11-2009

Why IT Managers Don't Go for Cyber-Insurance Products

Tridib Bandyopadhyay
Kennesaw State University, tbandyop@kennesaw.edu

Vijay S. Mookerjee
University of Texas at Dallas

Ram C. Rao
University of Texas at Dallas

Recommended Citation
Bandyopadhyay, Tridib, Vijay S. Mookerjee, and Ram C. Bao. "Why IT managers don't go for cyber-insurance products." Communications of the ACM - Scratch Programming for All 52, no. 11 (2009): 68-73.

DOI:10.1145/1592761.1592780

Proposed contracts tend to be overpriced because insurers are unable to anticipate customers' secondary losses.

BY TRIDIB BANDYOPADHYAY, VIJAY S. MOOKERJEE, AND RAM C. RAO

Why IT Managers Don't Go for Cyber-Insurance Products

DESPITE POSITIVE EXPECTATIONS, cyber-insurance products have failed to take center stage in the management of IT security risk. Market inexperience, leading to conservatism in pricing cyber-insurance instruments, is often cited as the primary reason for the limited growth of the cyber-insurance market. In contrast, here we provide a demand-side explanation for why cyber-insurance products have not lived up to their initial expectations. We highlight the presence of information asymmetry between customers and providers, showing how it leads to overpricing cyber-insurance contracts and helps explain why cyber insurance might have failed to deliver its promise as a cornerstone of IT security-management programs.

Technological controls often lag hackers' skills at circumvention. As a result, residual IT security risks cannot be completely eliminated through technological advancement alone. Investment models[9] of information security suggest that residual IT security risks are transferable to a willing party through cyber insurance. Academic research[2] also corroborates the economic value of cyber insurance in managing the cyber risks integral to a firm's operations. Cyber insurance refers to insurance contracts designed to mitigate liability issues, property loss and theft, data damage, loss of income from network outage and computer failures, Web-site defacement, and cyberextortion.[12] Current cyber-insurance products tend to provide three basic types of coverage: liability arising from theft of data; remediation in response to the breach; and legal and regulatory fines and penalties.[1]

The size of the U.S. cyber-insurance market (annual premiums) was expected to reach $2.5 billion by 2005,[11] and insurance giants like AIG and Chubb created numerous cyber-insurance products for managing IT risk. However, IT managers still show little interest in cyber insurance for their risk-management programs; in 2008, the size of cyber-insurance market was estimated at $450 million.[1] The 2006 CSI/FBI computer crime and security survey[8] reported that although firms use cyber insurance more than before, the annual rate of increase is not substantial; respondents indicating utilization of cyber-insurance products increased from 25% to 29% between 2005 and 2006.

Scant attack-loss data, lack of product-market experience, and accounting difficulties are the most commonly cited reasons for the market's slow growth. These factors have led to conservatism by providers that err on the safe side by overpricing their products. However, in a competitive market,

overpricing is generally corrected over time, as risks/uncertainties are better understood. However, even after more than a decade of commercialization, cyber-insurance products remain underutilized. Here, we argue that the demand-side problem with cyber insurance is deeper than the supply-side problem. Moreover, unless the former is addressed, it is unlikely to correct itself naturally over time.

We further highlight the difference between the way a cyber-insurance contract is structured and the way it is used by IT managers, exploring the decisions behind a disclosure and an indemnity claim of a breach. We differentiate the types of breach based on the way they affect firms. We also explain how they might alter the contract-intended claiming behavior of IT managers. When insurers are unaware of such off-contract behavior or choose to not incorporate such behavior in pricing their offerings, information asymmetry prevails in cyber-insurance contracts. The result is an overpriced cyber-insurance contract and less risk being transferred.

