Who Pays to Secure Shared Components?

A Pay-for-Performance Model for Cyber Insurance

John Chester · Adam Hastings
Fordham University · GameSec 2026

FORDHAM UNIVERSITY



Problem Statement:

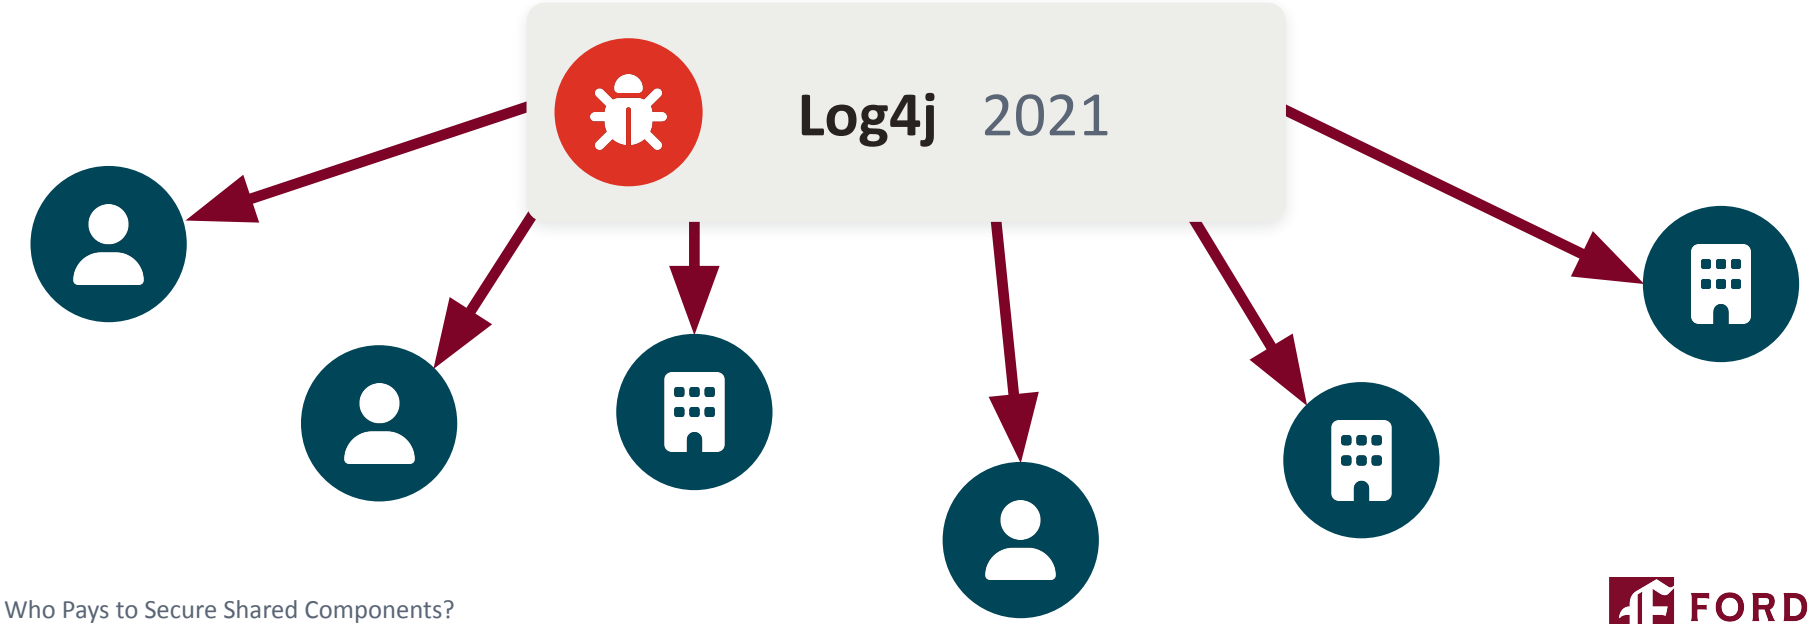
Technology is built by integrating upstream components

A bug in one upstream component can affect billions of devices

Problem Statement:

Technology is built by integrating upstream components

A bug in one upstream component can affect billions of devices



Problem Statement:

Technology is built by integrating upstream components

A bug in one upstream component can affect billions of devices



Problem Statement:

Technology is built by integrating upstream components

A bug in one upstream component can affect billions of devices



How can we improve the security of widely-used shared upstream technology components?

Let the market decide?

Security is not measurable
+ Massive info asymmetry
= Market for lemons

Make vendors liable for vulns?

Who is liable for OSS?

Impossible (?) to make bug-free products

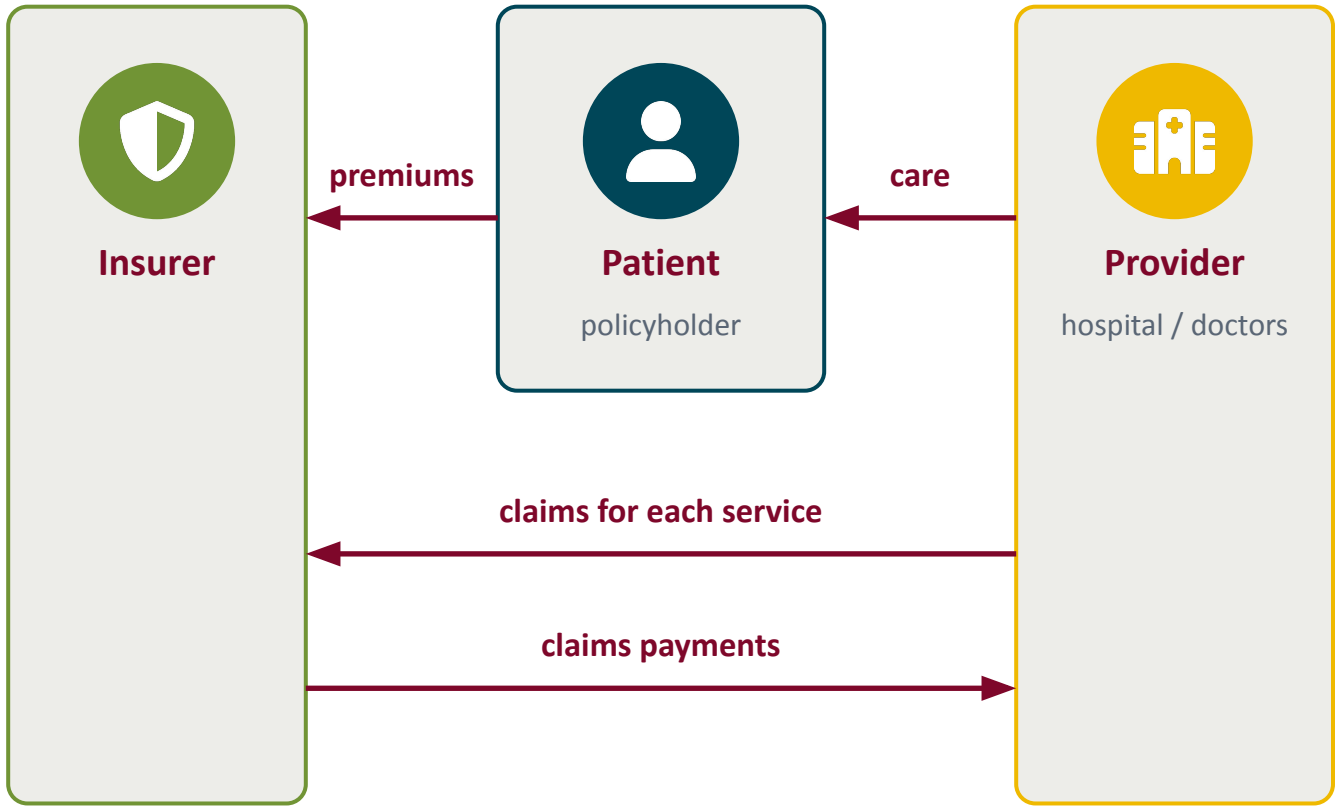
Let exposed users pay for security hardening themselves?

