Who Pays to Secure Shared Components? A Pay-for-Performance Model for Cyber Insurance

John Chester and Adam Hastings

Fordham University, New York, NY, USA
{adam.hastings,john.chester}@fordham.edu

Abstract. In health insurance, pay-for-performance (P4P) programs use targeted financial incentives to improve care quality and reduce downstream costs. We ask whether a similar logic can be applied to cyber insurance. Specifically, we study whether cyber insurers should directly invest in improving the security of shared technology components that are widely deployed across their policyholder base. We construct a model of component-level security and demonstrate that component-level security is a public goods problem that can be addressed by a custom pay-for-performance mechanism. A numerical analysis reveals that direct insurer investment into component security can be rational under plausible conditions. These results suggest that cyber insurers should be viewed not only as passive risk-pricers but as rational investors in shared component security as well. More broadly, we examine how P4P programs may serve as a blueprint for mechanisms for security investment.

Keywords: cyber insurance · software supply chain security · hardware security · component security · security economics · pay-for-performance

1 Introduction

Insecure digital technology continues to be a persistent harm to individuals and society, yet most systems are attacked using known weaknesses with known defenses. But why? It has long been established that the failures of cybersecurity are as much an issue of bad economics as they are of bad technology, and that fixing security requires fixing the structure of incentives surrounding security [1]. In this paper, we identify a system of incentives from an adjacent domain with similar incentive problems—pay-for-performance programs as used in health insurance—and examine how and where such programs may be used to improve security outcomes.

Pay-for-performance is a healthcare payment model where insurers offer financial incentives to care providers (like hospitals or physician networks) for meeting certain performance measures for quality and efficiency [8]. For example, an insurer and a hospital may negotiate a contract under which the hospital receives a bonus for meeting specified targets concerning hospital-acquired infections, readmissions, or adherence to recommended clinical practices. The central insight behind these programs is that the insurers who ultimately bear the costs of poor

health outcomes or cybersecurity breaches are well placed to push upstream agents (like hospitals) to adhere to standards of care they would otherwise underinvest in. For example, in the presence of patient information asymmetry, a negligent hospital may not appropriately prioritize patient healthcare outcomes or readmission rates since the hospital is paid irrespective of outcome. Likewise, consider how this same dynamic appears in cybersecurity: in the presence of technological information asymmetry, a negligent component vendor may not appropriately prioritize product security, since the adverse costs of insecurity are typically passed on to the user of the technology rather than the vendor. In both cases, the right upstream incentive or investment may reduce downstream costs and improve outcomes for patients, policyholders, and insurers alike. We provide a detailed overview of P4P programs in Section 2.

How then might the lessons learned from pay-for-performance programs as used in health insurance apply to cybersecurity and cyber insurance? To this end, we construct a model of a P4P system as it might be applied to cyber insurance and component-level security in Section 3. Importantly, the proposed cyber P4P program elicits private security information, giving component vendors an incentive to disclose their products' security—and responsiveness to investment—to an insurer (something that would otherwise remain a closely guarded secret). Using the solicited information, an insurer can then make rational investments into the third-party technology components that pose the most risk to policyholders. In Section 4, we derive closed-form solutions to this problem under a variety of plausible security investment functions and identify component-level security to be an instance of the public goods game. We demonstrate how the P4P mechanism solves the game by allowing the insurer to act as a social planner, and then evaluate the model numerically and find that insurer investment in third-party component security is rational under plausible model inputs. We discuss the limitations of our model and of P4P systems in general in Section 5 and outline future work directions in Section 6. Finally, Section 7 concludes.

2 Background

Pay-for-performance (P4P), also referred to as value-based care, is a healthcare payment model in which insurers provide targeted financial incentives to hospitals and doctors ("care providers") for measured quality and efficiency outcomes rather than for the volume of care delivered [7–9]. For example, an insurer may provide a bonus payment to a care provider if the care provider achieves an agreed-upon reduction in the rate of complications or readmissions for a particular type of surgery. This benefits all parties involved: If the target is met, care providers receive the bonus payment, patients receive better-quality care, and insurers reduce claims costs. An outline of a typical P4P system is provided in Figure 1.

Pay-for-performance is a well-established approach operating at national scale: the U.S. Centers for Medicare & Medicaid Services (CMS) has run P4P programs for nearly two decades—among them the Hospital Readmissions Reduction Program (HRRP) [7], Hospital Value-Based Purchasing (VBP) [8], and the Merit-

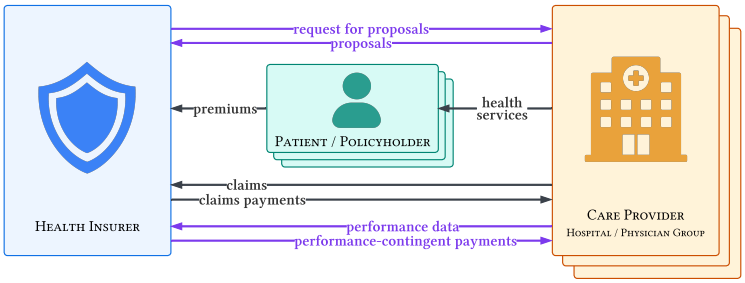


Fig. 1: A systems diagram of a typical pay-for-performance (P4P) program in healthcare. In a P4P system, an insurer makes payments to care providers based on the quality, efficiency, or outcomes of the services provided to patients, in an effort to incentivize higher-quality care and reduce the downstream liability of patients' health needs. P4P mechanics are shown in purple.

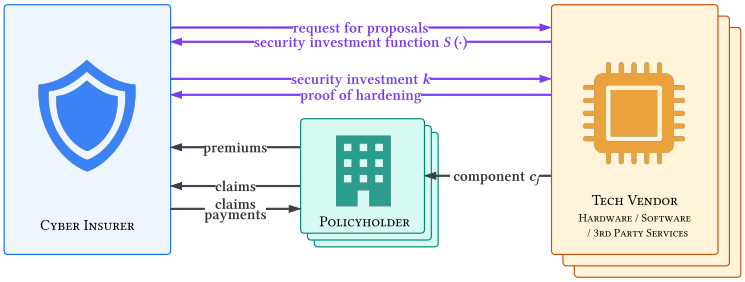


Fig. 2: An example of how P4P programs might be ported to cybersecurity and cyber insurance. An insurer might first determine which technology components are most widely shared and security-critical among policyholders. Then a cyber P4P program (shown in purple) may solicit proposals for security hardening from component vendors, allowing insurers to identify and address leverage points in the policyholder pool via targeted investments.

based Incentive Payment System (MIPS) [9]; private insurers run similar programs as well. These programs offer decades' worth of data and experience on aligning third-party incentives and provide a template for how similar incentive-aligning mechanisms may be used in security.

