept the ransom offer $r(s,b) = s$. Off-path, adversaries make offer $r(0,0) = b$ to uninsured firms, who accept the offer.

In the unique equilibrium, insurers in the perfectly competitive market are only compensated for the marginal cost of offering an additional insurance contract, $q\,M_r$. Since M_b is never paid out in equilibrium, the commitment value it provides is supplied at zero marginal cost. This implies that in the perfectly competitive equilibrium, risk neutral insurers are not compensated for the commitment value of insurance, but only for the risk-sharing component. For that reason, firms are strictly better off relative to the benchmark equilibrium without an insurer present, i.e. $U^* > U^0$.

4.3 Endogenous participation of adversaries

We can also verify that the equilibrium insurance terms (M_r^*, M_b^*) of Proposition 2 remain the ones offered in equilibrium if we extend the model to allow for endogenous participation of adversaries. Assume that each one among the continuum of adversaries has a participation cost c drawn from cdf F and that adversaries participate if and only if their expected revenue exceeds their entry cost. I assume adversaries do not observe the true contract offered to firms at the time they make their participation decision, rather only at the time they make their ransom demand.[14] This implies that the insurer does not directly influence hacker participation by the choice of insurance contract, but in equilibrium, participation will depend on the contract offered.

For given mass of participating adversaries $m \leq 1$, each firm is breached with probability $q\,m$ and according to Proposition 1 suffers net harm $(b - M_b)$. By the logic of Proposition 2, the unique equilibrium contract involves $M_r^* = s$, $M_b^* = b$, and we obtain the following:

Lemma 2. *If $b > s$, there is a unique equilibrium in the game with endogenous participation of adversaries. In equilibrium, ransom offers to insured firms are $r^* = s$ and firms accept. Expected revenue of active adversaries is $\pi^* = q\,s$ and participation given by $m^* = F(\pi^*)$. The contract offered is ($M_r^* = s$, $M_b^* = b$, p^*), where p^* that makes firms indifferent between buying insurance and not doing so.*

Since this is the contract that minimizes the expected payoff of each active hacker, it is also the contract that *minimizes participation.* The presence of an insurer reduces adversarial activity through an additional margin relative to the benchmark. This has another implication:

---

[14] As discussed in my introduction, I assume hackers obtain this information by gaining access to a particular firm's network and files, which is consistent with individuals not having this information at the time they decide to become ransomware hackers. This is also consistent with my assumption that firms cannot be targeted on the basis of having purchased insurance.

in equilibrium, the outside option of firms, U^N is greater than the benchmark utility U^0, since the mass of active hackers is smaller in the insurance equilibrium. In equilibrium, $U^I = U^N$, and insured firms are made strictly better off relative to the no-insurance benchmark, even if the monopolist insurer has full bargaining power vis-a-vis insurance buyers.

5 Unobserved Contracts

In this section, I explore the extent to which the results of Proposition 2 are sensitive to the assumption that adversaries can observe the precise insurance contract a breached firm has signed. In particular, I assume that the adversary cannot observe the contract terms offered by the monopolist and cannot observe whether a breached firm has signed a contract or not. In the main text, I present the main arguments for finding pure-strategy equilibria and defer presentation of the complete argument to the Appendix.

The adversary's best-response.

In any equilibrium, all firms purchase insurance, hence I restrict attention to this case. If the hacker anticipates that contract terms M_r, M_b are offered, the optimal ransom demand is again:

$$BR^A(M_r, M_b) = \begin{cases} r : r \geq s, & \text{if } r(M_r, M_b) < s \\ r(M_r, M_b), & \text{if } r(M_r, M_b) > s \end{cases}$$

where $r(M_r, M_b) := b + M_r - M_b$ is the maximum ransom that an insured firm is willing to pay.[15]

The insurer's best response.

Unlike the case of observed contracts, the insurer does not directly affect the adversary's offered ransom with the choices of contract terms and takes r as given when choosing the insurance contract to offer. For given ransom demand $r \geq s$, the insurer decides whether to set contract terms that induce acceptance or rejection of this offer by insured firms. If it wants to induce acceptance, it is optimal to fully insure the firm by offering $M_r(r) = r$, and the firm accepts the offer if and only if $M_b \leq b$. Since M_b is not paid in equilibrium and also does not affect the adversary's ransom offer, this constraint is *slack*, and the insurer's best-response correspondence

---

[15] In other words, a best response of the adversary is again $r(M_r, M_b) = \max\{r^{max}(M_r, M_b),\ s\}$. Similar to the case of observed contracts, the best-response of the adversary is not uniquely pinned down if he wants to induce rejection of the offer.

is:

$$BR^I(r) = \begin{cases} M_r = r, \\ M_b \in [0, b], \\ p : U^I(M_r, M_b, p; r) = U^N(r) \end{cases}$$

Firms' outside option of not buying insurance yields utility $U^N(r) = (1-q)u(w) + qu(w-r)$, which in the game with unobserved contracts now depends on the ransom demand r.

Equilibrium.

In equilibrium the insurer's and adversaries' conjectures must be correct, and putting the two best responses together: $M_r(r) = r$ and $r = b + M_r - M_b$ imply that $M_b = b$ is the only value of M_b consistent with an equilibrium in which ransom is paid. Finally, it must be that the adversary earns payoff higher than s, but that is not sufficient to uniquely pin down M_r. In the Appendix, I prove the following:

Proposition 5. *In the game with unobserved contracts there exists a multiplicity of pure-strategy SPNE. In every equilibrium, firms buy insurance.*

- *If $b > s$: there exists a continuum of equilibria in which the firm accepts the ransom offer. These equilibria can be indexed by the equilibrium value of M_r. In every such equilibrium, the insurer sets $M_b^* = b$, $M_r^* \in [s, b]$, and premium to fully extract firm's surplus, $U^I(p^*, M_r^*, M_b^*) = U^N(r^*)$. Adversaries make ransom offers $r^* = M_r^*$. Firms accept the ransom offer. Off-path, if a firm does not purchase insurance and is breached, it accepts the ransom offer, too.*

- *If $b < s$: the insurer sets $M_b^* = b$, $M_r^* \in [0, s]$, and premium to fully extract firm's surplus, $U^I(p^*, M_r^*, M_b^*) = U^N(r^*)$. Adversaries make ransom offers $r^* = s$. Firms reject the ransom offer. Off-path, if a firm does not purchase insurance and is breached, it rejects the ransom offer, too.*

For every value $M_r \in [s, b]$, there exists an equilibrium in which all firms buy insurance and breached firms accept the ransom offer $r = M_r$. Given the above contract terms, the adversary optimally asks for ransom $r = M_r \geq s$, and $r(M_r, M_b) = M_r = r$ so that the firm is *indifferent* between acceptance and rejection of the demand. The insurer cannot affect the ransom offer and best responds by offering a contract that provides full insurance, $M_r = r$. The *strategic complementarity* between M_r and r is the source of the equilibrium multiplicity: higher ransom demands increase the amount of M_r required to provide full insurance and higher M_r induces higher ransom demand by increasing insured firms' willingness to pay ransom, $r^{max}(M_r, M_b)$.

In the Appendix, I also show that under $b > s$, there also exist equilibria with *rejection* of ransom, in which the adversary makes "unreasonably" large demands. I offer arguments based on *trembling* and robustness to *counteroffers* to disregard these equilibria and focus on the ones presented in this section.

5.1 Welfare

Welfare comparison across equilibria. To compare welfare of firms across the multiple equilibria under $b > s$, we can simply compare the off-path expected utility uninsured firms in each equilibrium. Under unobserved contracts, an uninsured firm would be worse off in equilibria with higher M_r and higher r since it would receive the same ransom offers as insured firms. So, the equilibrium in which firm's utility is maximized is the one with $M_r = s$. However, the *opposite* holds for the insurer's expected profit. Higher ransoms allow for higher premiums that more than offset the increased expected payouts. In the Appendix, I prove the following:

Proposition 6. *For $b > s$: in the game with unobserved contracts, equilibrium expected profit of the monopolist insurer is higher in equilibria with higher M_r. In contrast, equilibrium expected utility of firms is higher in equilibria with lower M_r.*

In light of the above result, we can think of the insurer and adversaries as playing a coordination game; they are both better off when adversaries ask for high ransom and insurance against high ransom levels is provided, but that is at the expense of firms.

Comparison with observed contracts case. Focusing on the case of $b > s$, notice that there exists an equilibrium with $M_r^* = s,\ M_b^* = b$, which are also part of the contract offered in the (unique) equilibrium under observed contracts. Even though the ransom offer made to the insured firms is the same across the two equilibria, the *premium* charged by the insurer is lower under unobserved contracts. Under unobserved contracts, uninsured firms, off the equilibrium path, benefit from a *positive externality* from insured firms' decision to buy insurance, because the adversary's optimal ransom demand is lower for firms that are insured with the equilibrium contract. This externality implies that even if the monopolist extracts the *risk-sharing* value of insurance, firms nevertheless appropriate the entire *commitment* value it offers.