More info asymmetry
+ Software supply chain security is a mess
+ Incentive to free ride

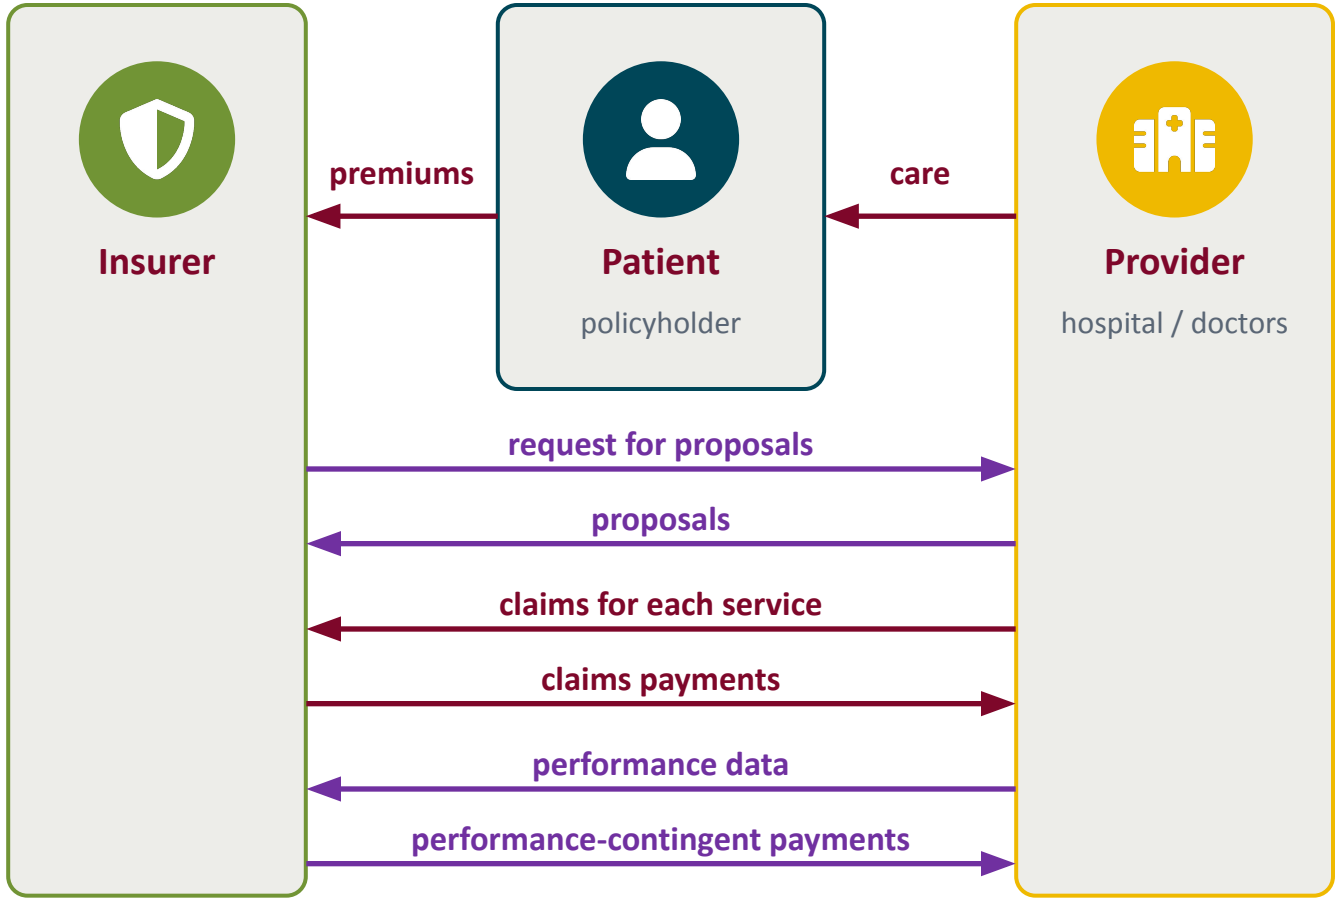
Something else?

Borrow P4P ideas from health insurance

Health insurance 101



Improving quality of care with P4P



Purple arrows are the P4P mechanism layered on top of the ordinary claims flow.

Pay-for-performance (P4P)

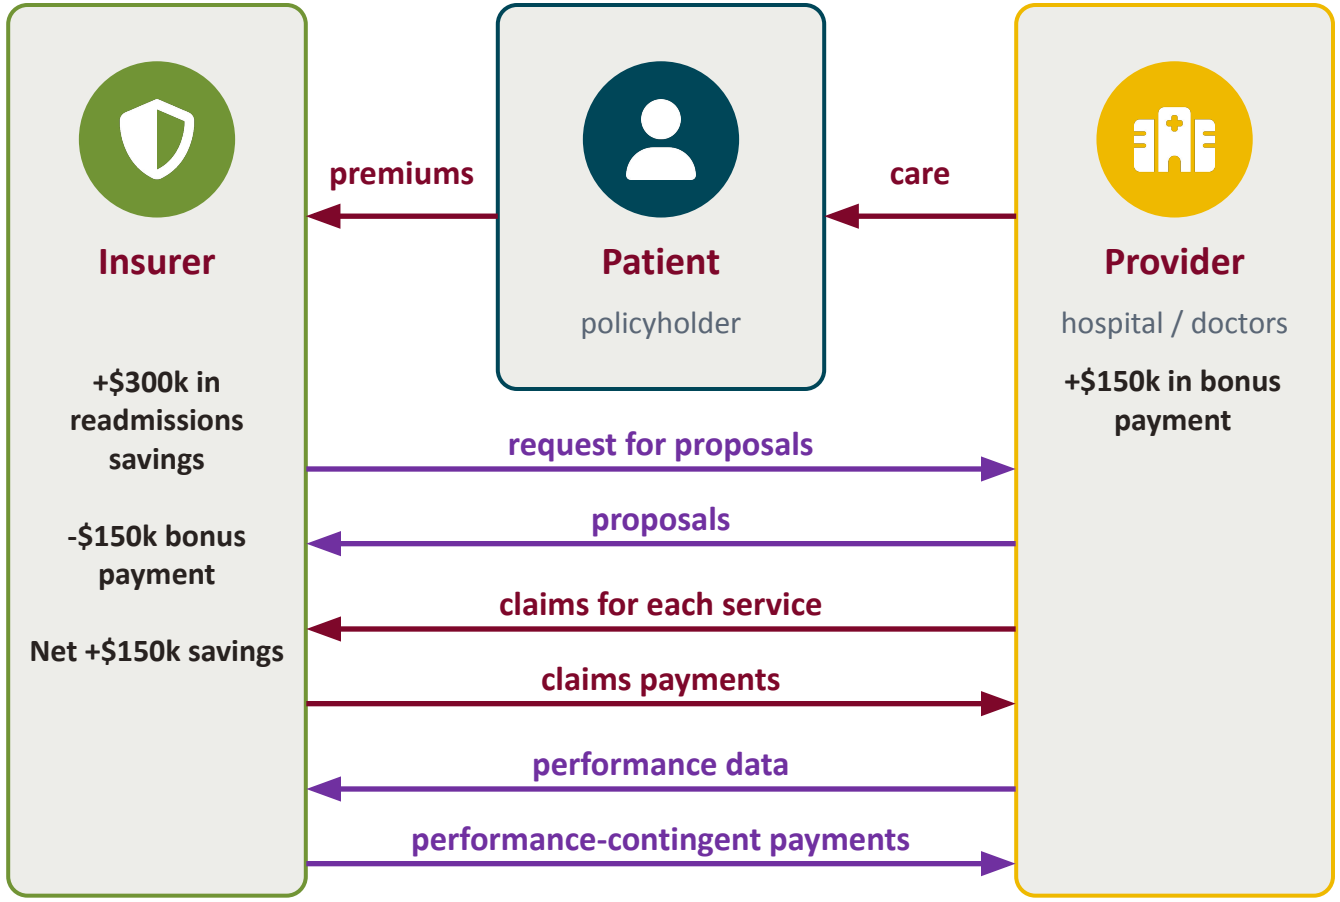
Insurers pay hospitals a bonus payment for achieving an agreed-upon level of healthcare quality

Upstream investment in patient care reduces downstream costs for insurers

~20 yrs

of national-scale P4P at CMS: HRRP, Hospital VBP, MIPS. Private insurers run their own.

Example: Knee replacement surgeries



Purple arrows are the P4P mechanism layered on top of the ordinary claims flow.

Suppose 10% of surgeries have complications and require hospital readmission

Readmissions cost the insurer extra money

P4P idea: Pay hospital a bonus if readmissions rate falls below 5%

Incentivize better outcomes and split the savings

Healthcare incentives mirror cyber incentives

- Insurers want to limit claims payouts
- Healthier policyholders = more profit for insurers
- Investing in upstream care quality can prevent downstream losses

This paper: Does P4P work for security?

Healthcare P4P

Health Insurers incentivize higher-quality care, reducing downstream costs

Cyber P4P?