Disclosure and Claim of a Realized Breach

With the help of an event study, H. Cavusoglu et al.[6] showed that publicly disclosed IT security breaches reduce breached firms' stock prices, at least in the short term, because breaches convey questionable health of an IT security program to stakeholders, who then downgrade their risk perception of the firm. Elsewhere, K. Campbell et al.[5] showed that investors discriminate against the type of breach in valuing a breach's economic effect. It is not surprising that the CSI/FBI computer crime and security survey[8] found that only a fraction of the realized breaches are publicly disclosed. Firms apparently use discretion in disclosing realized breaches, depending on the requirements of legal compliance, types of breach, professional norms, and accounting materiality.

Suppose there is no regulatory requirement for disclosure. When a firm lacks cyber-insurance coverage, the information flow regarding a realized breach remains strictly internal to the firm (see Figure 1). On the other hand, if the firm has a cyber-insurance contract in place, it is able to claim its losses from a breach, but the claiming process involves additional external organizations. The increased information flow through external firms greatly affects the firm's ability to keep breach information private. Integrating these ideas with insight from H. Cavusoglu et. al.[6] and K. Campbell et. al.,[5] consider the following observations about claiming indemnity from IT security breaches:

The grapevine. Word of an undisclosed breach can reach stakeholders indirectly via interorganizational grapevines and independent analysts;

Stakeholder perception. As a subsequent effect of the breach, a firm may also suffer secondary loss in terms of reduced stakeholder (investors and customers) valuation; and

Managers' decisions. Because breach information might trigger further secondary losses, IT managers' decisions (whether or not to file a claim) depend on the primary and secondary losses, as weighed against the contract's potential indemnity payout.

Breaches and Losses

Because the process of reclamation through cyber-insurance contracts involves compromise, post-breach definitions are pertinent, starting with the breach:

Symptomatic. A breach is symptomatic when a firm is breached through exploitation of firm-specific vulnerabilities (such as hackers in 2005 accessing the T.J. Maxx stores database of customer credit and debit card information, an exploitation of the vul-

Figure 1. Information flow in a cyber-insurance claim process.

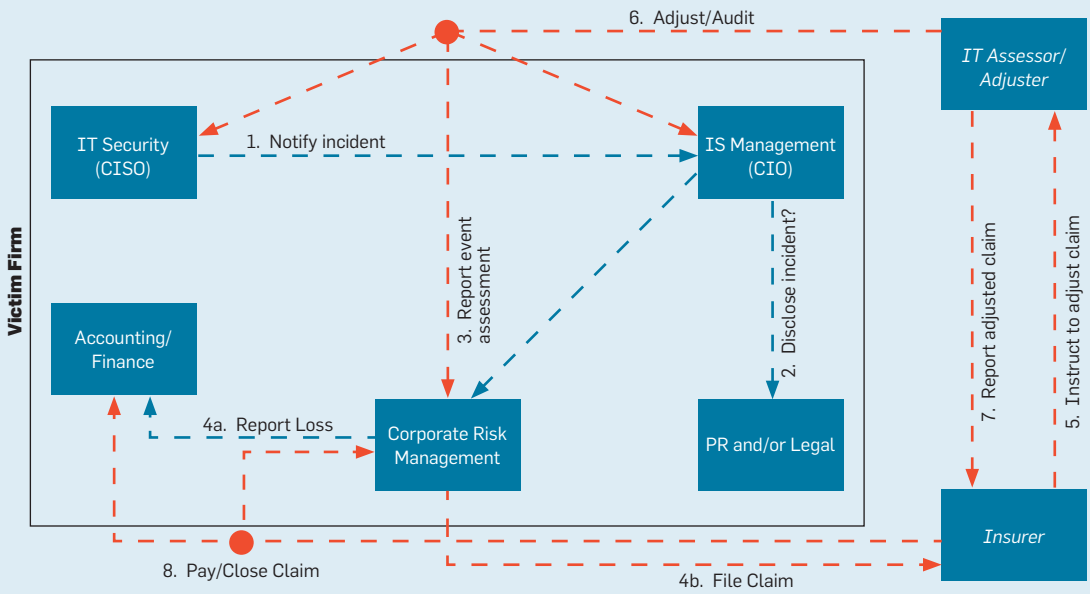

Red lines indicate information flows (entities) external to the victim firm's organizational boundaries.