Two ideas in particular are worth noting: First, when the ultimate outcome is noisy, delayed, or hard to attribute to a single decision, healthcare P4P anchors payment to *measurable proxies*—process and structural measures, such as whether a guideline-recommended treatment was administered—in place of the outcome itself. Second, care providers actively propose concrete interventions to the payer, who funds those promising the largest reduction in downstream liability.

This allows for a fairly straightforward mapping from existing P4P systems to how they might be applied to security, shown in Figure 2. The insurer solicits proposals not from care providers but from technology vendors, who supply widely used technology components (like software, hardware, cloud services, etc.) to the

insurer's policyholders. "Quality of care" becomes "security of a technology component", and a patient's future health needs become the policyholder's expected breach loss. In both settings a single payer (the insurer) aggregates exposure across a large population and can profit from funding upstream interventions—better care, or more secure components—that lower its downstream liability. Both payers also face a targeting problem: just as a health system directs spending to the interventions and patient cohorts where it buys the most health, a cyber insurer can direct hardening dollars to the components most widely shared and most security-critical across its policyholder pool.

One feature, however, does not easily translate: in healthcare, P4P rewards are paid after outcomes are measured against target metrics (like lower readmission and complication rates) using longitudinal patient data, often provided by CMS. In contrast, cybersecurity typically lacks such clear signals: a security enhancement that is highly effective, marginally effective, or entirely superfluous produces the same observable result, namely, no breach occurring. This "no news is good news" property means the absence of an incident cannot be attributed to any particular investment, while a breach may reflect a zero-day or a suppressed disclosure rather than a defensive failure—so there is no ground truth against which performance can be scored after the fact. An outcome-conditioned reward would confound genuine improvement with luck and under-reporting, inviting moral hazard, so we anchor incentives to measurable inputs instead—adoption of security standards, remediation of known vulnerabilities, or hardening of critical shared technology components, including software, hardware, and cloud computing services.

2.1 Related Work

Cyber Insurance: Contract Design and Incentive Modeling. The cyber insurance literature has predominantly adopted a contract-centric paradigm focusing on optimizing the terms of the contract to incentivize security investment [17, 18]. Prior work has modeled the insurer as a passive risk-pricer who optimizes terms (deductible, premiums, and coverage limits) for each insured firm. Moral hazard, adverse selection, and correlated loss from network effects and shared service providers are addressed with the central focus being the firm-level relationship between the insurer and the insured [3, 5, 16, 23].

Security Economics, Interdependent Security, and Public Goods Provision. Individual firm investment in security has long been established as socially insufficient due to market failure and free-riding [1,2,12,13]. Typically this presents as the classic problem of the private provisioning of public goods [4,22,24]. More recent work examines incentive mechanisms such as vendor liability and audits [14].

Software Supply Chain and Dependency-Level Risk. Software dependency attacks like Log4Shell and the XZ Utils backdoor have brought new urgency to software supply-chain security. Software supply chains have been studied empirically [21, 26]. Meanwhile, the policy and standards communities have largely taken a compliance-oriented approach [6, 10, 20] but do not specify when security hardening should be funded, nor by whom.

Bug Bounties and Market Approaches to Security Funding. Bug bounties are a form of an existing pay-for-performance mechanism in security where organizations pay external researchers for valid vulnerability reports affecting their own products or services [11, 19, 25]. Our setting differs in both the funder and the target of investment: bug bounties are typically funded by the vendor or platform whose product is tested, whereas our mechanism is funded by an insurer that aggregates downstream losses across many policyholders. Furthermore, bounties primarily reward vulnerability discovery whereas our work asks when a third party should fund the hardening of a shared component.

Our contribution. Each of these literatures leaves part of the problem open. Contract design stops at the insured firm, security economics explains underprovision without naming a funder, supply-chain work responds with visibility and compliance rather than a funding model, bug bounties draw on the vendor's own bounded stake, and healthcare P4P depends on an outcome signal that cybersecurity lacks. This paper models insurer-funded hardening at the component level and derives closed-form conditions under which aggregate portfolio exposure makes it rational, including for dependencies that belong to no single insurable firm.

3 Model Description

We make the following assumptions in our model:

1. Policyholders are denoted by the finite set Π, with typical element π_i.
2. Technology components are denoted by the finite set $\mathcal{C}$, with typical element c_j. These comprise both the third-party products supplied by technology vendors and the bespoke in-house systems operated by individual policyholders.
3. Each policyholder π_i operates a *tech stack* $\mathcal{C}_i \subseteq \mathcal{C}$, the set of components it deploys. Among these is its in-house technology $c_i^{\text{int}} \in \mathcal{C}_i$, which no other policyholder uses; the rest are third-party products that other policyholders may deploy as well. Both kinds contribute to π_i's breach risk, but only components shared by two or more policyholders give rise to the collective-action problem we study in Section 4.2.
4. Each component c_j has an associated security level s_j with $0 \leq s_j \leq 1$, where $s_j = 1$ represents perfect security while $s_j = 0$ represents perfect insecurity. Its complement, the component *vulnerability* $v_j = 1 - s_j$, is the component's probability of compromise, $\Pr(c_j \text{ compromised}) = v_j$. We treat vulnerability as an intrinsic property of the component that does not vary between policyholders.
5. A *breach* is a loss event where a policyholder π_i's security has failed, causing a loss of amount ℓ_i.
6. For all π_i, each component $c_j \in \mathcal{C}_i$ has a conditional breach probability $\beta_{ij} = \Pr(\text{breach of } \pi_i \mid c_j \text{ compromised})$, representing the probability that a compromise of component c_j leads to a full breach of policyholder π_i. The breach probability β_{ij} is specific to each policyholder π_i for each component c_j,

reflecting how deeply the component is integrated within their technology stack.

7. The *component risk* ρ_{ij} that component c_j poses to policyholder π_i is the product of its intrinsic vulnerability and conditional breach probability, and is defined as $\rho_{ij} \;=\; \beta_{ij}\, v_j$. This formulation decouples the inherent security of a component from the structural role it plays within a particular policyholder's technology stack[1].
8. We assume a single cyber insurer for simplicity, though this may be extended to multiple insurers in future work.
9. There exists a *security investment function* $S_j(\cdot) : [0, \infty) \to [s_j, 1]$ that maps a capital investment $k \geq 0$ in component c_j to an enhanced security level, and with $S_j(0) = s_j$. Specific functional forms of $S_j(\cdot)$ are studied in §3.1.
10. The probability that policyholder π_i suffers a breach is:

$$\Pr(\text{breach of } \pi_i) = 1 - \prod_{c_l \,\in\, \mathcal{C}_i} (1 - \rho_{il}) \tag{1}$$

This follows from Assumptions 4 and 7, treating breach events across components as independent.