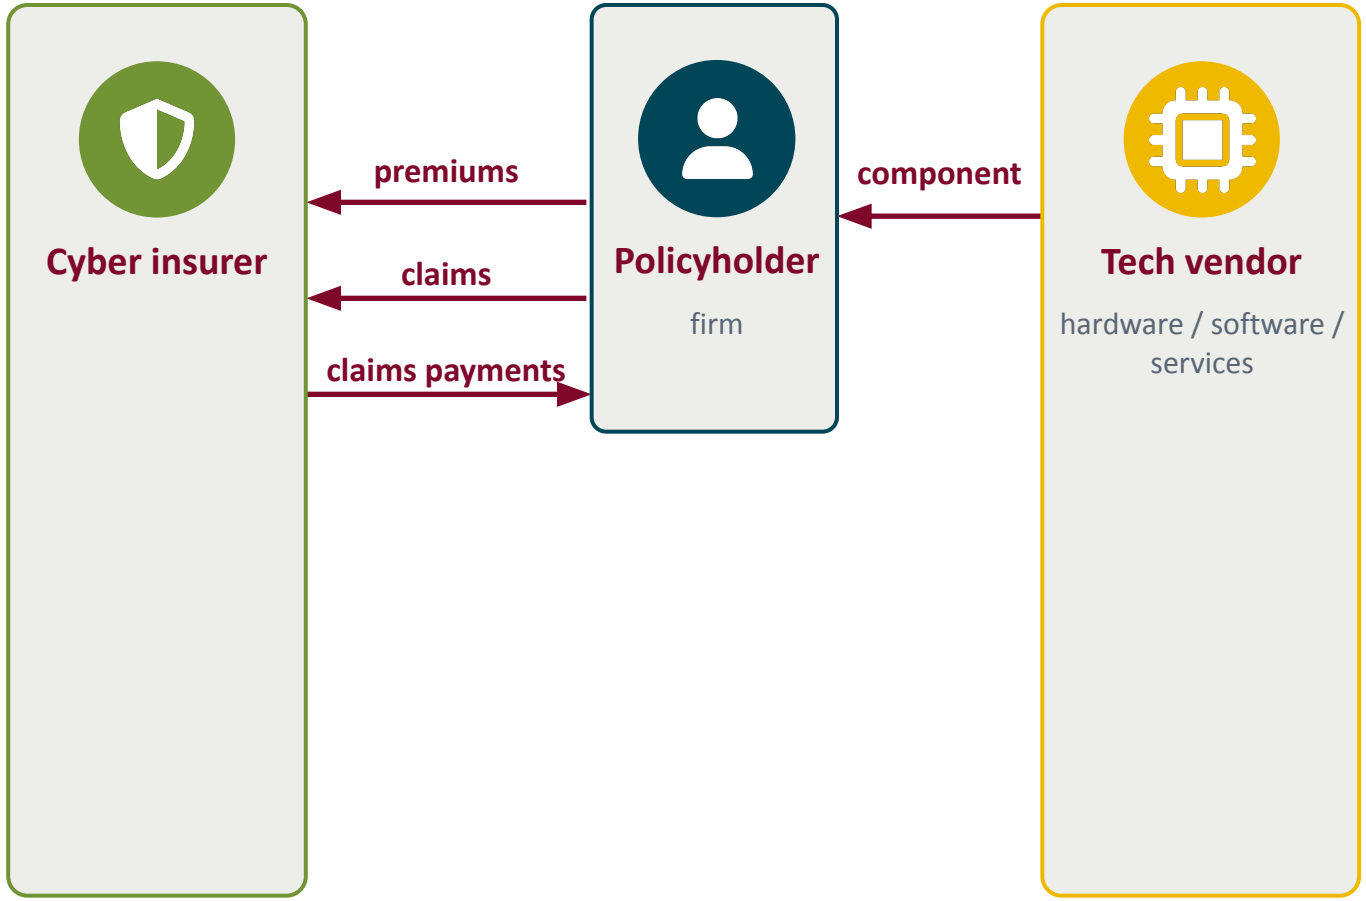
Cyber insurers incentivize more secure upstream tech components, reducing downstream costs

Research questions:

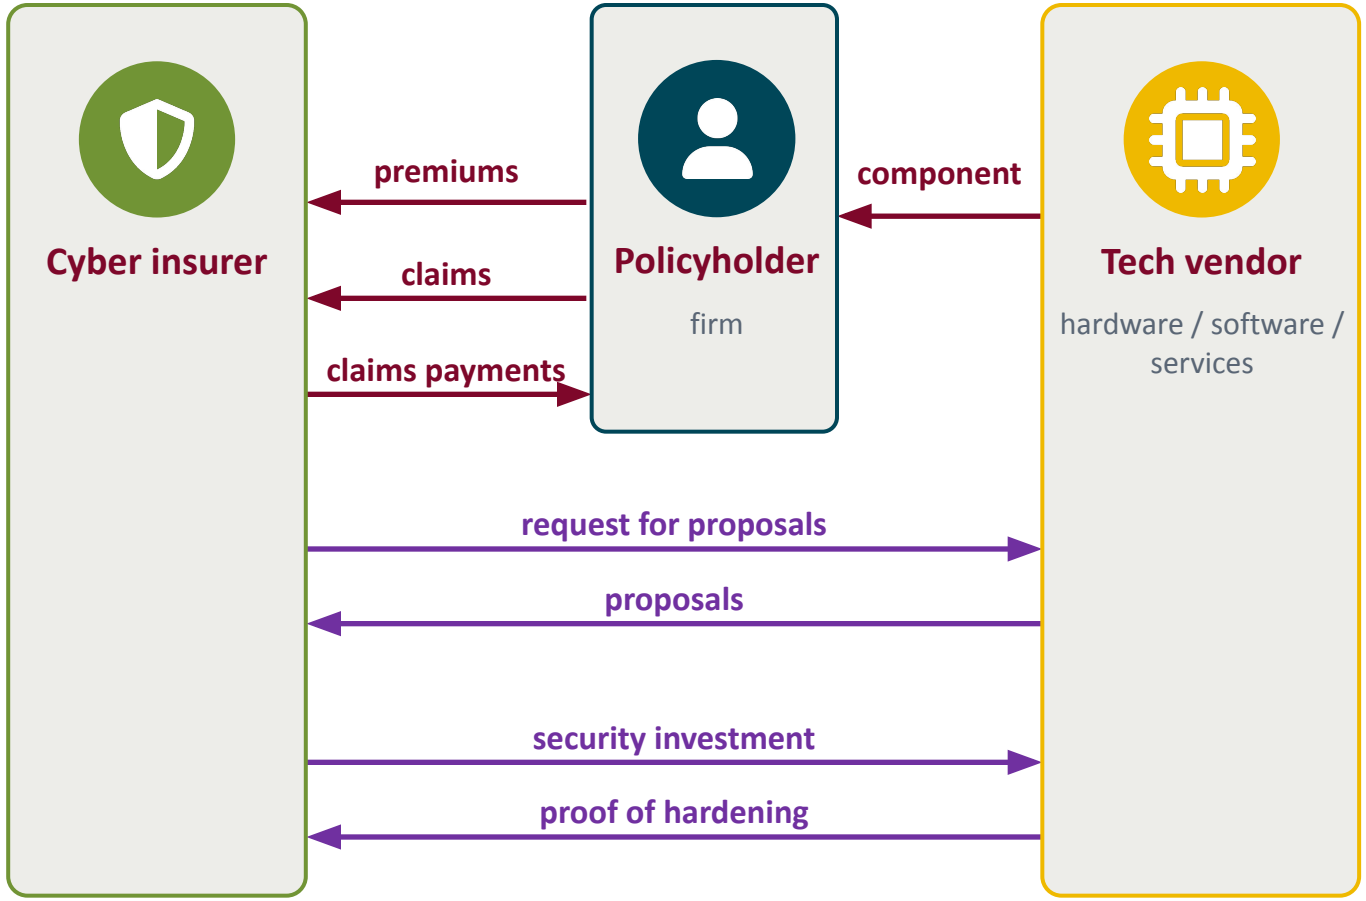
1. **When is it rational for cyber insurers to invest in upstream technology components?**