nerabilities of the stores' data-storage arrangements). Such compromises suggest questionable health of a firm's security program. Consequently, stakeholders downgrade their perception of the firm's IT security;

Systemic. A systemic breach occurs when the affected firm has no reasonable or even known way to defend itself against a new threat vector, especially when the threat is transmitted through the business networks; for example, in January 2004, the MyDoom virus spread primarily via email and severely slowed or shut down email servers with excess traffic (http://reviews.cnet.com/4520-6600_7-5118745-1.html). In this case, stakeholders do not alter their perception of a firm's IT security for systemic breaches, IT security programs plan only for known threats, and firms are all understood to be part of the internetworked global economy where such unknowns are always possible;

Public. A breach is public if it is publicly observed (such as a Web page being defaced), or an observable distributed denial-of-service attack disables a firm's e-commerce transactions or is disclosed through legal requirements or accounting norms (such as the California disclosure requirement for loss of customer data and accounting material loss). By this definition, breaches that are not made public are private.

Here are the potential losses:

Primary. Breaches lead to primary loss (such as direct loss of information or data and operating loss). As an uncontrollable first-degree effect arising from the unuse, disuse, abuse, and misuse of information assets, a primary loss arises under all breach scenarios, or under all combinations of public/private and systemic/symptomatic breaches; and

Secondary. A secondary loss is a second-degree effect, indirectly triggered by information concerning a firm's security inflicting further losses[14] under certain contingent scenarios (see Figure 2). Such losses include indirectly lost or diminished reputation, goodwill, consumer confidence, competitive strength, credit rating, and/or customer churn.

With a cyber-insurance contract in place, the compromised firm claims the primary loss from a public breach, though the secondary loss occurs anyway. The firm also claims its losses for systemic breaches but does not incur secondary loss. For a private symptomatic breach, the secondary loss could occur but only if the firm chooses to file a claim, as in Figure 2. IT managers thus realize that situations could arise where the claiming decision and hence potential indemnity payout must be weighed against the sum of the primary and secondary loss. The secondary losses could be subjectively estimated with the help of extant research outcomes[5,6,14] or assessed/perceived by experiential or other benchmarking processes by the firm's IT managers.

Armed with such foresight, the managers could revise the deductible for the optimal cyber-insurance contract, in turn influencing the amount of insurance that is purchased.

Altered Claiming Strategy

Considering claiming strategy, we begin with a basic insurance contract characterized by cyber-risk-specific circumstances of potential breach and loss scenarios. The firm faces an arbitrarily distributed primary loss x we assume is transparently known to the insurer. Note that this assumption is significant, as it neutralizes all cited and accepted difficulties of cyber insurance in our treatment. That is, our insight is valid irrespective of the correctness of the friction in the cyber-insurance market.

The insurer presents a deductible (d)-based cyber-insurance contract that can be bought for an up-front premium (P), with the promise that primary losses greater than d will be compensated in full by the indemnity payment (I) (see Figure 3). The premium P decreases as the deductible d increases (see Figure 4), a standard observation concerning insurance contracts. In practice, the premium P also includes a market-loading factor that takes care of contract-writing costs, as well as other overhead, including profit margins, if there are any. Figures 3 and 4 together depict the presented contract and the contract-intended claiming behavior of the insured firm.