11. A breach of π_i costs a fixed amount ℓ_i, so its expected loss is $\Pr(\text{breach of } \pi_i)\, \ell_i$. Writing L for the insurer's network-wide loss, a random variable, and taking the insurer to be risk-neutral, its expected value sums across all policyholders:

$$\mathbb{E}\,[L] = \sum_{i \in \Pi} \Pr(\text{breach of } \pi_i)\; \ell_i \tag{2}$$

12. We now isolate the part of the insurer's expected losses that component c_j governs. Fix c_j and hold the security of every other component constant. Separating the factor for c_j from the rest of the stack in Eq. (1) gives, for each policyholder,

$$\Pr(\text{breach of } \pi_i) = 1 - \big(1 - \rho_{ij}\big)\sigma_{ij} \;=\; \big(1 - \sigma_{ij}\big) \;+\; \sigma_{ij}\, \rho_{ij}, \tag{3}$$

where $\sigma_{ij} \;=\; \prod_{c_l \,\in\, \mathcal{C}_i \setminus \{c_j\}} \big(1 - \rho_{il}\big)$ is the probability that π_i escapes a breach through every component *other than* c_j. Because σ_{ij} is determined entirely by those other components, it is unaffected by any investment in c_j. Writing $\rho_{ij} = \beta_{ij}\, v_j$ (Assumption 7) and summing over the policyholders, the insurer's network-wide expected loss then separates exactly into a part that c_j cannot affect and a part that it controls. Since v_j is a property of c_j alone (and does not vary across policyholders), it factors out of the second sum:

$$\mathbb{E}\,[L] = \underbrace{\sum_{i \in \Pi} \big(1 - \sigma_{ij}\big)\ell_i}_{R_j} \;+\; v_j \cdot \underbrace{\sum_{i \in \Pi} \sigma_{ij}\, \beta_{ij}\, \ell_i}_{A_j} \tag{4}$$

[1] For example, a component with high vulnerability but low conditional breach probability (e.g., a poorly secured but isolated service) may pose less overall risk than a moderately vulnerable component that is deeply integrated into critical infrastructure.

For a policyholder that does not use c_j we have $\beta_{ij} = 0$, so it contributes nothing to A_j and its full expected loss is carried by R_j. We call the scalar A_j the *aggregate component exposure* of c_j. Each policyholder contributes the exposure it carries through c_j, namely $\beta_{ij}\,\ell_i$, weighted by σ_{ij} so that a loss is counted only in those states where π_i would not already have been breached by some other route. A_j is therefore the expected loss genuinely *attributable* to c_j, and R_j is the loss the insurer bears irrespective of c_j's security. Since a component with security $S_j(k)$ has vulnerability $1 - S_j(k)$, the attributable term after investment k becomes $\big(1 - S_j(k)\big)A_j$.

13. The insurer's total cost of investing $k \geq 0$ in component c_j, accounting for both the reduction in expected losses and the direct cost of investment, is the *component objective*:

$$J_j(k) \;:=\; \big(1 - S_j(k)\big)A_j + k, \qquad k \geq 0 \tag{5}$$

This is the expected loss of Eq. (4) plus the cost of investment, with the residual R_j omitted[2]. The insurer then seeks to minimize $J_j(k)$ over all non-negative investment amounts. Because A_j is a fixed scalar, this reduces to a single-variable optimization for any chosen $S_j(\cdot)$. Closed-form minimizers k_j^* are derived in Section 4.

3.1 P4P Investments

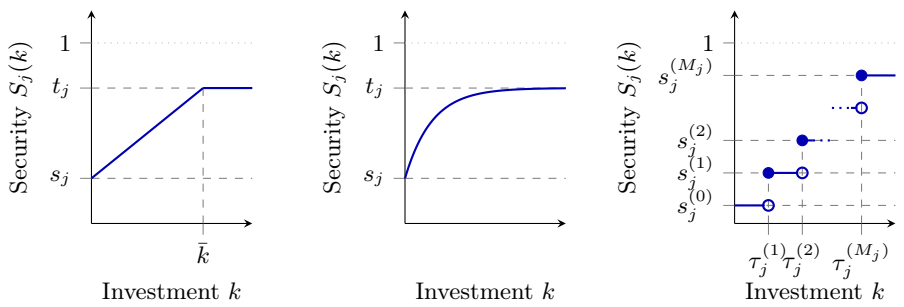
Rather than assert the shape of the security investment function $S_j(\cdot)$ from Assumption 9, we instead choose to study a variety of plausible security investment functions $S_j(\cdot)$. In practice, the insurer need not estimate this shape in a vacuum: as in healthcare P4P, where care providers pitch insurers with concrete investment proposals, component vendors competing for funding can disclose the security improvement a given investment would yield, making $S_j(\cdot)$ elicitable rather than purely hypothetical. Crucially, this channel does not arise in the ordinary market: a vendor has little incentive to disclose its products' weaknesses to prospective customers, for whom such disclosure depresses sales [1], yet every incentive to disclose them to an insurer who will fund their remediation—so the insurer can learn $S_j(\cdot)$ precisely because it pays to improve security rather than to purchase it. We return to the strategic limits of such disclosure in Section 5.

Bounded Linear Investment. The simplest possible security investment function $S_j(\cdot)$ is a linear relationship between investment k and security $S_j(k)$. Since our model requires $S_j(k) \leq 1$ for all k, and since hardening may plateau short of perfect security, we take

$$S_j(k) = \min\{t_j,\; a_j k + s_j\} \tag{6}$$

for some rate $a_j > 0$ and security ceiling $t_j \in (s_j, 1]$, shown in Figure 3a.

[2] Since the insurer bears R_j no matter what it spends on c_j, it is an additive constant that shifts J_j without moving its minimizer. Dropping it leaves J_j measuring only what investment in c_j can influence.



(a) Bounded linear (b) Diminishing returns (c) Stepwise

Fig. 3: Three plausible security investment functions $S_j(\cdot)$.

Diminishing Returns. Perhaps a more realistic security investment function $S_j(\cdot)$ is a "diminishing returns" form, where each additional dollar invested yields decreasing security improvements (Figure 3b). We represent this as a saturating exponential starting at s_j and asymptotically approaching the ceiling $t_j \in (s_j, 1]$:

$$S_j(k) = t_j - (t_j - s_j)e^{-\lambda_j k} \tag{7}$$

for some convergence factor $\lambda_j > 0$.

Stepwise Investment. Security investments might also be discontinuous, where investments produce no effect unless they reach some threshold amount. In general, $S_j(\cdot)$ may even be a *series* of such stepwise improvements with component-specific threshold points $0 < \tau_j^{(1)} < \tau_j^{(2)} < \cdots < \tau_j^{(M_j)}$, represented as:

$$S_j(k) = \begin{cases} s_j^{(0)} & \text{if} \quad k < \tau_j^{(1)}, \\ s_j^{(1)} & \text{if} \quad \tau_j^{(1)} \le k < \tau_j^{(2)}, \\ s_j^{(2)} & \text{if} \quad \tau_j^{(2)} \le k < \tau_j^{(3)}, \\ \quad\vdots & \\ s_j^{(M_j)} & \text{if} \quad k \ge \tau_j^{(M_j)} \end{cases} \tag{8}$$

where $s_j^{(0)} = s_j$, $0 \le s_j^{(0)} < s_j^{(1)} < \cdots < s_j^{(M_j)} \le 1$, and $M_j \ge 1$. We write $\Delta s^{(m)} := s_j^{(m)} - s_j^{(0)}$ for the security gained by reaching the mth threshold from the baseline. This is shown in Figure 3c.