2. **Should cyber insurers take an active role in patching security holes?**

Cyber insurance 101



A proposed cyber P4P system



Purple arrows are the P4P mechanism

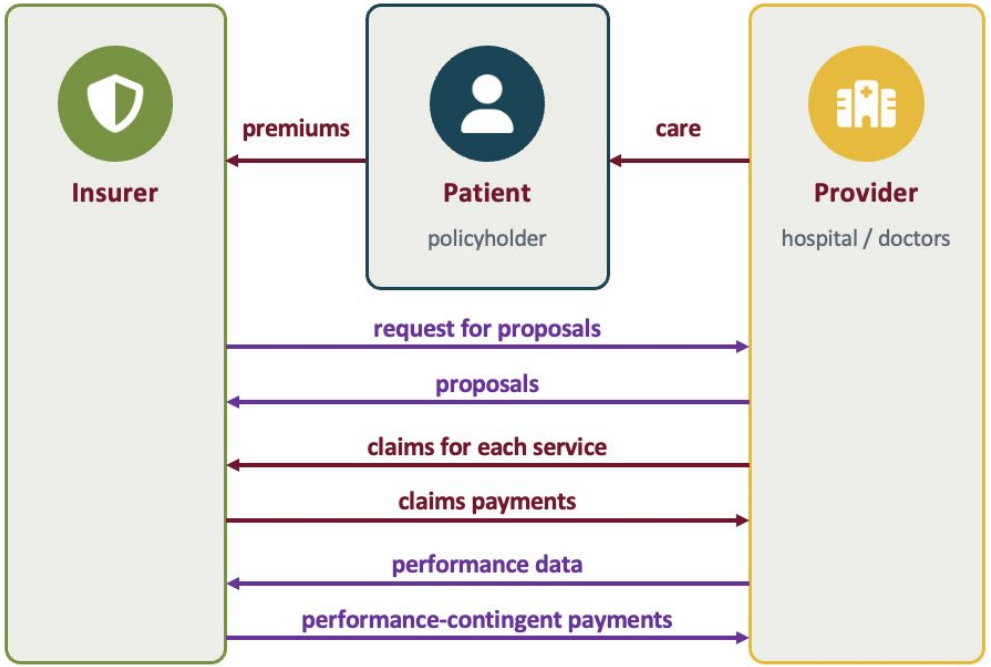
Cyber P4P

Insurers invest in widely-used technology components like OSS

Upstream investment in shared component helps prevent systemic losses

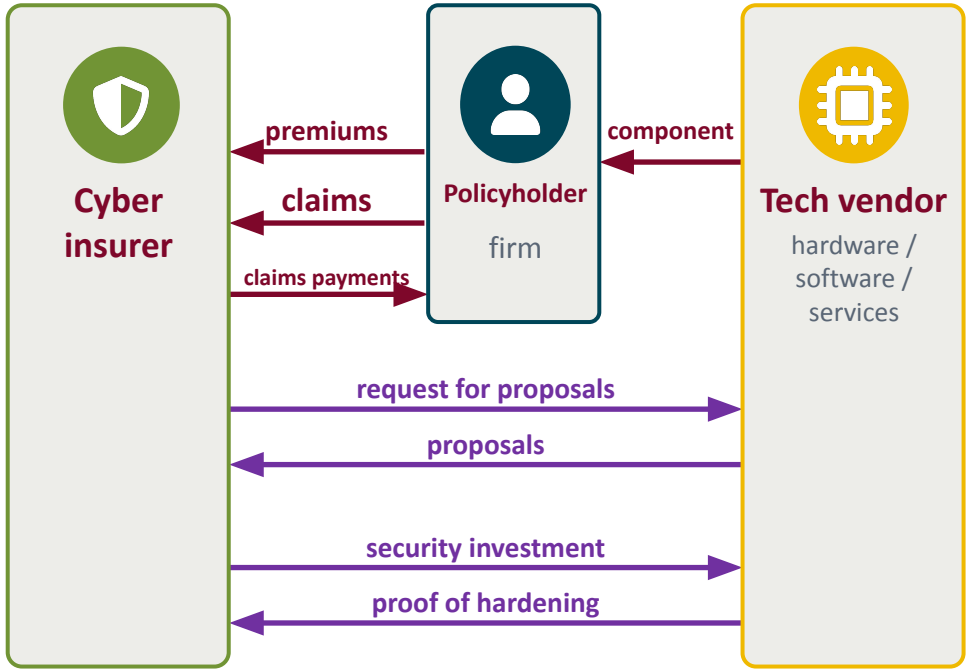
Healthcare P4P and Cyber P4P

Similar but not identical



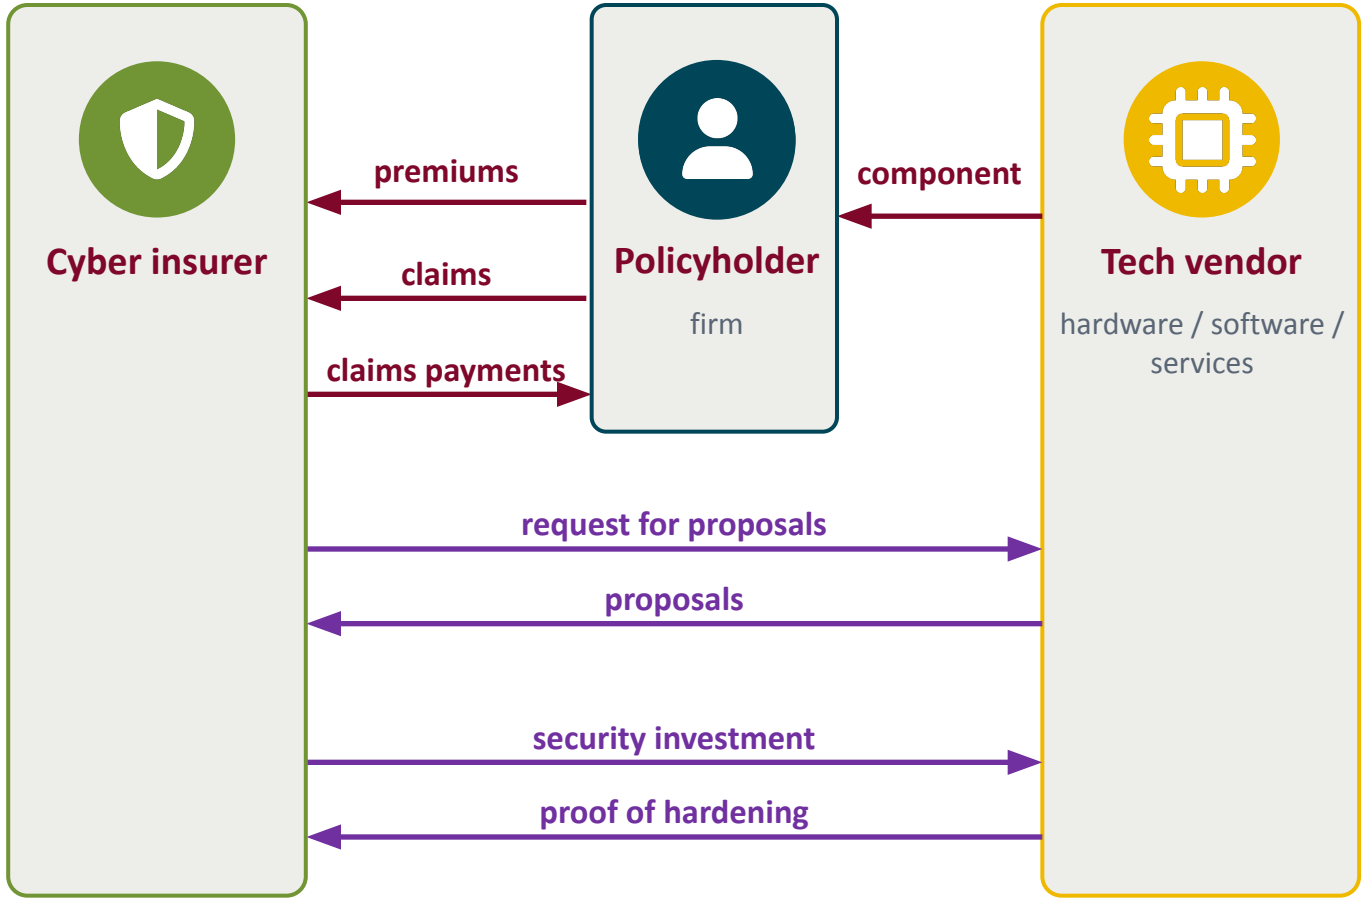
Purple arrows are the P4P mechanism layered on top of the ordinary claims flow.

Healthcare P4P



Cyber P4P (this work)

Example: Apache Log4j vulnerability (2023)



Purple arrows are the P4P mechanism

Vuln in open source Log4j Java library allowed RCE in millions of affected devices

~$100M in insured losses

Cyber P4P idea: Could this loss have been prevented by giving money to Apache?

Pay for security audits of Apache OSS, provide funds for bug bounties, etc.

What makes the insurer an interesting leverage point?