The prospect firm must optimize and communicate a unique optimal deductible (d^*) to the insurer. The deductible then fixes the premium (P^*) to be paid up front. However, an attempt to arrive at a unique optimal deductible d^* must consider and consolidate several scenarios:

Systemic and Public breach. The firm could submit a claim as per the contract, in case it realizes a systemic breach (no secondary loss) or public breach (the secondary loss occurs automatically); and

Symptomatic breach. Because a symptomatic breach suggests lack of awareness, inadequate technology control, failure to observe policy or procedure, lack of manager oversight, or insider breach, one of the following options is pertinent:

- If the firm discloses the breach, it has a symptomatic public breach for which a decision to submit a claim could follow;
- If the firm decides not to disclose the breach, the decision criterion is further binary:
- It could receive a claim for primary losses from the breach but incur the expected secondary loss or
- It might not claim the primary loss, avoid the secondary loss, or forgo the indemnity payout.
- Assume the act of claiming the symptomatic private breach reveals the breach to the stakeholders with probability p, and that the resultant (adverse) revision of the IT security risk costs[a] the firm y. Thus the expected

Figure 2. Secondary losses from a realized breach.

	Private breach	Public breach
Symptomatic breach	Depends on claiming decision (Probabilistic?)	Yes
Systemic breach	No	No
	Secondary Loss Scenario	

a Here we assume that an IT security breach yields a fixed amount of downward risk revision by stakeholders. We also separately analyzed the case (not included here) in which

secondary loss is py. Clearly, the firm has no reason to claim for losses up to $r = d + py$. By the second binary criterion, all symptomatic private breaches causing primary loss of magnitude between d and r are now likely to be unclaimed (see Figure 5). Note that r, not the contracted deductible d, is the de facto deductible for the symptomatic private breaches. Assuming that a portion of the realized breaches would be symptomatic private breaches, the unique optimized deductible d^* lies somewhere between d and r ($d < r$). This happens for any arbitrary deductible d the firm might choose. The overall optimized deductible d^* the firm must optimally use is always greater than d. More important, whenever a cyber-insurance contract with an arbitrary deductible d is operationalized at d^*, the insured firm stands to lose part of the expected indemnity payout over the contract horizon; see Figure 5 for the location of this unique deductible.

▸ Only when the firm faces no secondary loss or symptomatic private breaches (or both), d and d^* coincide, and the insured firm exhibits contract-intended behavior under all circumstances. Two interesting observations follow when a firm selectively uses the contracted or de facto deductible depending on the type of realized breach:

▸ The higher the secondary loss, the farther apart are d and r, meaning d and d^* are farther apart as well; and

▸ A greater proportion of symptomatic private breaches over the contract horizon increases the relative frequency when the de facto deductible r (not the contract intended d) are used. This proportionally raises the amount of indemnity to be lost in the process, as in Figure 5.

In effect, IT managers looking toward the contract horizon anticipate too little expected indemnity from cyber-insurance products, so the contract appears overpriced. For the same premium, the firm must use a higher overall deductible (see Figure 6).

Figure 7 outlines the complete cyber-insurance utilization scenario. No behavior and underclaiming behavior

downward risk revision is a function of the magnitude of the primary loss. The fundamental insights from each case are the same.

Figure 3. Relationship among loss, deductible, and indemnity in a cyber-insurance contract.

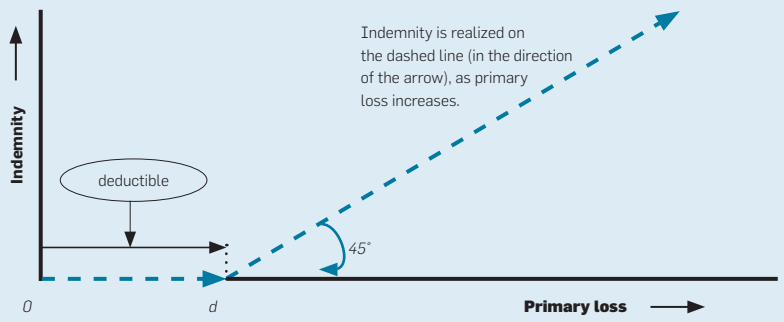

Figure 4. Relationship between premium and deductible in a cyber-insurance contract.