In the case of observed insurance contracts, there is no such externality because ransom offers are made on the basis of each firm's true contract. Since in any of the aforementioned equilibria firms are indifferent between purchasing insurance and not doing so, firms are better off *in any* equilibrium with unobserved contracts, relative to the unique equilibrium with observed contracts. In contrast, adversaries are worse off in the case of observed contracts, in which the insurer acts like a Stackelberg leader that pre-commits to the firm's bargaining position. I combine this discussion with the result in Corollary 1 to obtain the following:

Corollary 2. *Under $b > s$: in every equilibrium with unobserved contracts, firms are better off than both (a) under the unique equilibrium with observed contracts and (b) the unique equilibrium without an active insurer,* $U^{unobs}(M_r^*) \geq U^{obs} = U^0, \forall\ M_r^* \in [s, b]$. *Adversaries are better off than in the observed contracts case, but remain worse off than in the no-insurance benchmark,* $\pi^{obs} \leq \pi^{unobs}(M_r^*) \leq \pi^0, \forall\ M_r^* \in [s, b]$.

According to the above result, even if contracts are unobserved, the creation of an insurance market remains welfare-enhancing for firms and welfare-reducing for hackers across all equilibria. However, unlike in the case of observed contracts, the multiplicity of equilibria provides grounds for regulating insurance for ransom payments to ensure that equilibria with low M_r and thus low ransom transfers are selected. This observation is discussed in more detail in section 8.

Remark. The comparison of Propositions 2 and 5 echoes the results of Bolton and Scharfstein [1990], who find that unobservable contracts between a principal and an agent are less effective as pre-commitments than are observable contracts.[16] In their equilibrium with observed contracts, an investor (principal) can effectively use the financing contract it signs with a startup to deter predatory behaviour by an incumbent (who is analogous to the adversary in my context). But this is not possible in equilibrium when contracts are unobservable, because the principal has a *dominant strategy* when the contracts are unobserved by the incumbent, and it invites preying by the incumbent. In contrast, the provision of ransom insurance and the adversary's ransom demand of the adversary are strategic complements, leading to the multiplicity result I obtain in Proposition 5.

5.2 Random contract observability

A reader may wonder at this point whether the multiplicity of equilibria persists when the adversary observes the contracts with *exogenous* probability $k \in [0, 1]$. If a firm is not insured, the adversary observes nothing[17]. For interior k, insured firms that are breached face uninformed ransom demand r^u with positive probability $(1 - k)$, and with complementary probability informed ransom demand $r^i = \max\{M_r + b - M_b, s\}$. I make the natural assumption that the

---

[16] In their setting, an investor who chooses a financing contract for a startup faces a trade-off between providing strong incentives to the start-up company and encouraging predatory behavior by an incumbent competitor. The former requires the probability of funding continuation to be dependent on the start-up's first period performance. When the competitor observes contracts, the investor can deter predation by using a financing contract that is less sensitive to performance. In equilibrium, the investor commits to a lower reward for high-performance firms and deters preying.

[17] But firms obtain insurance in all equilibria, so that adversaries do not account for uninsured firms in their beliefs.

insurer cannot condition M_r, M_b on whether the adversary has discovered the contract or not. In the Appendix, I show that the multiplicity persists for every $k < 1$. For every $k \in [0,1]$, and for $r(k) := k\,s + (1-k)\,b$, any ransom $r^u = r^i \leq r(k)$ can be supported in equilibrium. Insured firms accept *both* informed or uninformed offers.

Proposition 7. *Assume $b > s$. For any value of k, there exists a multiplicity of equilibria in all of which insured firms accept the ransom offers they face. In equilibrium, $r^u = r^i = M_r^*$, with $M_r^* \in [s, r(k)]$, $M_b^* = b$, and p^* is set to make the insured firms indifferent.*

If $r^u > r(k)$, then the best response of the insurer is to induce *rejection* of the uninformed ransom demand, and no such equilibrium exists under $b > s$. Note that $\lim_{k\to 1} r(k) = s$ which corresponds to the case of observed contracts and $\lim_{k\to 0} r(k) = b$, which corresponds to the case of unobserved contracts. Greater values of k, i.e. greater probability that contracts are observed, monotonically reduce the extent of equilibrium multiplicity, but equilibrium multiplicity arises *for any* $k < 1$.

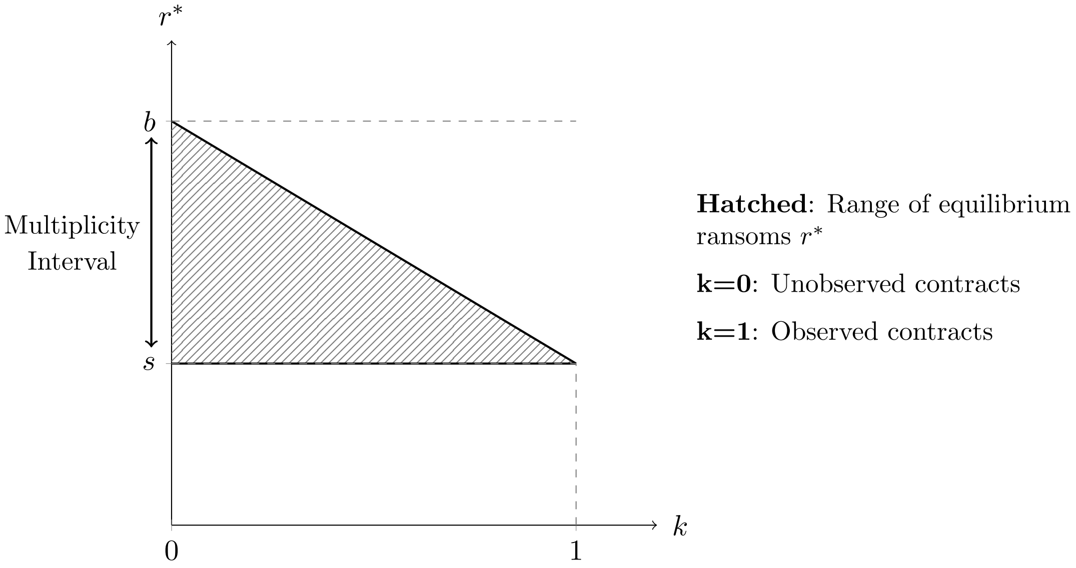

Figure 2: For any $k < 1$, the strategic complementarity between coverage and ransom demands sustains a multiplicity of equilibria. The upper bound of ransoms that can be supported in equilibrium monotonically approaches the hacker's outside option, s, as the probability of contracts being observed increases.

Remark. For any $k > 0$, $M_b = b$ is dominant for the insurer: At $k = 0$, even though $M_b = b$ is the sole equilibrium outcome, the insurer is *indifferent* across $M_b \in [0, b]$ in equilibrium (this was discussed in the equilibrium derivation of Section 5). This indifference the indifference is broken in favour of $M_b = b$ as soon as there is positive probability of contracts being observed.

6 Empirical Predictions

The above results, and in particular Propositions 2 and 5, generate empirical predictions for the case in which firms are not liquidity constrained.

First, the presence of the insurance provider does not change whether ransom is paid in equilibrium. Ransom is paid if and only if the business interruption to firms exceeds hackers' profit from monetizing stolen data, $b > s$, which is the same condition as in the absence of an insurance market. Second, conditional on breach, the occurrence of ransom payment is unaffected by whether firms are insured. Off the equilibrium path, uninsured firms pay ransom if and only if $b > s$, and the same is true for insured ones. Thus, the model predicts no difference in ransom payment rates between insured and uninsured firms.[18] Third, however, conditional on ransom payments occurring, uninsured firms pay higher amounts, similar to those observed in markets without active insurance providers. Fourth, when $b > s$, the equilibrium contract includes full coverage of rejection/interruption costs. The insurance market also offers insurance for ransom payments in equilibrium, $M_r > 0$, and in particular $M_r \geq s$. As shown in Proposition 7, as the probability that hackers observe insurance contracts increases, the equilibrium reimbursement limit converges toward s, the adversaries' outside option from bargaining breakdown.

In the following section, I extend the model to account for liquidity constrained firms and discuss the extent to which the above empirical predictions are affected.

7 Liquidity Constraints

An alternative reason why firms demand insurance for ransom payments may be liquidity constraints: firms that face ransom offer r may simply not have enough available funds to pay the adversary, even if they would find it optimal to do so. This is an additional reason why ransomware insurance, and insurance more generally, may be valuable, on top of risk aversion. It is plausible that liquidity constraints matter in the context of ransom payments. Firms must make funds available within a very short time frame, as well as understand how to make payments according to hackers' technical specifications, and while under pressure to contain the impact of the ongoing ransomware attack. As we will see, the presence of liquidity constraints will *qualitatively* affect how the equilibrium welfare of adversaries and firms compares across equilibria with and without an active insurance market. Additionally, this qualitative difference will depend on whether insurance policies are observed or not by the adversaries.

---

[18] There are no uninsured firms in equilibrium; this prediction would be generated by a simple extension of the model in which a fraction of (otherwise identical) firms simply never considers insurance, a natural assumption to make in a nascent market such as cyber and ransomware insurance.

I assume that every firm has access to liquidity $b > \ell > 0$ and an uninsured firm can only pay ransom r if $\ell \geq r$. On the other hand, if an insurance market exists, an insured firm can pay ransom demand r if $\ell + M_r \geq r$. I assume that liquidity constraints do not impact the firm's ability to pay for insurance. To keep things simple, assume that ℓ is a known scalar, known both to the insurer and the adversary. For brevity, I focus on the case of $b > s$.