Access to data

Insurers get policyholders' IT asset inventory when underwriting

→ Can identify which shared tech components pose the most systemic risk

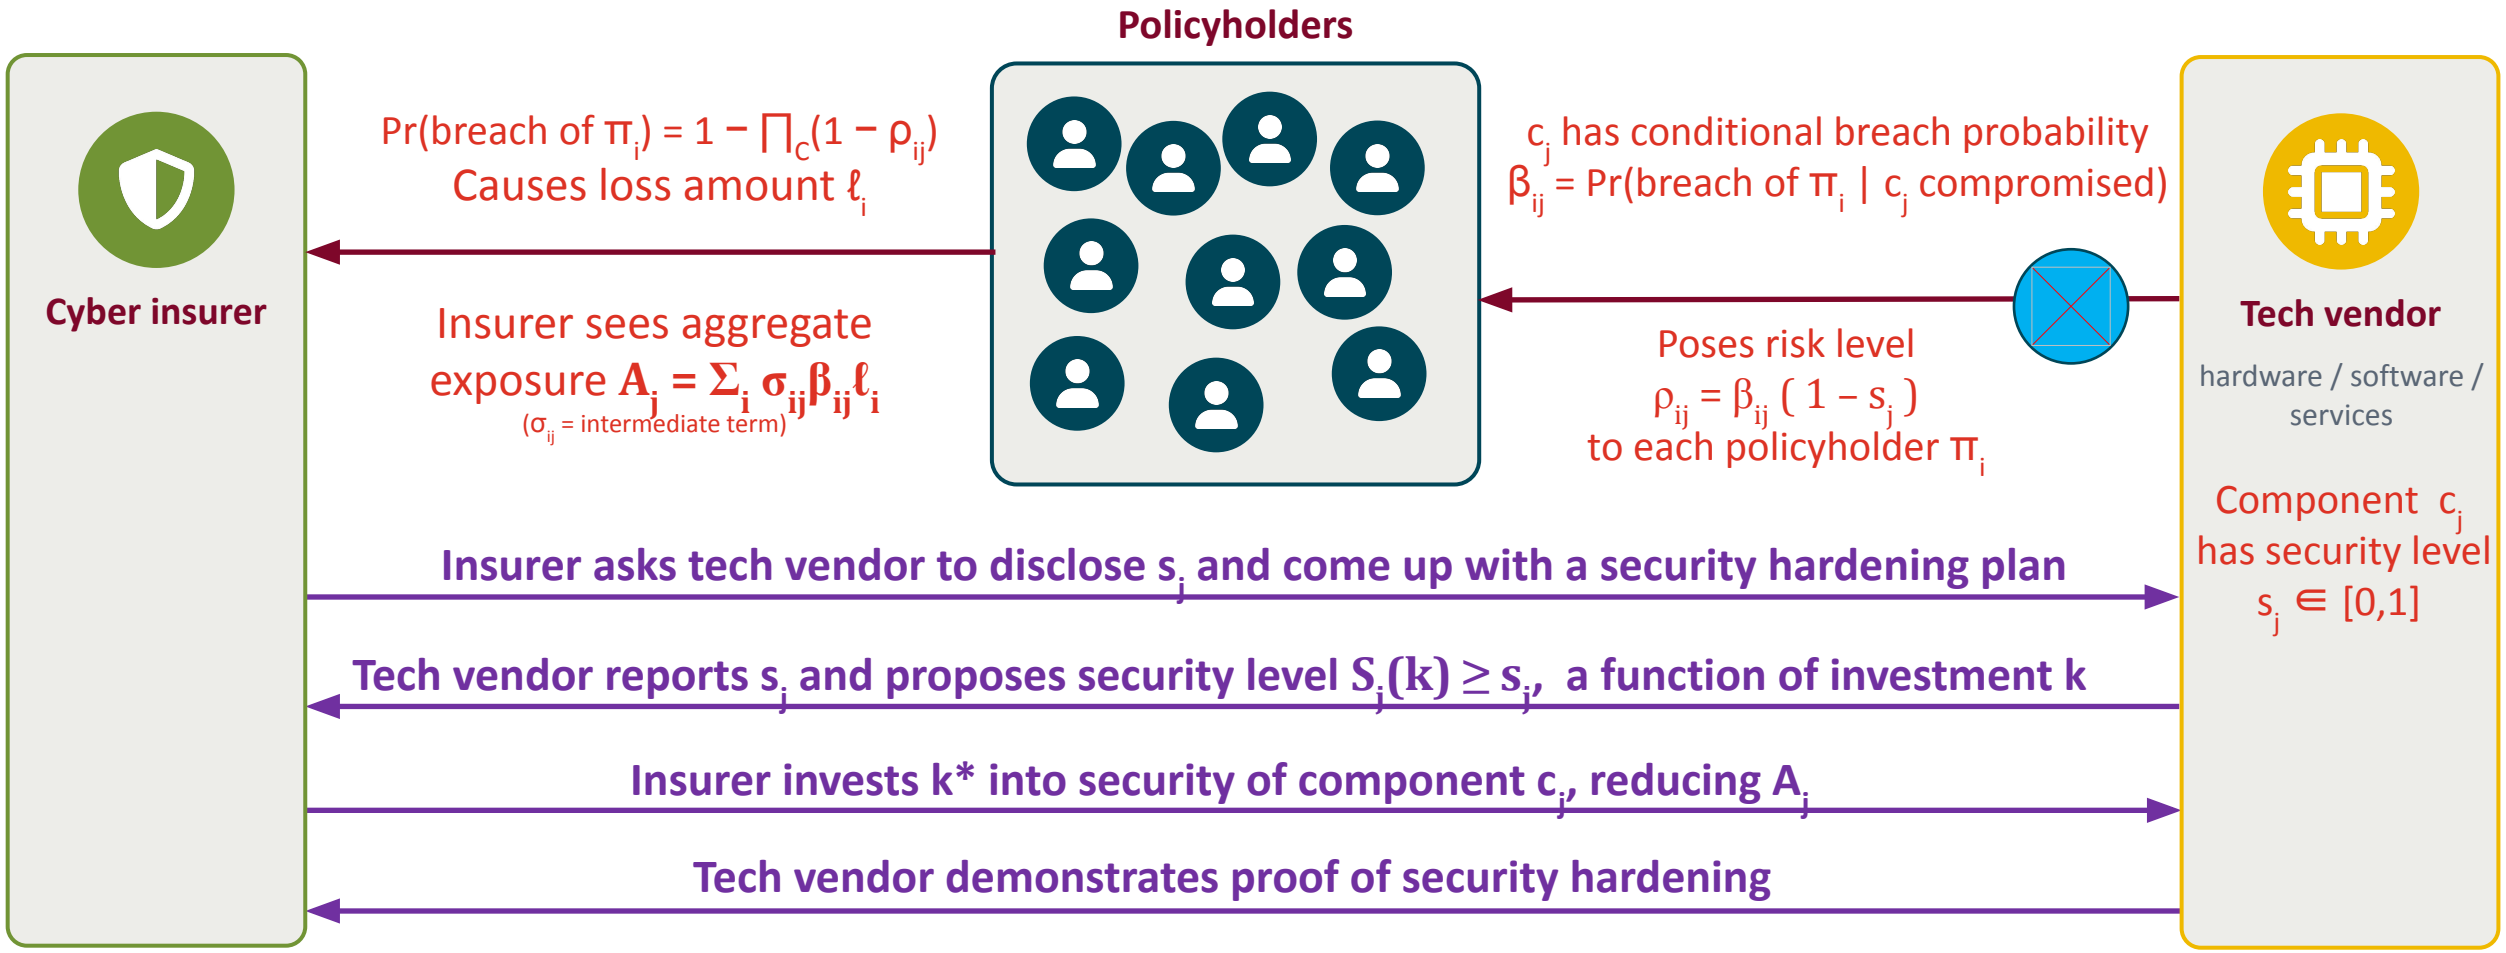
Incentivized to mitigate systemic + catastrophic vulns

Creates a new info channel for disclosing and patching and vulns!

Previously, tech vendors might not want to disclose if their products are buggy + full of vulns → info asymmetry