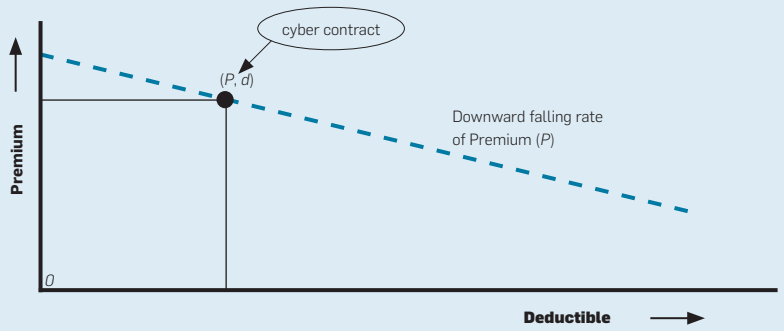

Drawn for uniform primary loss and constant market loading factor

Figure 5. Relationship between de facto deductible *r* and realized indemnity *I*.

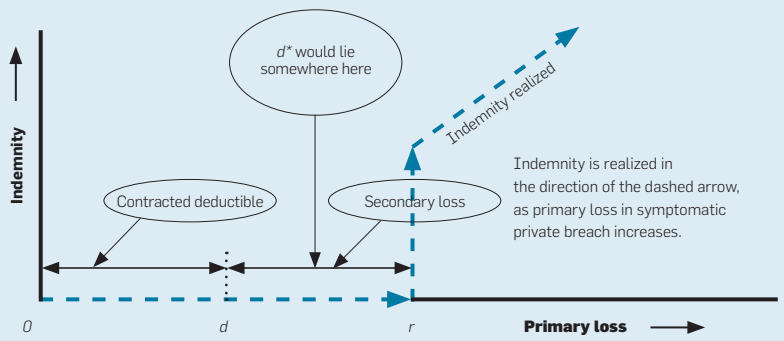

Figure 6. Perceived overpricing of a cyber-insurance contract.

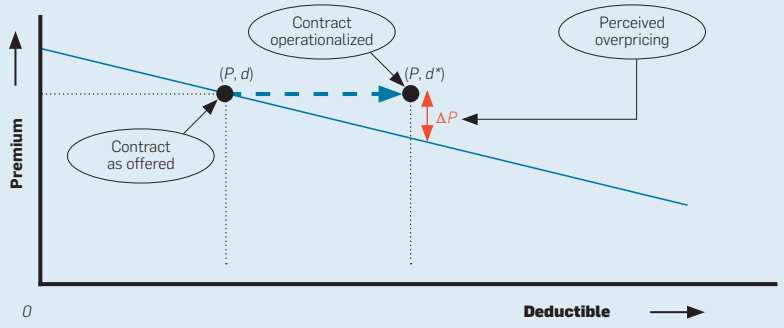

Drawn for uniform primary loss and constant market loading factor

Figure 7. Decisions, costs, and payoffs under contingent scenarios involving cyber insurance.