4 Model Analysis

4.1 Optimal Single-Component Investment

For a fixed component c_j, the insurer minimizes the component objective $J_j(k)$ from Assumption 13. Closed-form minimizers k_j^* are derived below. Throughout, an agent indifferent between investing and not is taken not to invest.

Bounded Linear Enhancement. For the bounded linear form, define the saturation level $\bar{k} = (t_j - s_j)/a_j$ when $a_j > 0$. For $0 \le k \le \bar{k}$, we have $J_j(k) = (1 - s_j - a_j k)A_j + k$, $\frac{dJ_j}{dk} = 1 - a_j A_j$, and for $k > \bar{k}$, we have $J_j(k) = (1 - t_j)A_j + k$, $\frac{dJ_j}{dk} = 1$. Hence,

$$k_j^* = \begin{cases} 0, & a_j A_j \le 1, \\ \dfrac{t_j - s_j}{a_j}, & a_j A_j > 1. \end{cases} \tag{9}$$

Diminishing Returns. For the bounded diminishing-returns form with asymptote $t_j \in (s_j, 1]$, we have

$$S_j(k) = t_j - (t_j - s_j)e^{-\lambda_j k}, \qquad \lambda_j > 0.$$

Then

$$J_j(k) = \big[(1 - t_j) + (t_j - s_j)e^{-\lambda_j k}\big]A_j + k,$$

so $\frac{dJ_j}{dk} = 1 - \lambda_j(t_j - s_j)e^{-\lambda_j k}A_j$ and $\frac{d^2 J_j}{dk^2} = \lambda_j^2(t_j - s_j)e^{-\lambda_j k}A_j \ge 0$. Solving $\frac{dJ_j}{dk} = 0$ and imposing $k \ge 0$ gives

$$k_j^* = \begin{cases} 0, & \lambda_j(t_j - s_j)A_j \le 1, \\ \dfrac{1}{\lambda_j}\ln\big(\lambda_j(t_j - s_j)A_j\big), & \lambda_j(t_j - s_j)A_j > 1. \end{cases} \tag{10}$$

Stepwise Investment. For the stepwise form, $S_j(\cdot)$ is piecewise constant, so the only candidate minimizers are $k = 0$ and the threshold breakpoints $\tau_j^{(1)}, \ldots, \tau_j^{(M_j)}$. Evaluating the objective at each,

$$J_j(0) = (1 - s_j^{(0)})A_j, \qquad J_j(\tau_j^{(m)}) = \tau_j^{(m)} + (1 - s_j^{(m)})A_j,$$

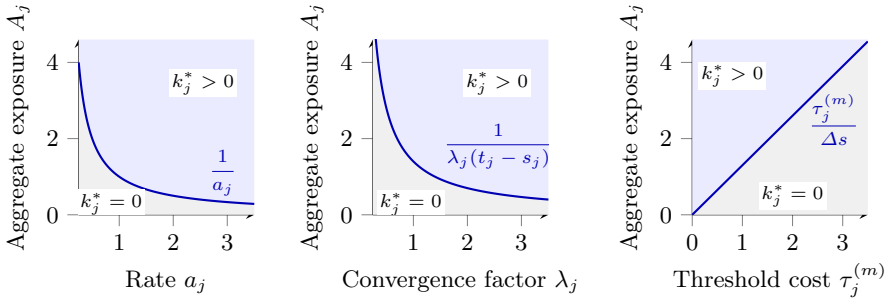
the insurer chooses the smallest candidate in $\{0, \tau_j^{(1)}, \ldots, \tau_j^{(M_j)}\}$ that attains the minimum across these $M_j + 1$ values.

Figure 4 summarizes the insurer's investment boundary for each form of $S_j(\cdot)$.

4.2 Component Security as a Public Goods Game

Using our model, we now demonstrate that the provision of component security is an instance of the public goods game. Investing in a component c_j raises its security level and lowers its vulnerability level for every user of c_j simultaneously[3], so component security is non-rival and non-excludable across the policyholders that share it, making it a public good.

[3] We do not consider the adjustment costs of patching in this work.



(a) Bounded linear (b) Diminishing returns (c) Stepwise

Fig. 4: Decision boundaries for insurer investment in a single component. When $k_j^* > 0$, it is rational for the insurer to pay for the security hardening of third-party components.

Consider the policyholders that depend on component c_j, i.e., those with $\beta_{ij} > 0$, and write each one's component-specific exposure as

$$A_{ij} := \sigma_{ij}\beta_{ij}\ell_i, \qquad A_j = \sum_{i \in \Pi} A_{ij}.$$

Each σ_{ij} is determined by the rest of π_i's stack, which π_i alone controls and already internalizes, so σ_{ij} plays no strategic role in what follows.

Decentralized provision. Absent an insurer, suppose policyholders could fund hardening themselves: each π_i contributes $k_i \geq 0$, total spending $K = \sum_l k_l$ sets the component's security, and every dependent enjoys the resulting security improvement. In this contribution game, each policyholder internalizes only its own exposure A_{ij}, and the strongest decentralized incentive comes from the highest-exposure policyholder, $\max_i A_{ij}$. In the bounded linear case, hardening saturates at $\bar{k}$ or remains zero, while in the diminishing-returns case, hardening is set by the most exposed policyholder's ideal level while the rest free-ride. Stepwise hardening is different: cost-sharing can fund a collectively worthwhile threshold, but no-provision remains an equilibrium when no policyholder would fund the threshold alone. For a full analysis, see Appendix A.1.

Insurer provision. Unlike any individual policyholder, the insurer bears losses across the whole insured portfolio. It therefore values hardening c_j by the aggregate exposure $A_j = \sum_i A_{ij}$, and so its optimal investment is the k_j^* derived in Section 4.1. The gap between $\max_i A_{ij}$, the largest decentralized stake, and A_j, the pooled stake, is the *wedge*. For continuous investments this wedge creates an insurer-only provision region; for stepwise investments it marks profitable thresholds that no policyholder funds alone.

Proposition 1 (Pooled provision of a shared component). *For the bounded linear and diminishing-returns forms, there is a non-empty set of parameter values—the* insurer-only region*—in which decentralized provision is zero in every equilibrium, yet insurer-funded provision is cost-rational.*

1. (Bounded linear.) *The region is where* $\max_i A_{ij} < 1/a_j < A_j$.
2. (Diminishing returns.) *The region is where* $\max_i A_{ij} < 1/[\lambda_j(t_j - s_j)] < A_j$.