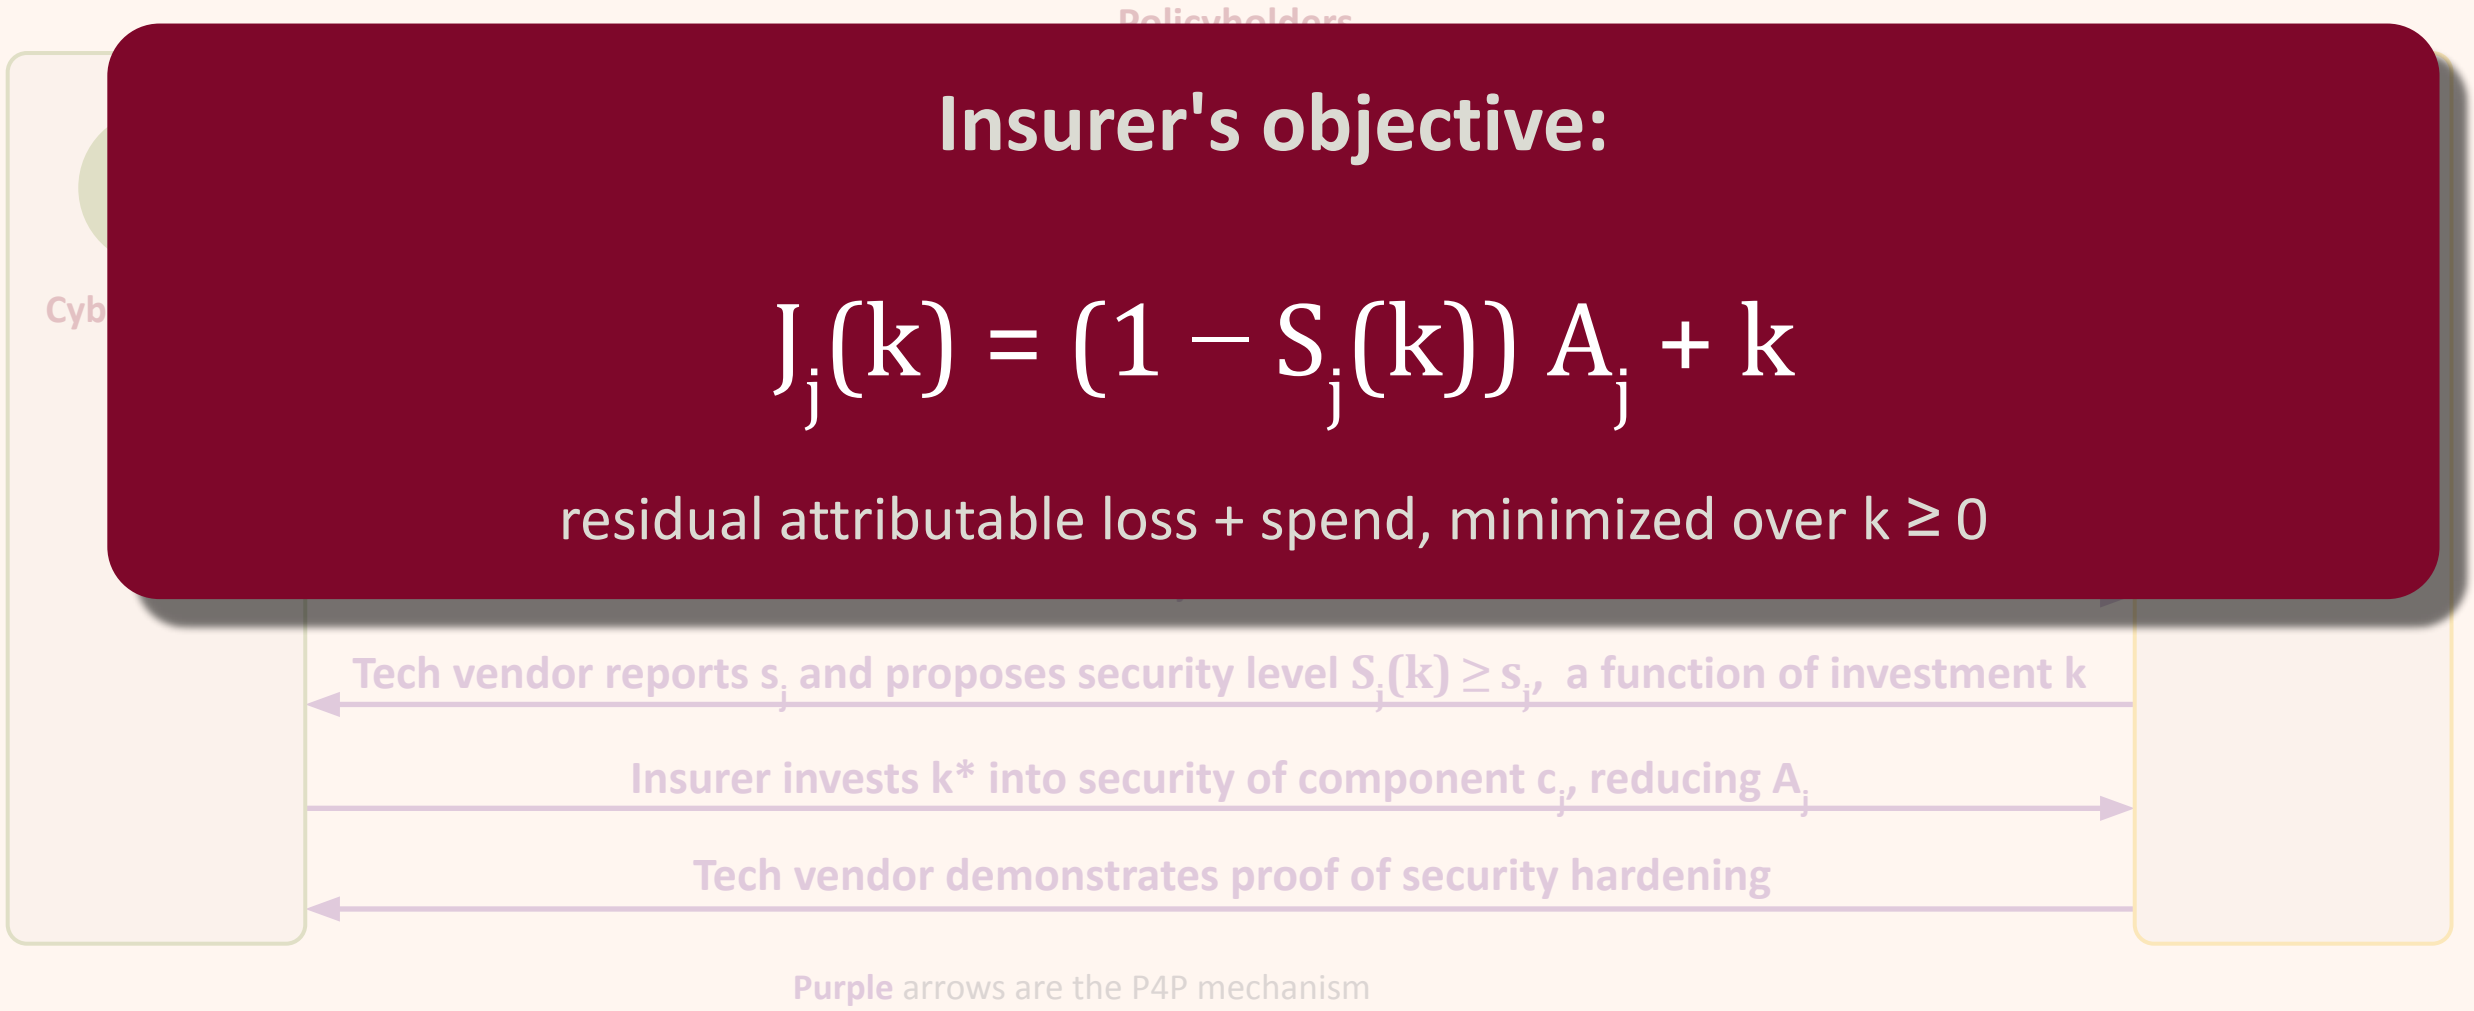
With Cyber P4P, vendors disclose and patch vulns in exchange for payments → reduces info asymmetry

Cyber P4P Model



Purple arrows are the P4P mechanism

Cyber P4P Model



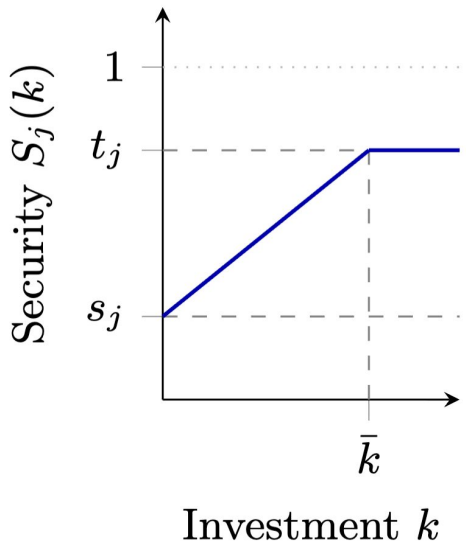
Insurer's objective:

$$J_j(k) = (1 - S_j(k))\, A_j + k$$

residual attributable loss + spend, minimized over $k \geq 0$

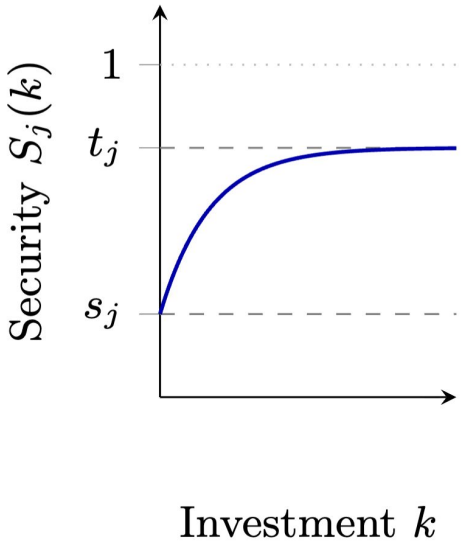
What might a plausible security investment function look like?