7.1 Benchmark: without an insurance market

Given the above assumptions, in equilibrium without an insurance provider, firms pay to adversaries as much as their liquidity allows, *if* their liquidity suffices to compensate the adversary for not selling the stolen data. In equilibrium, if $\ell > s$, firms transfer ransom ℓ to adversaries, but if $\ell < s$, no ransom is paid: firms suffer harm b and adversaries sell the data for revenue $s \in (\ell, b)$.

Lemma 3. *Without an active insurance provider, there is a unique equilibrium. If $\ell > s$, firms pay ransom equal to their available liquidity, and payoffs are $\pi^0 = \ell$, and $U^0(\ell) = q\,u(w - \ell) + (1 - q)u(w)$. Otherwise, they pay do not pay ransom and suffer harm b. Payoffs are $\pi^0 = s$, and $U^0 = q\,u(w - b) + (1 - q)u(w)$.*

Thus, a successful adversary earns $\pi^0 = \max\{s, \ell\}$, increasing in ℓ. The firm obtains expected utility U_0, which increases in ℓ once ℓ exceeds s, and the firm no longer suffers b in equilibrium. But for $\ell \geq s$, U_0 becomes *decreasing* in ℓ, since the ransom paid is $r = \ell$.

Remark. For any liquidity $\ell < s$, firms would be better off if their liquidity increased to any level $\ell \in [s, b)$ that allows it to pay ransom and avoid the business interruption. This direct, positive impact of the ransom-paying option is a key mechanism in the paper of August et al. [2025].

7.2 Observed contracts

For the case of $b > s$, it is easy to verify that the optimal provisions $M_r^* = s$ and $M_b^* = b$ of Proposition 2 are again part of the profit-maximizing insurance policy: the insurer will sell a contract that depresses the ransom demand to the adversary's indifference point and also provides full insurance: the equilibrium ransom demand is $r^* = s < M_r^* + \ell$, so that the liquidity constraint *never binds* for insured firms. In equilibrium, adversaries earn s regardless of whether $s > \ell$.

Lemma 4. *When insurance contracts are observed, and $b > s$, there is a unique equilibrium and in equilibrium, ransom is paid for any ℓ. The insurance contract offered is $M_r^* = s$, $M_b^* = b$ and p^* such that $U^I = U^N$.*

Simple comparison with the previous Lemma reveals that adversaries are strictly *worse off* in the equilibrium with insurance if $\ell > s$. If stand-alone liquidity is high, the insurer strips the adversary of their ability to extract the firm's liquidity, forcing the ransom down to the hacker's outside option s. By the same argument of Corollary 1, firms are equally well off relative to the benchmark, under the presence of a fully extracting monopolist.

7.3 Unobserved contracts

I look for equilibria of the game with unobserved contracts in *pure strategies*, in which the firm pays ransom. The same arguments I appeal to in the proof of Proposition 5 apply again, and I select only equilibria in which $r \leq b$, i.e. equilibria in which ransom is paid in equilibrium. The logic that leads to multiplicity in Proposition 5 applies again, and in the Appendix, I show the following:

Proposition 8. *If $b > s$, there exists a continuum of pure-strategy equilibria in which the firm accepts the ransom demand made by adversaries. For every value of $M_r^* \in [s, b]$, there exists one such equilibrium and in every such equilibrium, $M_b^* = b$ and $r^* = M_r^*$. In every equilibrium, the insurance premium is set to make firms indifferent between buying insurance and not doing so. Off the equilibrium path, uninsured firms accept the ransom offer if and only if $\ell \geq r^*$.*

Welfare comparisons. Same as in the case of observed contracts, there can be no equilibrium in which the liquidity constraint binds for insured firms, because the monopolist finds it optimal to offer full insurance. For that reason, the value of liquidity ℓ does not affect which values of ransom insurance M_r^* and $r^* = M_r^*$ can arise in equilibrium.

However, the value of ℓ affects the off-path utility of firms if they are uninsured, and thus affects firms' expected utility U^{unobs} (via the equilibrium premium p^*) at the particular equilibrium (M_r^*, r^*) being played. Fixing the value of r^*, the welfare of uninsured firms jumps discontinuously as their liquidity reaches $\ell = r^*$ and remains constant on either side of r^*. Additionally, the value of liquidity affects the firm's and adversary's payoffs in the absence of an insurance market, hence affects how the presence of an insurance market changes their payoffs. Whether an adversary is better off or not depends on how the equilibrium payoff compares to that of the equilibrium without insurance supply, $\pi^0 = \max\{\ell, s\}$, and the answer will depend on *which* of the above equilibria is being played.

Corollary 3. *For the case of $b > s$.*

- *If $\ell < s$, $\pi^0 = s$ and an adversary becomes better off in* ***every*** *equilibrium of Proposition 8 relative to the equilibrium without an insurance market.*

- *If* $\ell > s$, $\pi^0 = \ell$ *and the adversary becomes (weakly) better off in an equilibrium of Proposition 8 if and only if* $M_r^* \geq \ell$.

If $\ell < s$, the firm's liquidity constraint is so severe that it cannot even compensate the hacker for the value of their outside option, and there is ransom rejection in the equilibrium without an insurance market, $\pi^0 = s$. Under the presence of an insurer, the revenue of the successful adversary is $r^* \geq s$, so that the adversary becomes (weakly) better off in every equilibrium. If $\ell > s$, then there is acceptance in the equilibrium without an insurer, adversaries are better off only in equilibria in which the amount of insurance for ransomware payments is greater than firms' initial liquidity. Since in such equilibria, $r = M_r$, this implies that adversaries are better off under the presence of an insurer *if and only if* the equilibrium ransom demanded is greater than firms' standalone liquidity ℓ.

Notice that this result comes in stark contrast to the case without liquidity constraints in Proposition 5: in that case, the adversaries are *never* better off under an active insurance market. In fact, this is the first model variant examined with an equilibrium in which hackers are better off relative to the no-insurance profit π^0.

How does the welfare of *firms* compare to U^0, i.e., to the case without an insurance market? With a monopolist insurer who extracts all surplus, firms' expected utility U^* is the same as the off-path expected utility of uninsured firms, U^N. Because insurance policies are unobserved, off-path, firms face the same ransom as on the equilibrium path, $r^* = M_r^* \in [s, b]$ and thus accept if and only if they have sufficient liquidity, i.e. if $\ell > M_r^*$.

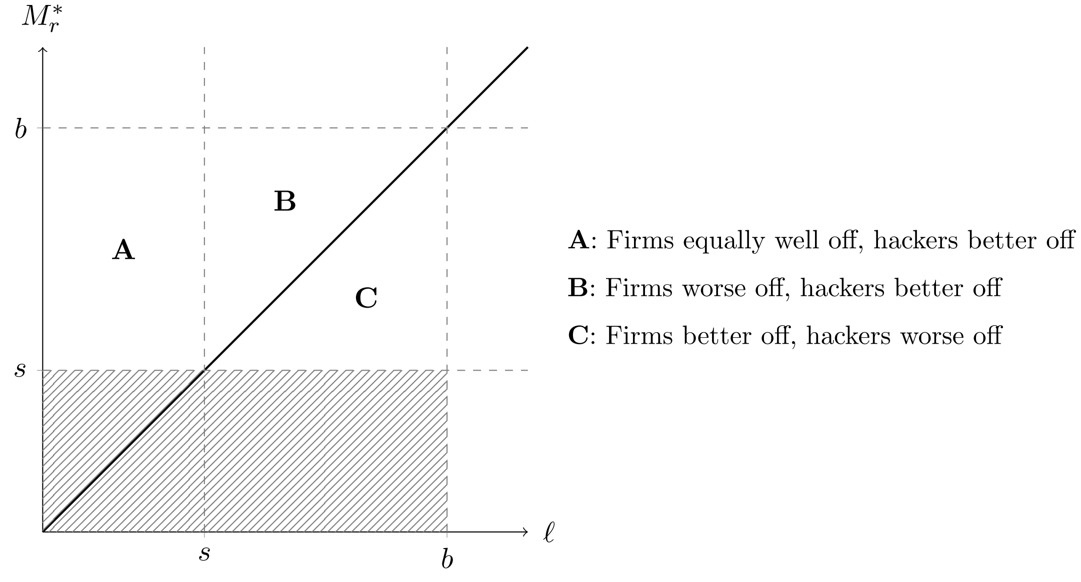

Figure 3: The comparison between the welfare of firms and hackers in an equilibrium with unobserved contracts relative to the equilibrium without insurance depends on firms' liquidity, ℓ and the equilibrium value of insurance for ransom payments, M_r^*.

Figure 3 shows the three possible cases for the comparison of U^N and U^0. In region A, firms are severely liquidity constrained in the absence of insurance, $\ell < s$, and breached firms suffer harm b.[19] With an active insurance market, off path, uninsured firms also cannot pay any ransom in $[s, b]$ and suffer harm b. Hence $U^N(M_r^*) = U^0$ and $U^I(M_r^*) = U^0$, too, for any M_r^*. Firms are equally well off, but hackers are strictly better off in equilibria with $r^* = M_r^* > s$.

In region C, $M_r^* \in [s, \ell]$ and $\ell > s$, and firms are **better off**, i.e. $U^I = U^N > U^0$. The reason is that the equilibrium ransom is *lower* than that of the equilibrium without insurance: the commitment value of insurance contracts, as measured by the difference $(M_b - M_r)$, which acts to reduce equilibrium ransom demands, is sufficiently large relative to compensate for the increase in insured firms' liquidity.