The wedge, and hence the insurer-only region, is non-empty precisely when $A_j > \max_i A_{ij}$*, i.e., when at least two policyholders carry positive exposure.*

The proof is given in Appendix A.2.

Stepwise investments yield a weaker threshold-coordination result. Writing $\Delta s^{(m)} = s_j^{(m)} - s_j^{(0)}$, insurer funding is cost-rational while no policyholder funds any threshold alone when every ratio $\tau_j^{(m)}/\Delta s^{(m)}$ exceeds $\max_i A_{ij}$ and at least one falls below A_j. No-provision then remains a decentralized equilibrium, though cost-sharing equilibria may also exist.

In each condition the middle quantity is the exposure at which hardening c_j just pays for itself: $1/a_j$ for bounded linear investment, $1/[\lambda_j(t_j - s_j)]$ for diminishing returns, and $\tau_j^{(m)}/\Delta s^{(m)}$ for a stepwise jump. The continuous forms therefore exhibit classical underprovision of a public good, while the stepwise form adds threshold coordination [4, 22, 24]. Because the cyber insurer bears the pooled loss, its private incentive to harden a shared component coincides with the social optimum for its portfolio. For a normalized view under the assumption of symmetric policyholders, see Section 4.4.

4.3 Comparative Statics

We now ask how the model's behavior changes with the component's baseline security, aggregate exposure, and the security investment function, and use those comparative statics to derive a targeting rule.

Baseline security and headroom. Does a component's baseline security affect whether hardening is worth funding? For bounded linear investment, the answer is no: the intervention boundary $a_j A_j = 1$ is independent of s_j, because the pre-saturation slope $dJ_j/dk = 1 - a_j A_j$ does not involve the baseline. Baseline security affects how much the insurer spends conditional on intervening, but not whether the first dollar pays.

For diminishing returns, baseline security matters directly. The relevant quantity is

$$\Theta_j := \lambda_j(t_j - s_j)A_j,$$

the return on the first dollar of hardening, so intervention is optimal exactly when $\Theta_j > 1$. A higher baseline reduces the remaining headroom $t_j - s_j$, making

investment less attractive. For stepwise investments, the analogous object is the threshold gain $\Delta s^{(m)} = s_j^{(m)} - s_j^{(0)}$: funding threshold m is worthwhile when

$$\Delta s^{(m)} A_j > \tau_j^{(m)}.$$

Thus higher aggregate exposure favors intervention in all three cases, while more cost-effective hardening—larger a_j, larger λ_j, larger threshold gains, or lower threshold costs—widens the region in which intervention is rational.

Intervention style. Conditional on intervention, the security investment function determines the insurer's procurement mode. Bounded linear hardening is all-or-nothing: aggregate exposure determines whether the component is hardened, but the conditional budget is just the cost of reaching saturation, $k_j^* = (t_j - s_j)/a_j$. Diminishing returns yields a tuned budget, $k_j^* = (1/\lambda_j) \ln \Theta_j$, that rises with aggregate exposure and remaining headroom. Its nonmonotone comparative static is that increasing effectiveness λ_j raises the optimal budget near the intervention boundary but lowers it once $\Theta_j > e$, because each dollar buys security faster. Stepwise investment is lumpy: the insurer funds the threshold with the largest positive net gain $\Delta_j^{(m)} = (s_j^{(m)} - s_j^{(0)})A_j - \tau_j^{(m)}$, choosing the smallest maximizer in a tie; if no threshold has positive net gain, it spends nothing.

Component targeting. Taken together, these comparative statics suggest a triage rule: rank candidate components by portfolio net gain, not exposure alone. Aggregate exposure A_j scales the benefit of any security improvement, but it must be paired with available security gain and the cost-effectiveness of hardening. Thus the strongest targets are widely used components for which hardening is both technically productive and cheap relative to the expected loss it avoids.

The insurer's distinctive role is clearest when exposure is broadly distributed. For the continuous forms, the added value of insurer provision is largest where $A_j - \max_i A_{ij}$ is wide: many policyholders depend on the component, but no one policyholder has enough at stake to internalize the full benefit. When one user dominates, the wedge nearly closes. For stepwise investments, broad exposure plays the analogous role by making a threshold collectively profitable even when no individual policyholder funds it alone.

4.4 A Symmetric-Policyholder Benchmark

The analysis above identifies conditions under which pooled insurer exposure makes component hardening rational even when no policyholder would fund it alone. To see how large this region can be, consider a symmetric benchmark. Suppose n policyholders all depend on component c_j with identical component-specific exposure $A_{ij} = A_0$. Then $\max_i A_{ij} = A_0$ and the insurer's aggregate exposure to c_j is $A_j = nA_0$. For the bounded linear model, Proposition 1 gives an insurer-only region when

$$A_0 < \frac{1}{a_j} < nA_0.$$

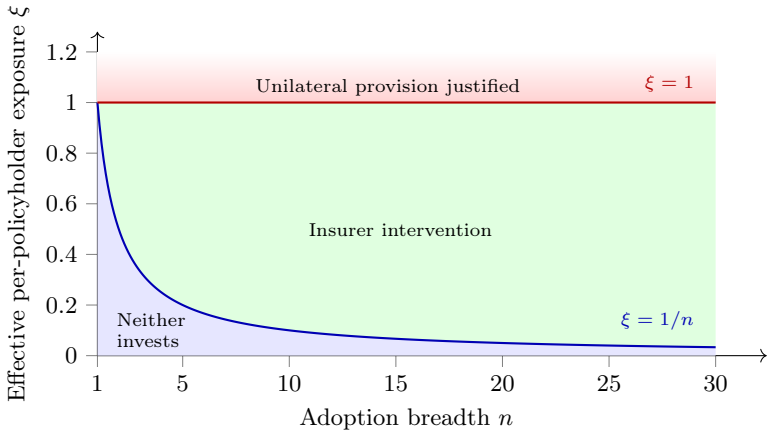


Fig. 5: In the symmetric benchmark of §4.4, ξ represents the per-policyholder return on hardening for the n policyholders who depend on the component. Below $\xi = 1/n$, pooled exposure is too small to justify insurer intervention; above $\xi = 1$, unilateral policyholder provision is already privately justified. For the bounded linear and diminishing return forms, the middle region is the insurer-only region of Proposition 1 (decentralized provision is zero while insurer provision is rational), while for stepwise investments, the region identifies thresholds that no policyholder funds alone, so zero private provision remains possible even when pooled insurer funding is rational.

Since $a_j > 0$, this is equivalent to $a_j A_0 < 1 < n a_j A_0$. Defining the dimensionless private return $\xi := a_j A_0$, the region becomes $\frac{1}{n} < \xi < 1$, and is non-empty whenever $n > 1$ and widens with adoption breadth n. The same normalized region applies to diminishing returns with $\xi := \lambda_j(t_j - s_j)A_0$. For a fixed stepwise threshold m, using $\xi := \Delta s^{(m)} A_0 / \tau_j^{(m)}$ gives the analogous no-unilateral-funding region, though cost-sharing equilibria may also exist.