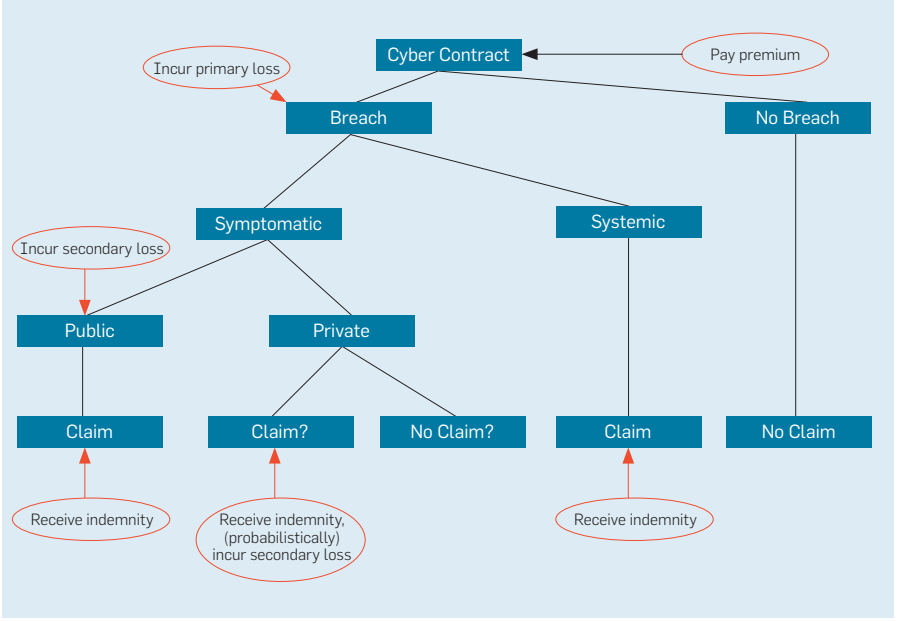

are thus rational for IT managers under certain circumstances; their lack of interest in cyber-insurance products is rational as well. More important, an underclaiming strategy remains off-contract, possibly heralding information asymmetry between insurers and insured firms in the cyber-insurance market.

Information Asymmetry

Figure 8 outlines how the cyber-insurance market could move through possible scenarios of information asymmetry. Initially, the market could begin in naïve symmetry (quadrant I) where neither the insured nor the insurer knows the existence, nature, or magnitude of the secondary loss. As such, a cyber-insurance contract is written with business prudence in light of other established insurance markets. As the insured firm utilizes information assets in its business processes, the value of asset unuse, disuse, abuse, and misuse become clearer. The insured firm realizes there could be attendant secondary losses following direct losses, as stakeholders reassess the firm's post-breach security. The insured firm now internalizes the ex-post definitions of the types of breach discussed earlier, and managers formalize their optimized claiming strategy for symptomatic private breaches also discussed earlier. This differs from the contract-intended behavior, and the market moves from naïve symmetry to information asymmetry (quadrant II).

Under information asymmetry, either the insured firm fails to credibly signal its off-contract behavior or the insurer ignores the signal while structuring the cyber-insurance contract. Either way, the market is in a state of information asymmetry, and the insured firm pays for the ensuing inefficiency.

The cyber-insurance market is, in part, locked in a state of information asymmetry. Only when the insurer considers the fact that the insured firm selectively uses the contracted and de facto deductibles when pricing the contract, does the market move to information symmetry (quadrant III). When the insurer corrects its premium structure this way, the contract is no longer overpriced, and the cyber-insurance product is able to efficiently transfer more IT risk from insured to insurer.

Risk Transfer

Employing the underclaiming strategy for symptomatic private breaches has a profound effect on cyber insurance as an instrument for transferring IT risk. Applicable for only some realized breaches, it reduces the expected indemnity payout for a given level of premium, causing firms to find the instrument overpriced and hence unattractive. Since firms lack a credible way to communicate their off-contract claiming strategy under current contract provisions, they are forced to pay for information asymmetry.

A detailed analysis of our mathematical model suggests that a cyber-insurance contract optimally transfers a lower amount of IT security risk under information asymmetry. It also suggests that further reducing risk transfer depends on the level of secondary loss. It is important to realize that the major consumers of cyber-insurance products are IT-intensive firms that could face relatively high secondary losses.

Unfortunately, IT-intensive firms also likely find the proposed premium structure overpriced in the presence of information asymmetry, as in Figure 5. On the other hand, firms with low IT security exposure may find cyber-insurance products less overpriced in light of their lower secondary loss.

We have shown that in the presence of secondary loss in symptomatic private breaches, the optimal deductible d^* is between d and r, as in Figure 5. Further analysis shows the smaller

Figure 8. Information asymmetry and market transition.