Crucially, the presence of liquidity constraints creates the possibility of a welfare-reducing outcome for firms. This is the case in region B, defined by $M_r^* \in [\ell, b]$ and $\ell \in [s, b)$, and firms are **worse off** in the equilibrium with an insurance provider,[20] $U^I = U^N < U^0$. The outside option U^N becomes lower than U_0 in equilibria with $M_r^* > \ell > s$. When liquidity is not too low, $\ell > s$, breached firms face ransom ℓ in the absence of an insurance market. In the insurance equilibrium with $M_r^* \in [\ell, b]$, insured firms liquidity constraint is relaxed, and they pay higher ransom than ℓ, thus the ransom demanded by uninsured firms off-path is greater than ℓ, too. Off-path, uninsured firms cannot afford to pay the ransom and suffer business interruption $b > \ell > s$. Hence $U^N < U^0$.

In Proposition 5, we already identified equilibria in which M_r^* is too high relative to what minimizes the welfare of adversaries. However, the case of liquidity constraints is the only in which there exist equilibria with higher payoffs than π^0 for adversaries. This can only occur with unobserved contracts: the insurer does not fully control the commitment value of insurance contracts for the firm in its bargaining with the hacker, since it cannot affect hackers' conjectures, and unlike the case of observed contracts, there exist equilibria with $M_r^* > \ell > s$.

8 Managerial and Policy Implications

The main mechanism underlying the results of this paper is that cyber insurance shapes the firm's bargaining position relative to ransomware adversaries. Coverage for rejection/interruption costs strengthens refusal credibility (commitment channel), while high ransom reimbursement expands feasible transfers (liquidity channel). Recognizing these additional channels on top

---

[19] With an active insurer, this is the worst possible off-path outcome for uninsured firms, hence we immediately see that $U^N \geq U_0$ and firms must be weakly better off.

[20] The result is reminiscent of the main one in Balasubramanian [2021]. In his model, the commitment value of insurance is absent, and by relaxing the liquidity constraint insurance markets *always* lead to higher equilibrium ransom demands.

of risk transfer yields insights for optimal insurance design as well as guidance for regulatory interventions.

8.1 Optimal contract design for insurers

Across regimes, a robust recommendation for profit-maximizing insurance contracts is that business interruption insurance, M_b, and ransom reimbursement, M_r are **not** pooled into a single undifferentiated limit. In particular, insurers should should provide full coverage of rejection/interruption costs, $M_b = b$. This requires, prior to contracting, that insurance companies assess the extent of damage that a ransomware attack can threaten to cause firms, and provide credible insurance against such harms.

Additionally, optimal contracts should never induce $M_r > b$. Depending on the technological capabilities of adversaries, when it is more likely that insurance contracts are discovered at the bargaining stage, the profit maximizing reimbursement M_r should approach the attacker's outside option, to the extent that this can be estimated. This value will mainly depend on (1) the type of data they capture and how much they can profit from exploiting it or selling it and (2) potential direct benefits of causing harm to the firm, which should be particularly relevant for foreign state-backed cyber attacks. Excessively high ransom limits can sustain high-transfer equilibria.

To the extent that it is possible, insurers should try to provide insurance contracts that can be used by firms to credibly reveal $M_b = b$ to adversaries act as a commitment device. Doing so would be both profitable to insurance firms and welfare-improving for firms.

8.2 Regulatory interventions

The key regulatory insight from Figure 3 is that insurance for ransom payments should not exceed firms' baseline capacity to pay the ransom. In other words, equilibria in which firms are insured for making ransom payments are welfare-improving, as long as they could also afford to make those same ransom payments without insurance, i.e. as long as $r^* = M_r < \ell$. When this is violated, it is evidence that the equilibrium is in region A or B of Figure 3, and the presence of insurance markets is making adversaries better off.

More generally, and regardless of liquidity constraints, a robust result is that the equilibrium with $M_r^* = r^* = s$ is simultaneously firm-best and hacker-worst.[21] To the extent that contracts remain unobserved with high probability at the bargaining stage, there may be a multiplicity of

---

[21] Note that this is true regardless of whether $\ell > s$ holds. If $\ell < s$, firms suffer business interruption in equilibrium without an insurance provider, and a contract with $M_r = s$ and $M_b = b$ allows them to pay the minimum required ransom and raise their welfare.

equilibria and in that case *regulatory caps* that anchor M_r to s would benefit firms and reduce adversaries' profits.

9 Conclusion

In this paper, I have studied an equilibrium model of the market for ransomware insurance under the presence of strategic ransomware adversaries. The insurance contract specifies different coverage amounts for business interruption and ransom payments made by firms. Purchasing insurance with higher coverage for business interruption provides *commitment value* to the firm in bargaining with the adversary, and insured firms pay *lower* ransom than in the case without insurance markets. In equilibrium, ransom is paid so that firms avoid suffering severe business interruption and coverage for ransom payment is used to provide full insurance to firms.

I study the role of adversaries' ability to observe insurance contracts. With observed contracts, the insurer can directly manipulate ransom demands and in the unique equilibrium, ransom payments become equal to adversaries' outside payoff from selling firms' stolen data. With unobserved insurance contracts, the insurer cannot fully control the commitment value of insurance contracts for the firm, since it cannot affect hackers' conjectures. But even so, the equilibrium insurance contract still provides commitment value for the firm, and this value is greater in equilibria with *lower* coverage for ransom payments. From a policy perspective, these results suggest regulatory intervention that limit coverage for ransom payments can raise welfare and harm adversaries, but does *not* justify banning insurance of ransom payments or the payments themselves. Importantly, in any equilibrium, the welfare impact of insurance markets is always positive for firms and negative for hackers.

This conclusion can change if contracts are unobserved and firms face *liquidity* constraints. Because insurance relaxes those liquidity constraints, there are equilibria in which ransom is higher relative to the pre-insurance equilibrium, and adversaries are better off. In these equilibria with a monopoly insurance market, firms are also made worse off, and the policy implication of the previous section is reinforced: coverage for ransom payments should be limited, and in particular, not exceed firms' stand-alone liquidity. The latter ensures that firms' welfare is higher in equilibria with an active insurance provider.

Finally, this model can be used to study the interaction between *privacy* and *data-breach* regulations, like the EU GDPR, with firms' incentives in the bargaining subgame. Anecdotally, hackers understand that firms stand to face regulatory scrutiny if their customers' data is leaked and use this as leverage to raise the ransom demanded. This should influence the optimal design of policy concerning both data breaches and ransomware insurance markets.

References

G-J van Rooyen. What is double extortion ransomware?, jun 2024. URL `https://www.techmonitor.ai/technology/cybersecurity/double-extortion-ransomware?cf-view`.

National Cyber Security Centre. Global ransomware threat expected to rise with ai, may 2024. URL `https://www.ncsc.gov.uk/news/global-ransomware-threat-expected-to-rise-with-ai`.

The Guardian. M&s says some personal data was taken in cyber-attack, 2025. URL `https://www.theguardian.com/business/2025/may/13/m-and-s-personal-data-cyber-attack-marks-spencer-card-passwords`. Report on Marks & Spencer data breach: customer names, addresses and order histories exposed; card details and passwords unaffected.

Financial Times. M&s cyber insurance payout to be worth up to £100mn, 2025. URL `https://www.ft.com/content/723b6195-1ce7-4b5f-94f5-729e9152c578`. Coverage of Marks & Spencer's potential £100m cyber-insurance claim and implications for premiums.

Munich Re. Cyber insurance: Risks and trends 2025, 2025. URL `https://www.munichre.com/en/insights/cyber/cyber-insurance-risks-and-trends-2025.html#379179748`. Accessed: 2025-07-24.

Marsh. Ransomware, 2025. URL `https://www.marsh.com/en/services/cyber-risk/expertise/ransomware.html`. Insights into Marsh's ransomware readiness and cyber-insurance expertise.

Acronis. The role of cybersecurity insurance in ransomware protection. `https://www.acronis.com/en-eu/blog/posts/cybersecurity-insurance-role-in-ransomware-protection/`, 2025. Accessed: 2025-07-24.

Coalition, Inc. How hackers leverage insurance details in ransomware attacks, 2025. URL `https://www.coalitioninc.com/en-gb/blog/cyber-insurance/how-hackers-leverage-insurance-details-in-ransomware-attacks`. Accessed January 2026.

Policyholder Pulse. Ransomware hardbit and cyber insurance policies, 2023. URL `https://www.policyholderpulse.com/ransomware-hardbit-cyber-insurance-policy/`. Accessed January 2026.

Will Cong, Campbell Harvey, Daniel Rabetti, and Zong-Yu Wu. An anatomy of crypto-enabled cybercrimes. *Management Science*, 71(4):3622–3633, 2025. doi: 10.1287/mnsc.2023.03691. URL `https://doi.org/10.1287/mnsc.2023.03691`.

Matthew Gooding. Is ransomware insurance fuelling the ransomware boom?, jul 2024. URL `https://www.techmonitor.ai/technology/ransomware-insurance?cf-view`.

Euronews. Cybercrime: Insurance giant axa to stop covering ransomware payments in france, May 2021. URL `https://www.euronews.com/2021/05/07/cybercrime-insurance-giant-axa-to-stop-covering-ransomware-payments-in-france`. Accessed: 2025-07-24.