Figure 5 plots these regions and shows how the pooled-intervention region expands as adoption breadth n increases. For example, if $\xi = 0.20$, no individual policyholder invests alone, but the insurer invests once the component is used by at least six policyholders. If $\xi = 0.05$, the insurer intervenes once the component is used by at least twenty-one policyholders. Thus even modest per-policyholder component risk can rationally induce insurer action when the component is sufficiently widespread.

To put dollar-scale numbers on the bounded-linear case, we write the per-policyholder component exposure as $A_0 = \beta_0 \ell$, where ℓ is the loss severity common to the symmetric population and $\beta_0 = \sigma_{ij}\beta_{ij}$ is the *attributable* conditional breach probability of Assumption 12. We anchor ℓ using IBM's 2026 *Cost of a Data Breach* report, which finds a global average breach cost of \$5.0M [15]. For the bounded-linear security investment function, we use the benchmark $a_j = 2.5 \times 10^{-6}$ dollars^{-1}, which corresponds to a \$100,000 intervention that improves security by 0.25 before saturation in the bounded linear model. This

value is not meant as an empirical estimate of security productivity but instead is meant to be a transparent benchmark for translating the normalized wedge into dollar-scale examples. To keep the calibration conservative, we use illustrative low-to-moderate values of β_0 to represent the component's attributable breach contribution. Under this parameterization, $\xi = a_j A_0 = a_j \beta_0 \ell$. Then whenever $\xi < 1$, the minimum adoption breadth needed for insurer-only intervention is $n_{\min} = \left\lfloor \frac{1}{\xi} \right\rfloor + 1$. Such illustrative scenarios are shown in Table 1.

Table 1: Illustrative bounded-linear calibration with $a_j = 2.5 \times 10^{-6}\,\text{dollars}^{-1}$. The insurer-only region exists when $\xi < 1$ and $n \geq n_{\min}$; when $\xi \geq 1$, unilateral investment is already justified.

	ℓ	β_0	Exposure A_0	Private return ξ	Min. n
Conservative	\$1.0M	0.5%	\$5,000	0.0125	81
Moderate	\$5.0M	2%	\$100,000	0.25	5
Elevated	\$5.0M	5%	\$250,000	0.625	2
Severe	\$5.0M	10%	\$500,000	1.25	N/A

These scenarios suggest three qualitatively different outcomes. In the conservative case, insurer-only intervention requires very wide sharing: the component must appear across at least 81 policyholders before pooled exposure outweighs the cost of intervention. In the moderate case, however, the insurer-only region appears once the component is shared by only five policyholders. In the elevated case, sharing by just two policyholders is enough. Only in the severe regime does the insurer-only region disappear, because per-policyholder exposure is already large enough that decentralized actors would individually invest. Thus insurer-only intervention is neither universal nor is it knife-edge—it is unlikely for low-loss, weakly attributable components, but becomes plausible once shared components carry moderate attributable breach risk.

5 Limitations

5.1 Model Limitations

Our modeling choices are tradeoffs that come with necessary limitations. First, the relationship between security investment k and security outcome $S_j(k)$ is generally unknown. In response, we study three classes of plausible functions but also note that there may be an incentive for product vendors to truthfully report the shape of the investment function to insurers (discussed in §6). Regardless, our candidate investment functions are deterministic whereas real-world security investment may have non-deterministic outcomes. Second, for simplicity, we model the threats and product vendors as passive, though an active adversary or strategic vendor may use knowledge of the P4P system to their advantage. Third, to keep our model focused on P4P, we consider only a single insurer under

a policy of mandatory insurance, thereby removing typical insurance-modeling considerations such as adverse selection, premiums, and deductibles. Fourth, we do not consider the effects of catastrophic losses.

5.2 Strategic and Adversarial Behavior

Healthcare P4P also illustrates how cyber P4P might be gamed. For instance, insurers and care providers may use the mechanism as an avenue for collusion. Consider an insurer and care provider who agree to a 4% increase in care costs. An insurer may split this into a 2.5% service-rate increase plus a 1.5% P4P payment tied to a near-guaranteed performance target[4]. This can make it appear as if the insurer's plans are cheaper than their competitors, which may attract new customers while leaving open whether the bonus is meaningfully contingent on performance. Performance measures that are easy to satisfy can therefore be abused to turn risk-sharing for quality improvement into a back channel for less-visible payments. Healthcare P4P offers other examples of adverse behavior and effects as well:

Kingmaker effects. Through selective investment, the insurer can act as a kingmaker among vendors, especially when a few high-value policyholders dominate portfolio exposure and share homogeneous technology stacks. The insurer may fund vendors that matter to those policyholders rather than vendors whose components create the largest portfolio-wide security gains.

Monoculture. Relatedly, the scheme could induce a technology-stack monoculture, where policyholders are incentivized to adopt the preferred (or most widely used) components, raising their probability of payout should a widespread adverse event occur.

Misrepresentation and audit gaming. A central strategic risk is information asymmetry. Vendors can understate baseline security, overstate the marginal effect of investment, or stage temporary hardening that passes an audit but is later reverted. Policyholders can likewise misreport their technology stacks to reduce premiums or hide high-risk dependencies. We discuss this further in §6.

Fraud and moral hazard. Absent credible tracking of funded work, vendors can misallocate funds or solicit multiple insurers to pay for the same fix. Cyber insurance can also weaken policyholders' incentives to maintain their own internal controls, so component-level P4P should complement rather than replace ordinary underwriting, audits, and security requirements.

6 Future Work

Although we focus on pay-for-performance in this work, health insurance provides many more examples of mechanisms for shaping incentives. For example, health

[4] For example, in a genuine risk-sharing model, an insurer might agree to share savings with a care provider if they reduce, say, knee-replacement readmission rates by 10%. But in practice, such deals can devolve into "gimmes" pegged to all-but-guaranteed targets, like requiring an 80% physician handwashing rate, which is self-reported.

insurers sometimes establish a "preferred vendor" for certain services to secure a built-in discount for policyholders for a given period of time. Similar mechanisms may apply to cyber insurance as well: If insurers can incentivize better product security through the promise of increased sales volume, then component vendors can compete for preferred vendor status on the basis of security without requiring direct investment from the insurers at all. Policyholders, in turn, would favor preferred components to lower their premiums, so incentives would flow in both directions. We assume the agents in our model are appropriately motivated, and reserve the design of such a preferred-vendor mechanism for future work.