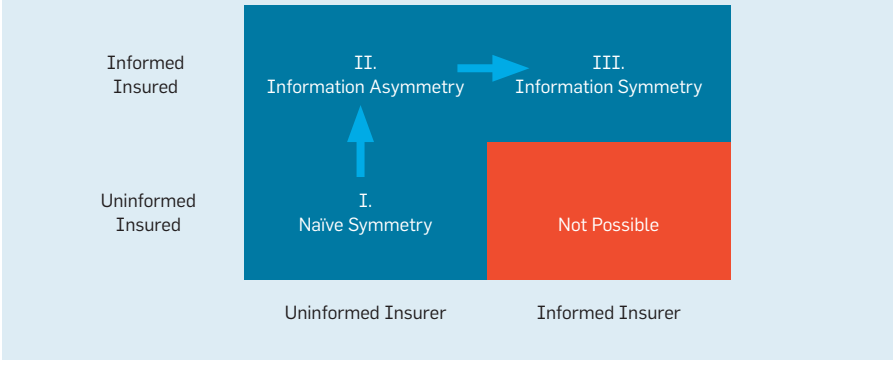

the ratio of secondary loss to the deductible, the lower is the relative overpricing of a cyber-insurance product. Thus, it appears that managers make a rational choice when using cyber-insurance products with high deductible d such that the effect of relative overpricing on their contracts is minimized. IT managers tend to self-insure the smaller losses yet attempt to provide assurance to their stakeholders of low-probability catastrophic breaches.[b]

This analysis suggests that firms with IT-intensive business processes find themselves better off self-insuring a high proportion of their cyber-risk, whereas those with low-intensity IT processes could find cyber-insurance products less pricey under today's market conditions. In light of these outcomes, it becomes apparent why cyber insurance, as a market instrument, has seen little utilization or growth as a financial instrument in managing firms' IT security risk.

Outlook

The cyber-insurance market is characterized by information asymmetry in contracts resulting in the suboptimal transfer of IT risk. From a market perspective, moving to information symmetry (Figure 8, quadrant III) is desirable. Because insured firms pay the price for information asymmetry (quadrant II), a move to information symmetry necessarily increases the utility of the insured firm, with other conditions the same. However, the same may not hold for the insurer. A detailed analysis of the contingency tree (see Figure 7) suggests that under certain conditions (such as significant secondary loss) the insurer is better off under information asymmetry. Under other conditions, the insurer could be better off under information symmetry. This means the insurer would find it beneficial to lower premiums and thus grow the market for cyber insurance.

Firms with a significant amount of IT in their core business processes largely constitute the demand side of the cyber-insurance market. The market is thus relatively homogeneous with respect to (high) secondary loss, and the insurer is better off in a market characterized by information asymmetry. This situation suggests that market mechanisms alone may not produce information symmetry in the cyber-insurance market. Because insured firms likely utilize high levels of deductible in cyber-insurance contracts and do not claim small yet frequent losses, the accumulation of claim data suffers. Lack of claim data may be one reason why after even the past 10 years, cyber insurance is not a major component of corporate IT security initiatives. On the other hand, the relatively small size of the market keeps the costs of writing cyber-insurance contracts high, forcing insurers to impose high margins on individual contracts. Unless it expands, insurers cannot gain more experience or accumulate significant actuarial data and feel no pressing motivation to move to information symmetry. This could mean the market stays locked in information asymmetry.

The structural problem with the market can be resolved if secondary loss were included in contracts. Exotic bundled contracts (individual contracts for primary and secondary losses designed in tandem and bundled together) could be a viable solution. It might take care of the fact that the primary (secondary) losses are determined before (after) the breach, so IT managers are able to take independent decisions concerning disclosures and claims. Even so, valuing secondary loss is more challenging than valuing primary loss, so there appears no easy solution, even if bundled contracts are written.