Financial Times. The ransomware battle is shifting — so should our response, 2024. URL `https://www.ft.com/content/3b172a2a-4be5-4ef4-87cb-7fdcdee2ad99`. Opinion piece on strategic shifts in ransomware defense.

Anirudha Balasubramanian. Insurance against ransomware, February 2021. URL `https://ssrn.com/abstract=3846111`. Available at SSRN: `https://ssrn.com/abstract=3846111`.

Anna Cartwright, Edward Cartwright, Jamie MacColl, Gareth Mott, Sarah Turner, James Sullivan, and Jason R.C. Nurse. How cyber insurance influences the ransomware payment decision: Theory and evidence. *The Geneva Papers on Risk and Insurance–Issues and Practice*, 48(2): 300–331, 2023. doi: 10.1057/s41288-023-00288-8.

Tom Meurs, Edward Cartwright, Anna Cartwright, Marianne Junger, Raphael Hoheisel, Erik Tews, and Abhishta Abhishta. Ransomware economics: A two-step approach to model ransom paid. In *2023 APWG Symposium on Electronic Crime Research (eCrime)*, pages 1–13, 2023. doi: 10.1109/eCrime61234.2023.10485506.

Toni Ahnert, Michael Brolley, David Cimon, and Ryan Riordan. Cyber risk and security investment. Staff Working Papers 22-32, Bank of Canada, Jul 2022. URL `https://ideas.repec.org/p/bca/bocawp/22-32.html`.

Terrence August, Duy Dao, Marius Florin Niculescu, and Kihoon Kim. The impact of cryptocurrency on cybersecurity. *Management Science*, 2025. Forthcoming.

Aron Laszka, Samaneh Farhang, and Jens Grossklags. On the economics of ransomware. In *Decision and Game Theory for Security: 8th International Conference, GameSec 2017, Proceedings*, Vienna, Austria, 2017. doi: 10.1007/978-3-319-68711-7_20.

Anna Cartwright and Edward Cartwright. Ransomware and reputation. *Games*, 10(2):1–14, June 2019. URL `https://ideas.repec.org/a/gam/jgames/v10y2019i2p26-d238459.html`.

Terrence August, Duy Dao, and Marius Florin Niculescu. Economics of ransomware: Risk interdependence and large-scale attacks. *Management Science*, 68(12):8979–9002, 2022. doi: 10.1287/mnsc.2022.4300.

Itay P. Fainmesser, Andrea Galeotti, and Ruslan Momot. Digital Privacy. *Management Science*, 69(6):3157–3173, June 2023. doi: 10.1287/mnsc.2022.4513. URL `https://ideas.repec.org/a/inm/ormnsc/v69y2023i6p3157-3173.html`.

Alexandre de Cornière and Greg Taylor. A Model of Information Security and Competition. *Marketing Science*, forthcoming, 2024. URL `https://pubsonline.informs.org/doi/10.1287/mksc.2023.0513`.

Michael L. Katz. Game-playing agents: Unobservable contracts as precommitments. Economics Working Papers 91-172, University of California at Berkeley, Jul 1991. URL `https://ideas.repec.org/p/ucb/calbwp/91-172.html`.

Patrick Bolton and David S. Scharfstein. A theory of predation based on agency problems in financial contracting. *The American Economic Review*, 80(1):93–106, 1990. URL `http://www.jstor.org/stable/2006736`.

Woodruff Sawyer. The new hacker playbook: Weaponizing the sec's cyber disclosure rules. Online, 2024. URL `https://woodruffsawyer.com/insights/sec-cyber-disclosure-rules`. Accessed: 2025-08-14.

Rainer Böhme and Galina Schwartz. Modeling cyber-insurance: Towards a unifying framework. In *Workshop on the Economics of Information Security (WEIS)*, Harvard University, Cambridge, MA, June 2010. Working paper; slides available online.

A Baseline Model

A.1 Proof of Proposition 2

The risk neutral monopolist solves:

$$\max_{p,M_r,M_b} \mathbb{E}\Pi(M_r, M_b, p) = p - \mathbb{E}\Pi(\text{insurance payment})$$

$$\text{subject to: } IC^{\text{hacker}}, IC^{\text{firm}}, IR^{\text{firm}}$$

According to Proposition 1, the expected profit function $\mathbb{E}\Pi(M_r, M_b, p)$ is discontinuous at points where the victim firm's equilibrium decision switches from "accept" to "reject" and thus the insurance payment changes between M_b (insurance payment following rejection) and M_r (insurance payment following acceptance).

Profit-maximizing contract that induces ransom payment.

Lemma 5. *The profit-maximizing contract that induces an equilibrium in which the firm pays ransom has* $(M_r^* = s, M_b^* = b)$ *and premium that is the unique solution to* $U^I(p^*, s, b) = U^N$*. The expected profit of the insurer is* $p^* - q\,s$*.*

First, we look for the most profitable contract that induces payment of ransom and the most profitable contract that induces rejection of the ransom offer.We want to identify the profit-maximizing insurance contract under the constraint that ransom is paid in equilibrium. By Proposition 1, the relevant IC constraint is $s \leq b + M_r - M_b$. The Lagrangian of the insurer's constrained maximization problem is:

$$\mathcal{L}(p, M_r, M_b) = p - qM_r - \lambda\big(s - \big(b + M_r - M_b\big)\big) - \mu(U^N - U^I(p, M_r, M_b)) \tag{4}$$

and I remind the reader that:

$$U^I(p, M_r, M_b) = (1-q)\,u(w-p) + q\,u(w - p - (b - M_b)) \tag{5}$$

and

$$U^N = (1-q)\,u(w) + q\,u(w-b) \tag{6}$$

First-Order Conditions:

$$\begin{aligned}
[M_r]: &\quad -q + \lambda + \mu \underbrace{\frac{\partial U^I}{\partial M_r}}_{=0} = 0 \implies \lambda = q > 0 \\
[M_b]: &\quad -\lambda + \mu \frac{\partial U^I}{\partial M_b} = 0 \implies \mu = \lambda\left(\frac{\partial U^I}{\partial M_b}\right)^{-1} > 0 \\
[p]: &\quad \mu = -\left(\frac{\partial U^I}{\partial p}\right)^{-1}
\end{aligned}$$

where by Proposition 1:

$$U^I(p, M_r, M_b) = (1-q)\,u(w-p) + q\,u(w - p - b + M_b) \tag{7}$$

and

$$U^N = (1-q)\,u(w) + q\,u(w-b) \tag{8}$$

There is no need to check complementary-slackness conditions, because the first-order conditions reveal that both constraints must bind at an optimal solution. In other words, the adversary must be made indifferent between monetizing via ransom or via the outside option of selling data. The f.o.c. imply:

$$\lambda = q \tag{9}$$

$$\mu = \frac{q}{q\,u'(w-p+M_b-b)} \tag{10}$$

$$\mu = \frac{1}{(1-q)\,u'(w-p) + q\,u'(w-p+M_b-b)} \tag{11}$$

Combining the last two f.o.c. yields:

$$\begin{aligned}
\frac{1}{(1-q)\,u'(w-p) + q\,u'(w-p+M_b-b)} &= \frac{1}{u'(w-p+M_b-b)} \\
u'(w-p+M_b-b) &= (1-q)\,u'(w-p) + q\,u'(w-p+M_b-b) \\
u'(w-p+M_b-b) &= u'(w-p) \\
M_b &= b
\end{aligned}$$

where the last step follows from $u'' < 0$. So, full-insurance against business interruption is profit-maximizing. This already pins down the optimal premium p^* since by Proposition 1, insured firms' utility does not depend on M_r. The optimal premium solves:

$$\begin{aligned}
U^I(p^*, M_r, b) = U^N &\iff \\
u(w-p^*) = U^N &\iff \\
p^* = w - u^{-1}(U^N) &\iff \\
p^* = w - u^{-1}\Big((1-q)\,u(w) + q\,u(w-b)\Big)&
\end{aligned}$$

Given $M_b = b$, the binding IC constraint ($\lambda > 0$) then reveals that $M_r = s$ must hold, which is intuitive: lower values of M_r would violate the IC and induce rejection. Offering a contract with a higher value of M_r would maintain acceptance of the ransom offer in equilibrium of the bargaining subgame, but, as shown above, would not increase the premium a firm is willing to pay. It would, however, increase the expected payment by the insurer, and would thus be less profitable to offer.

Profit-maximizing contract that induces ransom rejection.