A deeper question is the credibility of the elicited security investment function itself. Our model assumes that insurers can elicit $S_j(\cdot)$ from vendors, but in practice a vendor competing for funding may overstate the effectiveness of their proposed hardening in an effort to secure hardening funding over their competitors. Yet at the same time, vendors may also be *dis*incentivized from overstating the effectiveness of security investment if vendors are competing for preferred vendor status as well, since an insurer is unlikely to pick a preferred component known to be full of vulnerabilities. These two forces are then in opposition to one another and may be able to cancel each other out. Designing an instrument for the truthful elicitation of $S_j(\cdot)$ therefore may require some degree of co-investment, clawback, audits of verifiable outcomes, vendor competition, and reputation, and becomes a mechanism design problem in its own right. Our work brings this new challenge into view, and is a natural direction for future work.

Finally, our numerical evaluation assumes a symmetric population of policyholders. Characterizing the insurer-only region under heterogeneous exposure across policyholders, where the aggregate no longer factors as nA_0, is a natural extension of this work.

7 Conclusion

In this paper we asked whether the pay-for-performance logic that health insurers use to fund better care can be carried over to cyber insurance. To study this, we developed a component-level model in which an insurer aggregates exposure across its policyholder base and chooses security investments to minimize expected losses, and derived closed-form conditions for when such investment is rational.

We highlight three takeaways. First, a cyber insurer can be the missing provider of a shared security public good. Component hardening benefits every policyholder that depends on the component, but a policyholder deciding alone counts only its own avoided loss. The insurer counts avoided losses across its portfolio, which can make funding rational even when private incentives are weak. Second, this public goods problem has a simple max-versus-sum structure. Decentralized incentives turn on the largest individual exposure, $\max_i A_{ij}$, while insurer incentives turn on aggregate exposure, A_j. The insurer's role is clearest where exposure is broadly distributed rather than concentrated in one dominant user. For the continuous forms, this gives a true insurer-only region; for stepwise investments, it appears as a threshold-coordination problem. Third, the symmetric

benchmark gives the scale of the wedge. Insurer-only intervention is unlikely for low-loss, weakly attributable components, but becomes plausible once shared components carry moderate attributable breach risk. The security investment function then determines how hardening is procured, whether through saturation, a tuned budget, or a discrete threshold investment.

Beyond the specific conditions derived here, the broader contribution is a bridge between the world of cybersecurity with its market failures and misaligned incentives, and the world of healthcare and health insurance which has decades' worth of hard-earned lessons to pull from. Instruments like pay-for-performance offer paradigms for how agents like insurers can be more than passive risk-pricers and instead be active investors in the security of shared technology components.

Acknowledgments. We would like to thank the anonymous reviewers for their careful reading and constructive feedback.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Anderson, R.: Why information security is hard—an economic perspective. In: Seventeenth annual computer security applications conference. pp. 358–365. IEEE (2001)
2. Anderson, R., Moore, T.: The economics of information security. Science **314**(5799), 610–613 (2006)
3. Awiszus, K., Knispel, T., Penner, I., Svindland, G., Voß, A., Weber, S.: Modeling and pricing cyber insurance. European Actuarial Journal **13**(1), 1–53 (2023)
4. Bergstrom, T., Blume, L., Varian, H.: On the private provision of public goods. Journal of public economics **29**(1), 25–49 (1986)
5. Böhme, R., Schwartz, G.: Modeling Cyber-Insurance: Towards a Unifying Framework. In: Workshop on the Economics of Information Security (WEIS) (2010)
6. Boyens, J., Smith, A., Bartol, N., Winkler, K., Holbrook, A., Fallon, M.: Cybersecurity Supply Chain Risk Management Practices for Systems and Organizations (SP 800-161 Rev. 1). Tech. rep., National Institute of Standards and Technology (2022)
7. Centers for Medicare & Medicaid Services: "hospital readmissions reduction program" (2024), https://www.cms.gov/medicare/quality/value-based-programs/hospital-readmissions
8. Centers for Medicare & Medicaid Services: "value-based programs" (2024), https://www.cms.gov/medicare/quality/value-based-programs
9. Centers for Medicare & Medicaid Services: "traditional mips overview" (2025), https://qpp.cms.gov/mips/traditional-mips, accessed November 4, 2025
10. Cybersecurity and Infrastructure Security Agency: Software Bill of Materials (SBOM) (2024), https://www.cisa.gov/sbom, accessed March 13, 2026
11. Finifter, M., Akhawe, D., Wagner, D.: An empirical study of vulnerability rewards programs. In: 22nd USENIX Security symposium (USENIX Security 13). pp. 273–288 (2013)
12. Gordon, L.A., Loeb, M.P.: The economics of information security investment. ACM Transactions on Information and System Security (TISSEC) **5**(4), 438–457 (2002)

13. Grossklags, J., Christin, N., Chuang, J.: Secure or insure? A game-theoretic analysis of information security games. In: Proceedings of the 17th international conference on World Wide Web. pp. 209–218 (2008)
14. Huang, Z., Biczók, G., Liu, M.: Incentivizing Secure Software Development: The Role of Voluntary Audit and Liability Waiver. ACM Transactions on Privacy and Security **28**(4), 1–30 (2025)
15. IBM: Cost of a Data Breach Report 2026 (2026), https://www.ibm.com/reports/data-breach
16. Khalili, M.M., Liu, M., Romanosky, S.: Embracing and controlling risk dependency in cyber-insurance policy underwriting. Journal of Cybersecurity **5**(1), tyz010 (2019)
17. Lelarge, M., Bolot, J.: Economic incentives to increase security in the internet: The case for insurance. In: IEEE INFOCOM 2009. pp. 1494–1502. IEEE (2009)
18. Liu, M.: Cyber Insurance as an Incentive Mechanism for Cybersecurity. Morgan & Claypool (2021)
19. Maillart, T., Zhao, M., Grossklags, J., Chuang, J.: Given enough eyeballs, all bugs are shallow? Revisiting Eric Raymond with bug bounty programs. Journal of Cybersecurity **3**(2), 81–90 (2017)
20. National Institute of Standards and Technology: Secure Software Development Framework (SSDF) Version 1.1: Recommendations for Mitigating the Risk of Software Vulnerabilities (SP 800-218). Tech. rep. (2022). https://doi.org/10.6028/NIST.SP.800-218
21. Ohm, M., Plate, H., Sykosch, A., Meier, M.: Backstabber's knife collection: A review of open source software supply chain attacks. In: International Conference on Detection of Intrusions and Malware, and Vulnerability Assessment. pp. 23–43. Springer (2020)
22. Olson, M.: The Logic of Collective Action: Public Goods and the Theory of Groups. Harvard University Press, Cambridge, MA (1965)
23. Romanosky, S., Ablon, L., Kuehn, A., Jones, T.: Content analysis of cyber insurance policies: how do carriers price cyber risk? Journal of Cybersecurity **5**(1), tyz002 (2019)
24. Varian, H.: System reliability and free riding. In: Economics of information security, pp. 1–15. Springer (2004)
25. Zhao, M., Grossklags, J., Liu, P.: An empirical study of web vulnerability discovery ecosystems. In: Proceedings of the 22nd ACM SIGSAC conference on computer and communications security. pp. 1105–1117 (2015)
26. Zimmermann, M., Staicu, C.A., Tenny, C., Pradel, M.: Small world with high risks: A study of security threats in the npm ecosystem. In: 28th USENIX Security symposium (USENIX security 19). pp. 995–1010 (2019)