It is possible that along with increased regulatory compliance and oversight, the relative proportion of private breaches decreases, along with the information asymmetry between insurer and insured. Similarly, separating contracts on the basis of disclosure (compliance or discretionary) might also be a move in a positive direction. However, contracts offered by major insurers today are either architecture-oriented (such as a network breach), asset-based (such as a data breach), attack-specific (such as viruses and worms), or liability-focused. No offered contract considers secondary loss or accommodates the complexities of a firm's decision to file a claim in the face of secondary loss. It appears that without significant changes in the design of the contracts, there is little hope for the continued growth of the overall cyber-insurance market. C

b That firms buy cyber insurance with high deductibles was also pointed out by the IT director of a Dallas firm during a discussion with us at the University of Texas, Dallas. He explained that firms often buy cyber insurance to allay investors' fear of major losses from IT security breaches yet depend on the policies, procedures, and technical controls of IT security to manage more frequent but smaller losses.

References

1. Betterley Report. *Cyberrisk Market Survey 2008*; http://www.betterley.com.
2. Bohme, R. Cyber insurance revisited. In *Proceedings of the Workshop on the Economics of Information Security* (Boston, MA, June 2–3, 2005); http://infosecon.net/workshop/index.php.
3. Borch, K. *Economics of Insurance. Advanced Textbooks in Economics 29*, K.K. Aase and A. Sandmo, Eds. North Holland, Amsterdam, 1990.
4. Borch, K. *The Mathematical Theory of Insurance.* Lexington Books, Lexington, MA, 1974.
5. Campbell, K., Gordon, L.A., Loeb, M.P., and Zhou L. The economic cost of publicly announced information security breaches: Empirical evidence from the stock market. *Journal of Computer Security 11*, 3 (2003), 431–438.
6. Cavusoglu, H., Mishra, B., and Raghunathan, S. The effect of a security breach announcement on market value: Capital market reactions for breached firms and Internet security developers. *International Journal of Electronic Commerce 9*, 1 (Fall 2004), 69–104.
7. Gollier, C. and Pratt, J.W. Risk vulnerability and the tempering effect of background risk. *Econometrica 64*, 5 (Sept. 1996), 1109-1123.
8. Gordon, L.A., Loeb, M.P., Lucyshyn, W., and Richardson, R. *The 11th Annual CSI/FBI Computer Crime and Security Survey* (2006); http://i.cmpnet.com/gocsi/db_area/pdfs/fbi/FBI2006.pdf.
9. Gordon, L.A., Loeb, P.M., and Sohail, T. A framework for using insurance for cyber-risk management. *Commun. ACM 46*, 3 (Mar. 2003), 81–85.
10. Kesan, P.J., Majuca, R.P., and Yurcik, W.J. *The Economic Case for Cyber Insurance. Securing Privacy in the Internet Age.* Stanford University Press, Palo Alto, CA, 2005.
11. Majuca, R.P., Yurcik, W., and Kesan, J.P. *The Evolution of Cyber Insurance* (2006); http://arxiv.org/ftp/cs/papers/0601/0601020.pdf.
12. Marsh, Inc. E-Commerce E-Business; http://global.marsh.com/risk/ecommerce/.
13. Mossin, J. and Smith, T. Aspects of rational insurance purchasing. *Journal of Political Economy 76*, 4 (July/Aug. 1968), 553–568.
14. Ponemon Institute. *The Fourth Annual U.S. Cost of Data Breach Study* (2008); http://www.ponemon.org.
15. Raviv, A. The design of an optimal Insurance policy. *American Economic Review 69*, 1 (Mar. 1979), 84–96.
16. Schlesinger, H. The optimal level of deductibility in insurance contracts. *Journal of Risk and Insurance 48*, 3 (Sept. 1981), 465–481.

Tridib Bandyopadhyay (tbandyop@kennesaw.edu) is an assistant professor in the Department of Computer Science and Information Systems at Kennesaw State University, Kennesaw, GA.

Vijay S. Mookerjee (vijaym@utdallas.edu) is the Charles and Nancy Davidson Distinguished Professor of Information Systems and Operations Management in the School of Management at the University of Texas at Dallas.

Ram C. Rao (rrao@utdallas.edu) is the Founders Professor and a professor of marketing in the School of Management at the University of Texas at Dallas.