When the insurer wants to induce no payment of ransom, the constraint becomes $s \geq b + M_r - M_b$, and the Lagrangian of the constrained maximization problem is:

$$\mathcal{L}(p, M_r, M_b) = p - qM_b + \lambda\big(s \,-\, \big(b + M_r - M_b\big)\big) - \mu(U^N - U^I(p, M_r, M_b)) \tag{12}$$

First-Order Conditions:

$$\begin{aligned}
&[M_r]: \quad && \lambda \,+\, \mu \underbrace{\frac{\partial U^I}{\partial M_r}}_{=0} = 0 \implies \lambda = 0 \\
&[M_b]: \quad && -q \,-\, \lambda \,+\, \mu \frac{\partial U^I}{\partial M_b} \;=\; 0 \implies \mu = q\,(\frac{\partial U^I}{\partial M_b})^{-1} > 0 \\
&[p]: \quad && \mu = -(\frac{\partial U^I}{\partial p})^{-1}
\end{aligned}$$

First, the top condition reveals that the relevant IC constraint is now *slack*: since M_r affects neither the utility of the insured firm nor the expected insurance payment, it can be freely adjusted to satisfy the IC constraint for any value of s. Second, combining the last two conditions and repeating the algebraic step from the previous proof implies $M_b^* = b$. Third, and given the above, the optimal premium solves $U^I(M_r, b, p^*) = U^N$. Again, M_r does not directly enter the left-hand side, so p^* is again pinned down by $M_b = b$: This is enough to guarantee full insurance, since the agent's net harm is always $(M_b - b)$, by Proposition 1.

$$p^* = w - u^{-1}\Big((1-q)\,u(w) + q\,u(w-b)\Big) \tag{13}$$

When any of the contracts $(p^*,\ M_b = b,\ M_r \leq s)$ is used, the insurer makes expected profit $p^* - qb$.

Comparing profits

The above analysis reveals that p^* is the same in every candidate optimal contract, hence the comparison of profits simply requires a comparison of expected insurance payments. The expected payment for the contract that induces ransom rejection, $q\,M_b^* = q\,b$, is lower, if and only if $b < s$, leading to the stated Proposition.

As a final note, notice that the contract $M_r = s, M_b = b,\ p^*$ can be supported in equilibrium in either case $b > s$, or $b < s$, but the (insured) firm's equilibrium strategy differs between the two cases. When $b < s$, there can be no equilibrium in which this contract is offered and the firm accepts the ransom offer. The insurer would profitably deviate to instead offer $M_r = s - \epsilon$, inducing rejection and thus reducing the insurance payment without changing p^*.

B Extensions

B.1 Proof of Proposition 3

Once the hacker has successfully breached the firm and seized control of network/data, the equilibrium ransom is determined by the Nash Bargaining Solution, with β being the bargaining power of the ransomware gang. The insured firm's payoff from successful bargaining is $(-r+M_r)$. Formally:

$$
\begin{aligned}
r^* &= \arg\max_r \quad \big(r-s\big)^{\beta}\big[(M^r - r) - (M^b - b)\big]^{(1-\beta)} \\
&= \arg\max_r \quad \big(r-s\big)^{\beta}\big(r^{max}(M_r, M_b) - r\big)^{(1-\beta)}
\end{aligned}
$$

where $r^{max}(M_r, M_b) = b + M^r - M^b$ is the firm's willingness to pay for ransom, and the first-order condition yields:

$$
r^* = (1-\beta)s + \beta r^{\max} \tag{14}
$$

Ransom payment occurs in equilibrium if $s \leq r^* \leq r^{max}(M_r, M_b)$, and I look for such an equilibrium. By standard arguments, the profit-maximizing contract must offer **full insurance**, $M_r = r^*$, i.e.:

$$
\begin{aligned}
&M_r = (1-\beta)s + \beta(b + M^r - M^b) \iff \\
&\beta(M_b - M_r) = \beta\,(b-s) + s - M_r
\end{aligned} \tag{15}
$$

At the profit-maximizing insurance contract, M_b must be such that the attacker's IC constraint binds, i.e. $r^* = s$. Lowering the equilibrium ransom paid by insured firms raises the premium firms are willing to pay for insurance. At the same time, the insurer's cost is not increased, because M_b is not paid out in equilibrium with ransom payment. The adversary earns his

disagreement payoff when $r^* = s$ which is equivalent to:

$$\begin{aligned}&(1-\beta)s + \beta(b + M^r - M^b) = s \iff \\ &\beta(b-s) = \beta(M_b - M^r) \end{aligned} \tag{16}$$

Combining equations (15) and (16) shows that the only insurance terms that achieve both full insurance and maximal reduction of ransom demand are: $M_r = s$, $M_b = b$. Thus, these must be part of the profit-maximizing contract. The associated premium is given by

$$p^*(\beta) = w - u^{-1}(U^N(\beta)) \tag{17}$$

, where:

$$U^N(\beta) = q\,u(w - h(0,0)) + (1-q)u(w)$$

B.2 Proof of Proposition 4

I show that there is a *unique* insurance contract that maximizes expected utility of insured firms, subject to maintaining weakly positive profit for the insurer. Hence, in competitive equilibrium, all insurers will be offering that contract and equilibrium will be in symmetric strategies. Assume that the contract that solves this problem induces payment of ransom in equilibrium. To find the contract that maximizes expected utility subject to the incentive and break-even constraints, I write the relevant Lagrangian:

$$\mathcal{L} = U^I(M_r, M_b, p) - \lambda(s - b - M_r + M_b) - \mu(qM_r - p) \tag{18}$$

The associated first-order conditions are:

$$\begin{aligned}
[M_r]: &\quad \frac{\partial U^I}{\partial M_r} + \lambda - \mu q = 0, \\
[M_b]: &\quad \frac{\partial U^I}{\partial M_b} - \lambda = 0 \implies \lambda = \frac{\partial U^I}{\partial M_b} > 0 \\
[p]: &\quad \frac{\partial U^I}{\partial p} + \mu = 0 \implies \mu = -\frac{\partial U^I}{\partial p} > 0
\end{aligned}$$

The second condition reveals that as long as higher value of M_b reduces the ransom paid by the focal victim firm without increasing the insurance firm's expected cost of insurance, the IC constraint of the adversary will bind, and $s = b + M_r - M_b$. The third condition implies the break-even constraint will bind and $p^* = q\,M_r^*$. As before, by Proposition 1, the net harm to the firm is always $b - M_b$, and optimality of full insurance implies $M_b^* = b$. The binding IC constraint of the adversary implies $M_r^* = s$. Thus, the equilibrium utility of the (fully) insured firms is $U^I = u(w - p^*) = u(w - q\, s)$.

Finally, we must show that as long as $b > s$, no competitor can profitably undercut by switching to a contract that induces rejection of the offer by an insured firm. I do this by showing that a deviating insurer cannot offer a contract that provides utility greater than $u(w - q\, s)$ to firms while earning weakly positive profits. Even if the break-even constraint binds for the deviating insurer, so that $p' = q\, M_b'$, the highest expected utility he could offer to firms would be lower than $u(w - q\, s)$. With a binding break-even constraint, that expected utility is:

$$q\, u(w - q\, M_b - b + M_b) + (1 - q)\, u(w - q\, M_b)$$

In the event of a successful attack, firms reject the offer, suffer b and are reimbursed M_b. This expected utility is maximized at $M_b = b$ (full insurance). But even at that maximum level, it is equal to $u(w - q\, b) < u(w - q\, s)$, so there is no profitable deviation for the insurer.

C Unobserved Contracts

C.1 Proof of Proposition 5

Best response of insurer Assume the insurer wants to induce acceptance of the offer, in which case expected insurance payment is $q\, M_r$. In order to induce acceptance, the firm's relevant IC constraint that contract terms must satisfy is $r \leq r^{max}(M_r, M_b)$, and the resulting Lagrangian for the insurer's maximization problem is:

$$L(M_r, M_b, p) = p - q\, M_r - \mu(U^N - U^I(M_r, M_b, p)) - \lambda(r - (M_r + b - M_b)) \qquad (19)$$

where:

$$U^I(M_r, M_b, p; r) = (1 - q)u(w - p) + q\, u(w - p - r + M_r) \qquad (20)$$

$$U^N(r) = (1 - q)u(w) + q\, u(w - r) \qquad (21)$$

and the first-order conditions are:

$$\begin{aligned}
&[M_r] \quad \mu\frac{\partial U^I}{\partial M_r} + \lambda - q = 0 \\
&[M_b] \quad \mu\underbrace{\frac{\partial U^I}{\partial M_b}}_{=0} - \lambda = 0 \implies \lambda = 0 \\
&[p] \quad 1 + \mu\underbrace{\frac{\partial U^I}{\partial p}}_{<0} = 0 \implies \mu = -\Big(\frac{\partial U^I}{\partial p}\Big) > 0
\end{aligned}$$

The condition for M_b implies that $\lambda = 0$, which is intuitive: when M_b is unobserved, the insurer's choice of M_b does not directly affect the ransom offer hence it is costless to change it in order to satisfy the constraint. The last f.o.c. intuitively implies that the second constraint binds and the firm is made indifferent via the premium. The derivative with respect to the premium is:

$$\frac{\partial U^I}{\partial p} = -(1-q)u'(w-p) - q\,u'(w-p-r+M_r) \tag{22}$$

and combining the three f.o.c. yields:

$$\begin{aligned}
qu'(w-p-(r-M_r)) = q\big[(1-q)u'(w-p) + q\,u'(w-p-(r-M_r))\big] \implies \\
M_r = r
\end{aligned}$$

where the last step is implied by the concavity of u. Thus, in any equilibrium in which the insurer wants to induce payment of ransom, it also offers full insurance against ransom payments. Then, the incentive-compatibility constraint of the victim becomes $M_b \leq b$, i.e. insurance against business interruption must be (weakly) incomplete. For the case of $r \leq b$, this will be optimal and instead of inducing ransom rejection which causes greater harm b; in that case, full insurance provision would raise the insurer's cost, and yield lower profit, by the logic of Proposition 2. Thus, for the case of $r < b$, the best-response correspondence of the insurer is:

$$BR^I(r) = \begin{cases} M_r = r, \\ M_b \in [0, b], \\ p : U^I(M_r, M_b, p; r) = U^N(r) \end{cases}$$

In equilibrium the insurer's and adversaries' conjectures must be correct, and putting the two best responses together: $M_r = r(M_r, M_b) = \max\{b + M_r - M_b,\ s\}$. In any equilibrium in which ransom is paid, $\max\{b + M_r - M_b,\ s\} = b + M_r - M_b > s$. The crossing of best responses then implies:

$$M_b^* = b \tag{23}$$

Finally, in such an equilibrium, the adversary's IC constraint that must be satisfied to ensure ransom is paid in equilibrium is $M_r \geq s$.