A Supplementary Derivations

A.1 Decentralized Provision Equilibria

Fix component c_j and its dependents $\{\,\pi_i : \beta_{ij} > 0\,\}$, with individual exposures $A_{ij} = \sigma_{ij}\beta_{ij}\ell_i$ and aggregate $A_j = \sum_i A_{ij}$. In the decentralized game of Section 4.2, each dependent contributes $k_i \geq 0$, the total $K = \sum_i k_i$ sets the shared security $S_j(K)$, and π_i minimizes its own expected loss plus outlay, $J_i(k_i; k_{-i}) = \big(1 - S_j(K)\big)A_{ij} + k_i$, the per-policyholder analogue of the component objective (Eq. 5) with A_{ij} in place of A_j. Each policyholder weighs only

its own stake, while an insurer or planner weighs A_j. For the continuous forms, decentralized provision turns on $\max_i A_{ij}$; the stepwise form additionally has a threshold-coordination problem. Write $B_j := \lambda_j(t_j - s_j)$ and $\Delta s^{(m)} := s_j^{(m)} - s_j^{(0)}$.

Proposition 2 (Equilibria of the contribution game).

1. (Bounded linear.) *If $a_j \max_i A_{ij} < 1$, the all-zero profile is the unique equilibrium; at equality it is the selected outcome under our indifference convention. If $a_j \max_i A_{ij} > 1$, every equilibrium total equals the saturation level $\bar{k} = (t_j - s_j)/a_j$, and the equilibria are exactly the profiles summing to $\bar{k}$ whose contributors all satisfy $a_j A_{ij} > 1$.*
2. (Diminishing returns.) *The equilibrium total is unique and equals $\max_i K_i^\dagger$, where $K_i^\dagger = 0$ if $B_j A_{ij} \le 1$ and $K_i^\dagger = \frac{1}{\lambda_j} \ln(B_j A_{ij})$ otherwise. It is funded by the highest-exposure dependents while the rest free-ride. Provision is positive iff $\max_i A_{ij} > 1/B_j$.*
3. (Stepwise.) *If $\Delta s^{(m)} \max_i A_{ij} \le \tau_j^{(m)}$ for every m, the all-zero profile is an equilibrium—but not necessarily the only one. For a single jump $(\Delta s, \tau)$ with $\Delta s \max_i A_{ij} < \tau < \Delta s\, A_j$, there also exist cost-sharing equilibria that fund the jump.*

Dominance in the continuous forms. Both continuous forms have $\partial J_i/\partial k_i = 1 - S_j'(K)A_{ij}$ with S_j' non-increasing, so the marginal return to contributing is greatest at $K = 0$. If the largest stake is strictly below the break-even test there, withholding is dominant and no-provision is the *unique* equilibrium; equality is handled by the indifference convention. Thus cases 1 and 2 need no separate treatment of cost-sharing profiles, and the strict no-provision region used by Proposition 1 needs neither complete information nor equilibrium selection.

Bounded linear. Here $\partial J_i/\partial k_i = 1 - a_j A_{ij}$ below saturation and 1 above it, so π_i contributes only if $a_j A_{ij} > 1$, and then on every dollar up to $\bar{k}$. Thus the total saturates iff $a_j \max_i A_{ij} > 1$, while $a_j \max_i A_{ij} < 1$ makes no-provision unique; equality is selected as no-provision. When provision is positive, the split is indeterminate: withdrawing ε below saturation costs a contributor $a_j A_{ij} \varepsilon > \varepsilon$, so any profile summing to $\bar{k}$ whose contributors all clear the bar is an equilibrium.

Diminishing returns. Here $\partial J_i/\partial k_i = 1 - B_j e^{-\lambda_j K} A_{ij}$ and $\partial^2 J_i/\partial k_i^2 \ge 0$, so the first-order condition is sufficient. Policyholder π_i's best response is $k_i = \max\{0,\, K_i^\dagger - K_{-i}\}$, with $K_{-i} = \sum_{l \ne i} k_l$. No profile below $\max_i K_i^\dagger$ survives because a maximizer would add, and none above it survives because every contributor would withdraw; the equilibrium total is $\max_i K_i^\dagger$, funded by the highest-exposure stakeholder up to ties [4].

Stepwise investment. The stepwise form admits no dominance argument. It is a *threshold* public good: since a sub-threshold contribution is pure loss, the only profitable unilateral deviation from the all-zero profile is to fund a jump m single-handedly, $k_i \ge \tau_j^{(m)}$, which pays iff $\Delta s^{(m)} A_{ij} > \tau_j^{(m)}$. Hence when

$\Delta s^{(m)} \max_i A_{ij} \leq \tau_j^{(m)}$ for *every* m, no policyholder funds any jump alone and the no-provision profile is a Nash equilibrium.

For a single jump $(\Delta s, \tau)$, consider shares k_i summing to τ. A contributor's only useful deviation is to withdraw its share, saving k_i but losing $\Delta s\, A_{ij}$, so no withdrawal is profitable iff $k_i \leq \Delta s\, A_{ij}$. Strictly slack shares exist whenever $\Delta s\, A_j > \tau$. Thus when $\Delta s \max_i A_{ij} < \tau < \Delta s\, A_j$, no-provision and cost-sharing provision are both equilibria—coordination failure, not impossibility. With several thresholds, no-provision persists; funded-threshold equilibria must also rule out deviations to other thresholds.

Information structure. Only the stepwise case turns on coordination, and incomplete information makes sub-threshold contributions no easier. The positive-provision allocation claims assume complete information; the main text relies only on the no-provision and insurer-provision comparisons above.

A.2 Proof of Proposition 1 and Stepwise Claim

Proof. For the continuous forms, Proposition 2 gives zero decentralized provision in every equilibrium unless $\max_i A_{ij}$ exceeds the break-even exposure, $1/a_j$ or $1/B_j$. Section 4.1 gives insurer provision under the same test with A_j replacing $\max_i A_{ij}$. Thus the insurer-only region places the break-even exposure inside $(\max_i A_{ij},\ A_j)$, giving the two cases. This interval is non-empty iff $A_j > \max_i A_{ij}$, and is empty when a single dependent carries all exposure.

For the stepwise statement in Section 4.2, Proposition 2 gives no unilateral threshold funding and no-provision as an equilibrium when every ratio $\tau_j^{(m)}/\Delta s^{(m)}$ exceeds $\max_i A_{ij}$. Section 4.1 gives insurer funding when at least one such ratio falls below A_j. Cost-sharing equilibria may coexist, so this is threshold coordination rather than an insurer-only result in the same equilibrium sense. □