Consider an investment, k, to enhance the security of a particular component, $S_j(k)$



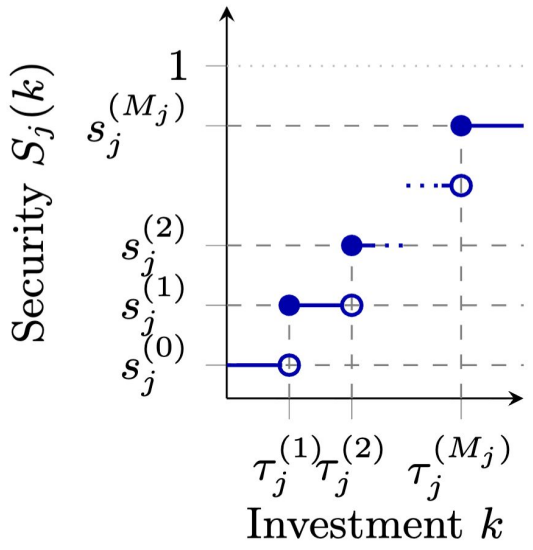
(a) Bounded linear

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



(b) Diminishing returns

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

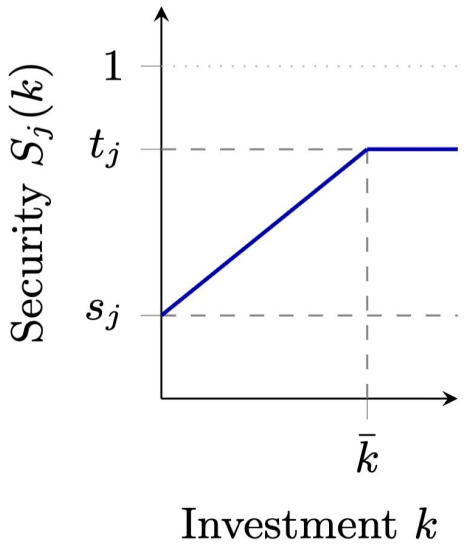


(c) Stepwise

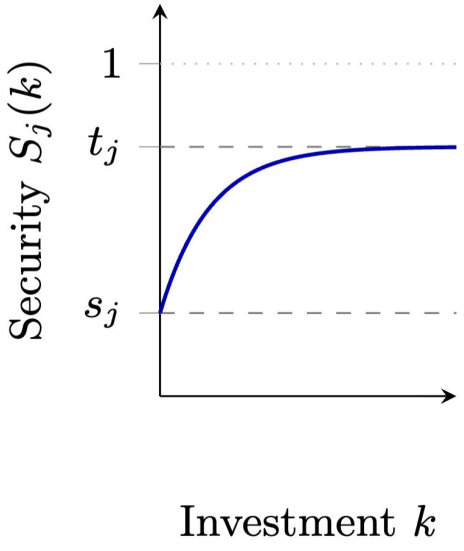
$$S_j(k) = \begin{cases} s_j^{(0)} & \text{if} \quad k < \tau_j^{(1)}, \\ s_j^{(1)} & \text{if} \quad \tau_j^{(1)} \le k < \tau_j^{(2)}, \\ s_j^{(2)} & \text{if} \quad \tau_j^{(2)} \le k < \tau_j^{(3)}, \\ \quad\vdots & \\ s_j^{(M_j)} & \text{if} \quad k \ge \tau_j^{(M_j)} \end{cases}$$

Model Analysis

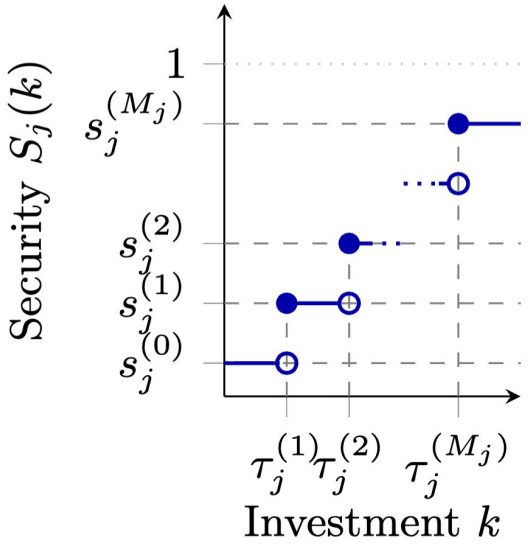
Insurer's optimal single-component investment



(a) Bounded linear



(b) Diminishing returns



(c) Stepwise

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

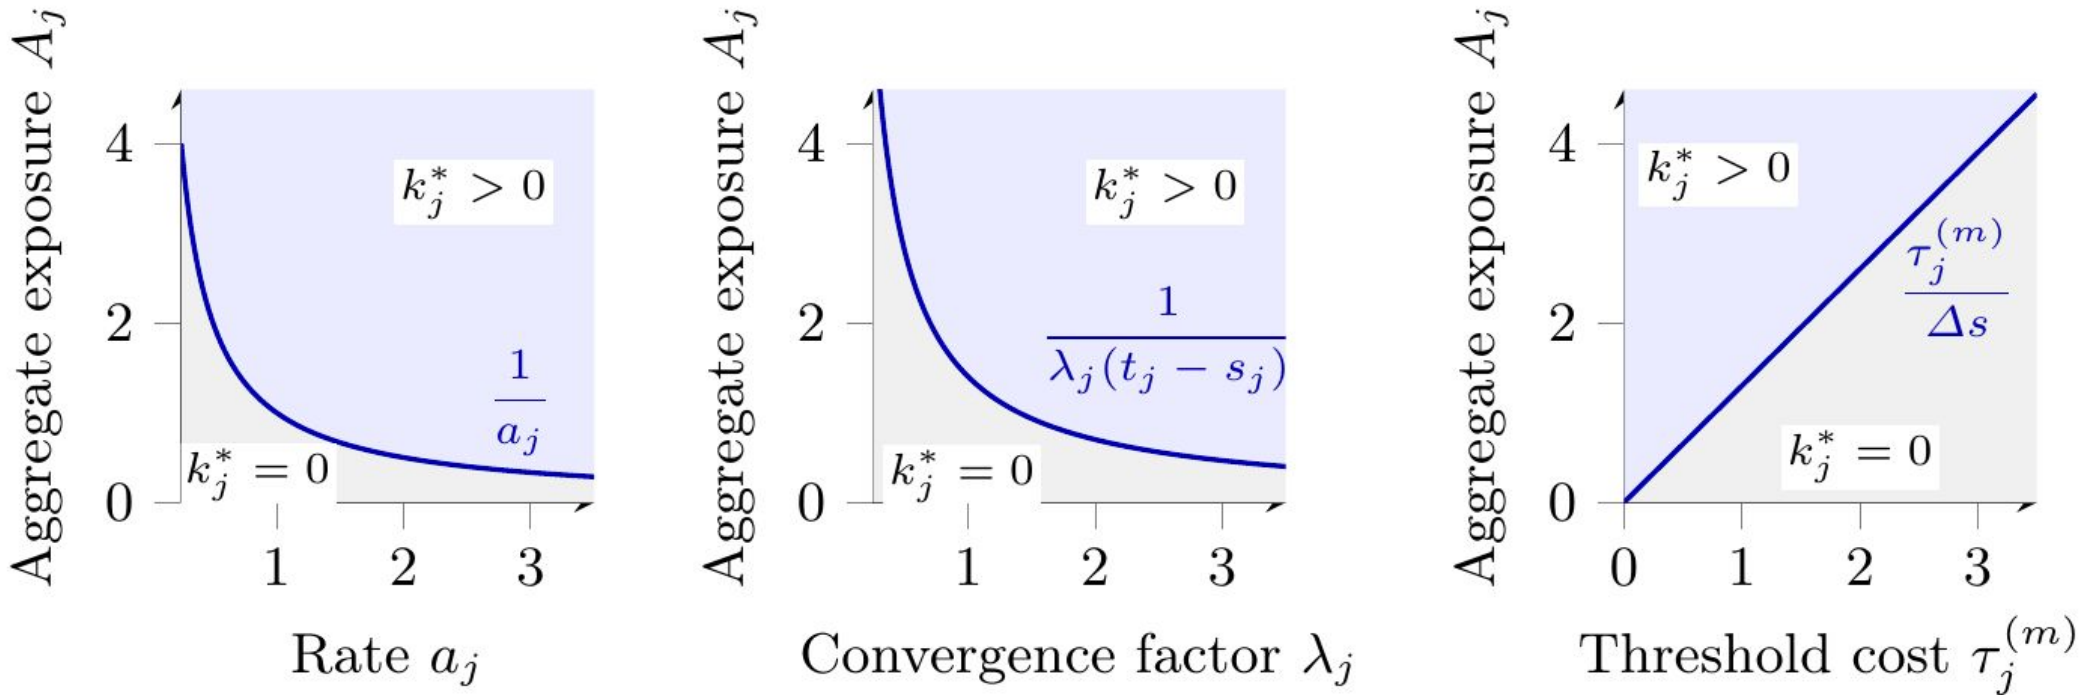
$$k_j^* = \begin{cases} 0, & \lambda_j (t_j - s_j) A_j \le 1, \\ \dfrac{1}{\lambda_j} \ln\big(\lambda_j (t_j - s_j) A_j\big), & \lambda_j (t_j - s_j) A_j > 1. \end{cases}$$