This proves that if $b > s$, then for every value $b \geq M_r \geq s$, the contract M_r, $M_b = b$, $p^* = w - u^{-1}(U^N(r))$ and ransom $r^* = M_r^*$ constitute an SPNE strategy profile. In each equilibrium, the insurer pays M_r every time a firm is breached. The adversary best-responds to the insurance contract, the insurance contract maximizes profits, given the ransom demand. Firms are indifferent between purchasing insurance or not and at the bargaining stage are indifferent between accepting and rejecting the ransom offer.

Next, I show that these are the *only* equilibria with acceptance of ransom. There are no equilibria with $M_r = r < s$, because the adversary would not be playing a profit-maximizing strategy. I also prove that there exists no equilibrium in which $r = M_r > b = M_b$. In this case, the insurer would have a profitable deviation. Suppose then that there exists an equilibrium in which $M_b = b$, $M_r = b + \Delta > b$ and all firms buy insurance, hence $r = M_r = b + \Delta$. An SPNE profile must specify sequentially rational strategies off-path, too. Off-path, uninsured firms optimally reject the ransom offer $r > b$, hence the expected utility from not buying insurance is $U^N(r) = qu(w-b) + (1-q)u(w) = U^N(b), \forall r \geq b$. In the candidate equilibrium, the premium is $u(w-p) = U^N(b) \iff p = w - u^{-1}(U^N(b))$ and the expected profit of the insurer is $E\pi = w - u^{-1}(U^N(b)) - q\,M_r < w - u^{-1}(U^N(b)) - q\,b$, for any $\Delta > 0$. If the insurer deviates to $M_r' = 0, M_b' = M_b = b$, the firm optimally rejects the (unchanged) ransom offer because $r^* - M_r' = b + \Delta > b - M_b' = 0$. Since the ransom offer is unchanged, the outside option remains of value $U^N(b)$. The optimal deviation contract also provides full insurance against the relevant risk, hence the optimal deviation premium is again given by $p' = p^* = w - u^{-1}(U^N(b))$. The insurer now pays expected insurance payments $q\,M_b$, thus, the deviation expected profit is $E\pi' = w - u^{-1}(U^N(b)) - q\,M_b > E\pi$, and we have obtained a profitable deviation.

Second, I prove that if $s > b$, then the equilibrium must feature rejection of the ransom offer. In any equilibrium $r \geq s$. Towards a contradiction, assume that $s > b$ and there is an equilibrium in which **ransom is paid**; then, it must be that the maximum ransom the firm is willing to pay is larger than s, i.e, $r = M_r + b - M_b \geq s > b$. We know that any profit-maximizing contract must feature full insurance, hence $M_r = r$. In equilibrium, the hacker asks for the highest possible ransom that is accepted $r = M_r + b - M_b$, hence the only value of M_b consistent with equilibrium $M_b = b$. This implies that $M_r = r \geq s > b$ must hold in equilibrium.

The insurance premium charged in such an equilibrium must satisfy $U^I(M_r,\, M_b,\, p\,, r) = U^N(r) \iff u(w-p) = U^N(r) \iff p = w - u^{-1}(U^N(r))$ and U^N is *decreasing* in the ransom r. The insurer's expected profit is $p - q\, M_r \leq p - q\, s$ and it is easy to see now that the insurer has a profitable deviation, given the adversary's strategy r, to any contract that induces rejection of the offer r and offers full insurance against business interruption, $M_b' = b$. To induce rejection, the contract must include any $M_r' < r$ and rejection becomes dominant for the insured firms since $r > M_r' + (b - M_b')$. Firms are fully insured again hence are willing to pay up to $p' = p = w - u^{-1}(U^N(r))$. Expected profit is $p - q\, M_b = p - q\, b \geq p - q\, s \iff s > b$, which is true. Hence, any equilibrium under $s > b$ must feature rejection of the ransom offer.

C.2 Eliminating additional equilibria

Under $b > s$, there also exist equilibria with ransom *rejection.* The adversary asks for ransom $r^* > b$, the contract offered has $M_b^* = b, M_r^* = s - \Delta > 0$, and the indifference inducing premium is charged, $p^* = w - u^{-1}(U^N(b))$. Nobody has a profitable deviation: insured firms find it optimal to reject since $r > b > M_r + b - M_b = s - \Delta$, and adversaries are indifferent over rejected ransom offers and would earn at most $s - \Delta$ by making a ransom offer that is accepted. The insurer earns $p^* - q\, M_b = p^* - q\, b$ on path and would at most earn $p^* - q\, r < p^* - q\, b$ if it were to offer a contract that induces acceptance of ransom. The counterfactual profit is this because (a) the most profitable candidate deviation is to offer the full-insurance contract that induces acceptance, $M_r' = r^* > b$ and (b) the associated deviation premium remains the same: it is pinned down by the off-path utility of uninsured firms, which remains the same since the ransom they face is unaffected by the monopolist's deviation (unobserved contracts). Hence, the suggested strategy profile is an SPNE.

Even though this equilibrium is **not** Pareto dominated by those of Proposition 5, there are natural reasons to disregard it.

Trembles: First of all, if with small probability ε, firms were to tremble and choose not to buy insurance, $r = b$ would yield strictly greater expected payoff to the adversary for every value of $\varepsilon > 0$.

Counteroffers: Second, in an extended version of the game, uninsured firms off-path have clear incentive to counter-offer to pay $r' = b$, which would make both them and the adversaries better off relative to equilibrium strategies.

For these two reasons, I disregard the pure-strategy equilibrium with the "unreasonably" high ransom demand $r > \max\{M_r + b - M_b,\, b\}$ and restrict attention to those of Proposition 5.

C.3 Proof of Proposition 6

In what follows, I evaluate partial derivatives at an equilibrium of the game with unobserved contracts identified in the first branch of Proposition 5. In such equilibria, $M_r = r$, $M_b = b$ and p^* is defined as the solution to $v^I(p) - v^N(M_r) = 0$, which is unique in every equilibrium. I use the ancillary definitions:

$$\begin{aligned} v^I(p) :&= U^I(M_r, M_b^*, p) = (1-q)u(w-p) + q\,u(w-p-r^*+M_r) = u(w-p^*) \\ v^N(M_r) :&= U^N(M_r, M_b^*) = (1-q)u(w) + q\,u(w - r^*(M_r, M_b^*)) \\ r^*(M_r, M_b^*) &= b + M_r - M_b^* = M_r \end{aligned}$$

Then, expected profit of the monopolist is given by:

$$E\pi = p^* - qM_r \implies \frac{\partial E\pi}{\partial M_r} = \frac{\partial p^*}{\partial M_r} - q \tag{24}$$

The implicit function theorem yields:

$$\frac{\partial p^*}{\partial M_r} = -\frac{\partial (v^I - v^N)}{\partial M_r}\Big(\frac{\partial v^I}{\partial p}\Big)^{-1} = \frac{\partial v^N}{\partial M_r}\Big(\frac{\partial v^I}{\partial p}\Big)^{-1} > 0 \tag{25}$$

where the last equality holds because as M_r is increasing in $[s, b]$ and the equilibrium ransom changes, higher M_r increases equilibrium ransom one-to-one and leaves U^I unchanged, holding the premium fixed. Evaluated at $M_b = M_b^* = b$, these partial derivatives are:

$$\frac{\partial v^I}{\partial p} = -u'(w-p^*) \tag{26}$$

$$\frac{\partial v^N}{\partial M_r} = q\,\frac{\partial u(w-r^*)}{\partial M_r} = -q\,u'(w-r^*) \tag{27}$$

hence, the slope of equilibrium profit is *positive* if:

$$\begin{aligned} \frac{q\,u'(w-r^*)}{u'(w-p^*)} - q > 0 &\iff \\ u'(w-r^*) > u'(w-p^*) &\iff \\ w - r^* < w - p^* &\iff p^* < r^* \end{aligned}$$

The equilibrium premium is smaller than r^* in all equilibria in which the firm pays ransom, i.e., for all $M_r \in [s, b]$, otherwise the firms would trivially deviate to not buying insurance. This completes the proof.

C.4 Proof of Proposition 7

First, note that Proposition 1 still holds. If the adversary finds out the true contract, the best response is $r(M_r, M_b) = \max\{s, M_r + (b - M_b)\}$ and the insured firm's net payoff is $(M_b - b)$. The adversary will rely on his conjectures with probability $1 - \kappa$, with which he does not find the contract. Fixing those conjectures, an insured firm's payoff is thus:

$$U^I = (1-q)u(w-p) + q\left[\kappa\, u(w-p+M_b-b) + (1-\kappa)\, u(w-p+\max\{M_r-\tilde{r}, M_b-b\})\right]$$

I am operating under the assumption that the insurer cannot condition payments on whether the adversary discovers the contract it or not. With a given choice of M_r, M_b, the insurer determines whether firms pay ransom in *either* contingency. If the hacker is *informed*, insured firms pay ransom if:

$$r(M_r, M_b) - M_r \leq b - M_b \iff r(M_r, M_b) \geq M_r + (b - M_b) \iff M_r + (b - M_b) \geq s$$

I call this constraint IC^i. If the hacker is uninformed, firms pay ransom if:

$$M_r - \tilde{r} > M_b - b \iff M_r + (b - M_b) \geq \tilde{r}$$

I call this constraint IC^u. Both are satisfied when $M_r + (b - M_b) \geq \max\{s, \tilde{r}\} = \tilde{r}$, i.e. if for some contract, firms pay ransom to uninformed hackers, firms also pay to informed hackers, but not vice-versa, i.e. IC^u implies IC^i. In equilibrium, the uninformed hackers also correctly anticipate the insurance contract, so $\tilde{r} = \max\{s, b + M_r - M_b\}$ and the two constraints coincide. This means that in equilibrium, either firms either pay ransom in both events, or in neither event.

Best response of insurer Suppose the insurer knows the ransom demand by *uninformed* adversaries is $\tilde{r} \in [s, b]$. The ransom demand by informed adversaries is $\max\{b + M_r - M_b, s\}$. By the argument of the last paragraph, the insurer can either use a contract that satisfies both IC^i and IC^u, only IC^i or neither of those.

Case A: Accept offers by both informed and uninformed adversaries.
If the insurer responds with such a contract, the insured firm's payoff is:

$$U^I = (1-q)u(w-p) + q\left[\kappa\, u(w-p+M_b-b) + (1-\kappa)\, u(w-p+M_r-\tilde{r})\right]$$

and the insurer's expected payout is $q\,M_r$. The insurer's Lagrangian is:

$$L = p - qM_r - \lambda\big(\tilde{r} - M_r - b + M_b\big) - \mu(U^N - U^I)$$

The first-order conditions are:

$$\begin{aligned}
&[M_r] \quad -q + \lambda + \mu\, q\,(1-k)u'(w - p - \tilde{r} + M_r) = 0\\
&[M_b] \quad -\lambda + \mu\, q\, k\, u'(w - p - b + M_b) = 0 \implies \lambda > 0\\
&[p] \quad \mu^{-1} = (1-q)u'(w-p) + q\big[ku'(w - p + M_b - b) + (1-k)u'(w - p + M_r - \tilde{r})\big] > 0
\end{aligned}$$

For any $k > 0$, i.e., any positive probability with which the hackers discover the contract, the IC constraint for firms facing uninformed adversaries must bind: If that IC is slack, marginally increasing M_b will increase the ex-post utility of insured firms in the event the contract becomes known (and thus their wtp for insurance), and not increase the insurer's payout, since firms are still accepting the ransom. Thus, for any expected $\tilde{r}$, the insurer optimally responds by setting $M_r = \tilde{r} + M_b - b \iff M_r - \tilde{r} = M_b - b$. Combining conditions $[M_r]$ and $[M_b]$ to eliminate λ yields:

$$\begin{aligned}
&\mu^{-1} = (1-k)u'(w - p - \tilde{r} + M_r) + k\, u'(w - p - b + M_b) \iff\\
&u'(w - p - b + M_b) = u'(w - p - \tilde{r} + M_r)
\end{aligned}$$

Using $[p]$, we obtain the familiar $M_r = \tilde{r}$. The insurance premium extracted is:

$$p^*(\tilde{r}) = w - u^{-1}(U^N(\tilde{r})) \tag{28}$$

and the insurer's profit is:

$$\mathrm{E}\pi^{both} = p^*(\tilde{r}) - q\, M_r = p^*(\tilde{r}) - q\,\tilde{r} \tag{29}$$

Case B: Accept offers by informed but not by uninformed adversaries.
The alternative for the insurer is to set a policy (M_r, M_b) such that $s \leq M_r + b - M_b \leq \tilde{r}$, so that the offer of uninformed adversaries is rejected. If the insurer responds with such a contract, the insured firm's payoff is:

$$\begin{aligned}
U^I &= (1-q)u(w-p) + q\,[\kappa\, u(w - p + M_b - b) + (1-\kappa)\, u(w - p + M_b - b)]\\
&= (1-q)u(w-p) + qu(w - p + M_b - b)
\end{aligned}$$

which is independent of k and M_r because the firm's net loss in both contingencies is $(b - M_b)$ and the insurer's expected payout is $q\,k\,M_r + q\,(1-k)\,M_b$. The insurer's Lagrangian is:

$$L = p - q\,[k\,M_r + (1-k)\,M_b] - \lambda^i\big(s - M_r - b + M_b\big) + \lambda^u\big(\tilde{r} - M_r - b + M_b\big) - \mu(U^N - U^I)$$

The first-order conditions are:

$$[M_r] \quad -q\,k + \lambda^i - \lambda^u = 0 \implies \lambda^i = q\,k + \lambda^u > 0$$

So that IC^i binds and $M_r = s + M_b - b$.

$$[M_b] \quad -q\,(1-k) - \lambda^i + \lambda^u + \mu\,q\,u'(w - p + M_b - b) = 0$$

and combining with the previous condition yields:

$$1 = \mu\,u'(w - p + M_b - b) \tag{30}$$

$$[p] \quad \mu = \frac{1}{(1-q)u'(w-p) + qu'(w-p+M_b-b)}$$

and combining with (30) yields:

$$M_b = b \implies M_r = s$$

where the value of M_r is implied by the binding IC^i. The insurance premium extracted is the same as in Case A:

$$p^*(\tilde{r}) = w - u^{-1}(U^N(\tilde{r})) \tag{31}$$

and the insurer's profit is:

$$\mathrm{E}\pi^{informed} = p^*(\tilde{r}) - q\,[k\,M_r + (1-k)\,M_b] = p^*(\tilde{r}) - q\,[k\,s + (1-k)\,b] \tag{32}$$

Compare to the previously derived profit under acceptance of both:

$$\mathrm{E}\pi^{both} = p^*(\tilde{r}) - q\,M_r = p^*(\tilde{r}) - q\,\tilde{r}$$

Putting everything together, the best response of the insurer is always to set $M_b = b$ and $p^* = w - u^{-1}(U^N(\tilde{r}))$. There are two candidate best-response values for M_r: $M_r = \tilde{r}$ if and only if $[k\,s + (1-k)\,b] > \tilde{r}$, to induce acceptance of the uninformed offer; otherwise set $M_r = s$, inducing rejection of the uninformed offer.

Equilibrium

An uninformed demand r^u can thus be part of an equilibrium if and only if $r^u \leq r(k)$, where $r(k) := k\,s + (1-k)\,b$. If $r^u \leq r(k)$, there is an equilibrium in which $M_r = r^u, M_b = b$ and insured firms accept either informed or uninformed offers. If $r^u > r(k)$, then the best response of the insurer is to induce *rejection* of the uninformed ransom demand, and use $M_r = s, M_b = b$. But as argued already, in equilibrium firms either accept both offers or reject both offers. Hence, there is **no equilibrium** with $r^u > r(k)$. Note that $\lim_{k \to 1} r(k) = s$ which corresponds to the case of **observed** contracts and $\lim_{k \to 0} r(k) = b$, which corresponds to the case of **unobserved** contracts. Greater values of k monotonically reduce the extent of equilibrium multiplicity.

D Liquidity Constraints

D.1 Proof of Proposition 8

I focus on equilibria in which the adversary demands ransom $r \in [s, b]$. In equilibria in which the insurance contract induces acceptance of the ransom, the best response of the insurer will always provide full insurance, so that $M_r = r$. Crucially, given this best response, insured firms' liquidity constraint *does not bind*. By assumption, they can use their insurance M_r to cover for the ransom payment.

Just as in the proof of Proposition 5, to see that the candidate strategy profile is indeed an SPNE, notice that (1) adversaries extract ransom equal to firms' willingness to pay, i.e. $r = b + M_r - M_b$, (2) this ransom exceeds adversaries' outside option, $r = M_r \geq s$, and (3) firms are indifferent between accepting or rejecting the ransom demand. Same as in the case without liquidity constraints, the only value of M_b consistent with equilibrium is $M_b = b$. The

The premium is chosen to make firms indifferent between buying insurance and staying uninsured. The expected utility of uninsured firms depends on the value of ℓ, and on the equilibrium ransom r. Define:

$$h(\ell, r) = \begin{cases} r, & \text{if } r \leq \ell \\ \\ b, & \text{if } r > \ell \end{cases}$$

The off-path expected utility of uninsured firms is $U^N(r; \ell)$ and the equilibrium premium satisfies:

$$u(w - p^*) = U^N(r; \ell) = (1-q)\,u(w) + q\,u(w - h(\ell, r)) \tag{33}$$

Given that the best-response of the insurer is always to provide full insurance, the arguments in Proposition 5 apply verbatim and these are the only equilibria that feature $r \in [s, b]$. There again exist equilibria in which ransom is rejected and $r > b$, but I consider those "unreasonable" and discard them by appealing to the trembling and counteroffer-robustness arguments I discuss in this Appendix.