$$k^*_j = \arg\min_k J_j(k)$$

FORDHAM UNIVERSITY

Model Analysis

Invest or Not-Invest decision regions



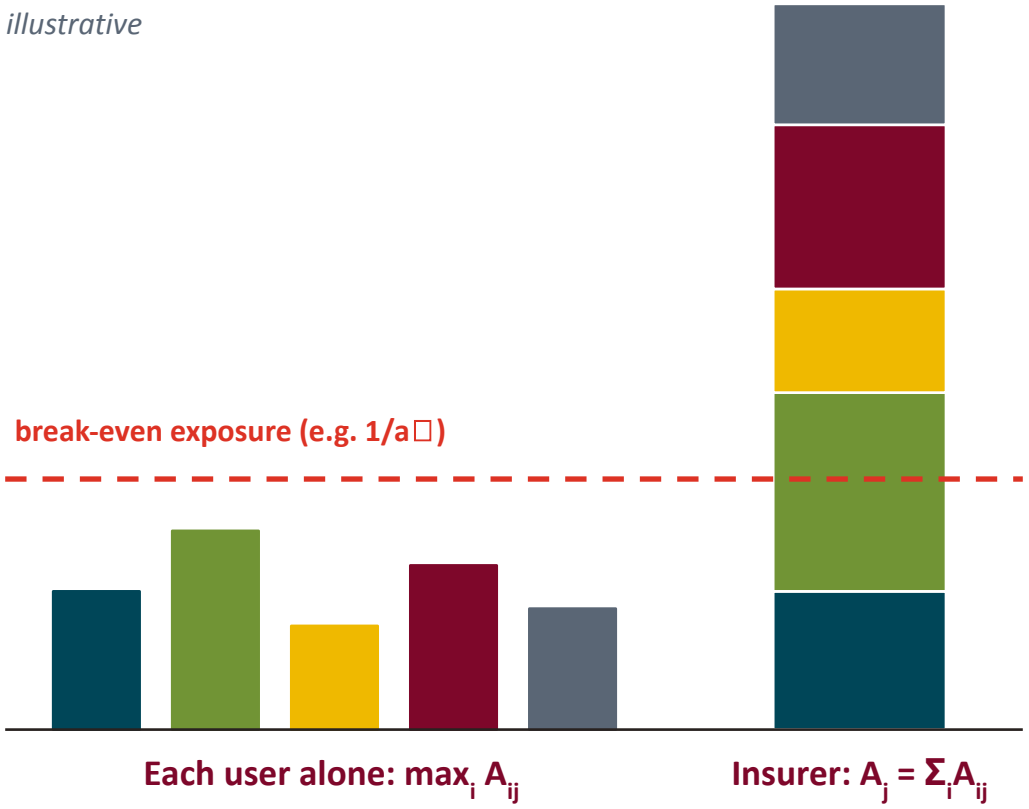
Invest when the return on the first dollar exceeds its cost:

$a_j A_j > 1$ $\lambda_j(t_j - s_j)A_j > 1$ $\Delta s^{(m)} A_j > \tau^{(m)}$

Higher exposure, cheaper hardening, and more headroom all widen the investment region.

Model Analysis

Component security as a public goods game



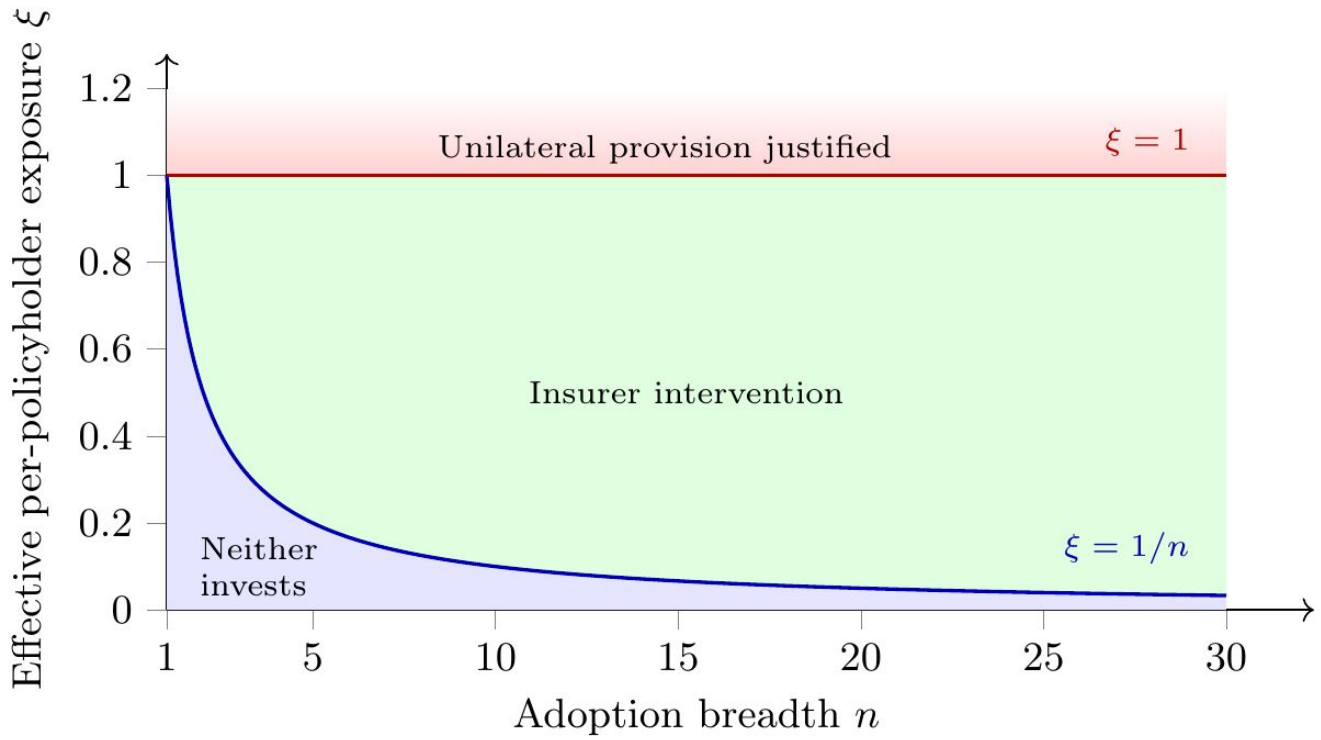
- Component security s_j is non-rivalrous and non-excludable. Hardening c□ helps every user at once.
- Without an insurer, each user weighs only its own stake.
- The most exposed user decides and the rest free-ride.

Proposition 1: an insurer-only region

For the bounded linear and diminishing returns security functions, there is always a non-empty set of parameter values where no individual user would invest yet is cost rational for the insurer.

Model Analysis

A symmetric policyholder benchmark



Symmetric benchmark

n users, each with exposure A_0
Private return $\xi = a_j A_0$

Insurer-only region: $1/n < \xi < 1$

$\xi = 0.20$

$n \geq 6$

insurer invests

$\xi = 0.05$

$n \geq 21$

insurer invests

Model Analysis

A symmetric policyholder benchmark with empirical data

Scenario	Loss ℓ	β_0	Exposure A_0	Private return ξ	Min. users 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**

Insurer intervention can be rational with as few as

5

exposed policyholders, using plausible model inputs

$a_j = 2.5\times10^{-6}$ per dollar: a $100k intervention buys +0.25 security before saturation. ℓ anchored to IBM's 2026 Cost of a Data Breach ($5.0M global average).

Modeling Limitations

- Is the security investment function $S_j(k)$ ever truly knowable?
- Only a single insurer considered (for now)
- Model does not use traditional insurance variables (premiums, deductibles) or consider effects like adverse selection

Known limitations from healthcare P4P

The “gimme” contract

Insurer and hospital agree to a 4% cost increase, booked as:

2.5% rate increase + **1.5%** “bonus”

...pegged to a near-certain target, like a self-reported 80% handwashing rate. Plans look cheaper; the bonus is a disguised payment.

Collusion

Easy targets turn risk-sharing into a back channel

Kingmaker effects

Funding what big clients use, not what helps most

Monoculture

Everyone adopts the preferred stack, correlating losses

Misreporting

Understated baselines, inflated S(·), audit-only fixes

Fraud

Same fix billed to several insurers; misused funds

Moral hazard

Policyholders relax internal controls

Model limits: single insurer · deterministic S(·) · passive adversaries and vendors · no catastrophic losses

Takeaways

1. Cyber insurers are typically seen as passive risk-pricers. Perhaps they should take a more active role in defense.

2. Healthcare P4P offers decades of real-world lessons for how an insurer might incentivize upstream quality/security

3. Figuring out “Cyber P4P”: many open game theory + mechanism design problems

john.chester@fordham.edu · adam.hastings@fordham.edu

FORDHAM UNIVERSITY

Thanks!

Any questions?

John Chester · Adam Hastings
Fordham University · GameSec 2026

FORDHAM UNIVERSITY

