A Comparison of Market Approaches to Software Vulnerability Disclosure

Rainer Böhme

Technische Universität Dresden
Institute for System Architecture
01062 Dresden, Germany

rainer.boehme@tu-dresden.de

Abstract. Practical computer (in)security is largely driven by the existence of and knowledge about vulnerabilities, which can be exploited to breach security mechanisms. Although the discussion on details of responsible vulnerability disclosure is controversial, there is a sort of consensus that better information sharing is socially beneficial. In the recent years we observe the emerging of "vulnerability markets" as means to stimulate exchange of information. However, this term subsumes a broad range of different concepts, which are prone to confusion. This paper provides a first attempt to structure the field by (1) proposing a terminology for distinct concepts and (2) defining criteria to allow for a better comparability between different approaches. An application of this framework on four market types shows notable differences between the approaches.

1 Introduction

Vulnerabilities are errors in computer systems which can be exploited to breach security mechanisms. They typically emerge during software development and some remain undiscovered in the final product, largely because common software testing methods are not designed to detect errors that require strategic interaction of a malicious party. However, widely deployed software is subject to public scrutiny that leads to the discovery of vulnerabilities. Information about new (i.e., recently discovered) vulnerabilities is highly valuable as it decides about the success of attack or defense in open computer networks: *malicious users* may use the information to launch attacks on vulnerable systems, whereas *honest users* have an interest to assess the security risks they are exposed to and to decide about appropriate countermeasures, such as demanding a patch from the vendor or switching to a competitor's product. Hence, as long as perfectly secure software is not available, the optimal distribution of vulnerability information is an important factor for the stability of a "network society" [1–4].

The distribution of vulnerability information, however, is rarely a technical problem but rather a result of rational decision-making of the parties involved: Why should a teenage computer freak report the outcome of his leisure-time efforts to the public if he can increase his pocket money by selling crucial information on the black market? Why should a software vendor invest time and

Final version for proceedings of ETRICS (March 19, 2006). To appear in LNCS.

money in secure programming, when his competitor does not, and his customers cannot measure the difference in quality? These questions motivate to regard computer security from the point of view of economics, a discipline studying rational decision-making of independent agents. A good introduction to the field of economics and information security can be found in Ross Anderson's seminal article [5].

The interest in "vulnerability markets" can be partly attributed to theoretical work in this interdisciplinary community. In addition, recent developments, such as vulnerabilities being offered on online auctions and security firms allotting rewards for vulnerability reports, contribute to the public attention. However, sometimes completely different concepts are referred to as "vulnerability markets", which is a source for confusion. Therefore, this paper aims to structure the area by presenting a typology of vulnerability markets. Moreover, a criteria-based framework for the comparison of different market types is proposed.

The remainder of this paper is structured as follows: Section 2 briefly summarises economic reasons for the deficit in nowadays computer security and explains how vulnerability markets could change this situation to the better. Section 3 presents a typology of vulnerability markets in the literature and discusses their similarities and differences. In Section 4, a set of criteria based on the anticipated positive effects (see Section 2) is defined and then applied for a systematic comparison of different market types. Section 5 concludes the paper with pointers to existing limitations and possible future research.

2 The Computer Security Market Failure

Before discussing the effects of vulnerability markets, we sketch two examples illustrating how the market currently fails in providing computer security.

The first example refers to the supply-side for security technology. Its theoretical background is George Akerlof's lemon market problem [6]. Akerlof studied the rules of a market with asymmetrical information between buyer and seller. For instance, the typical buyer of a second hand car cannot distinguish between good offers and bad ones (so-called "lemons"), because—unlike the seller—he does not know the true history of the car. So the buyer is not willing to pay more than the price of a lemon. As a result, used cars in good condition will be under-provided on the market. The same applies to computer security: security is not visible and thus becomes a trust good. Since the buyer is unable to differentiate secure from insecure products apart, the market price drops to the level for insecure products. Hence, vendors have little incentive to develop sound security technology and some might rather prefer to invest in more visible features, or to be first on the market to dominate the technological standard [7].

The second example targets to the demand-side of security. Its theoretical roots lie in the popular "tragedy of the commons", another economic theory published by Garrett Hardin [8]. Consider a computer network and the threat of botnets [9], where security is rather a property of the network than of its individual nodes: if the weakest node gets corrupted then the other nodes face a high risk of being attacked and consequently face higher expected loss. Therefore, the cost of security incidents is distributed among all nodes. On the other hand, if one node decides to invest in security, then all computers in the network benefit, because the now secure node is less likely to cause harm to others from forwarded malicious traffic. In brief, since both risk and benefits are socialised between all nodes, individuals lack the incentive to unilaterally invest in security. They prefer to remain "free riders" waiting for others to pay in their place (who'll never do so, because of the same rationale; see [10] for a rigorous analysis).

To sum it all up, the lemon market suggests that vendors under-provide security to the market, whereas the tragedy of the commons can explain why users demand less security than appropriate. A common notion for this deadlock is *market failure.*

The collection of reasons for the market failure is by far incomplete[1] but it is enough to characterise the problem and to derive objectives to mitigate it. To counter the lemmon effect, security has to become measurable [13]. The free rider problem can be solved by redistributing the costs in a way that nodes are made responsible to bear all costs and receive full utility of their own decisions. In micro-economic terms this corresponds to an "internalisation of externalities"; or, as we might frankly say, tax bad security [14].

There are two ways to fix a market failure. At first, regulation—which is least desirable as there are numerous examples where regulation renders the situation even worse. Indeed, good regulation is really difficult since it often implies a trusted third party (TTP) as "social planner", whom to make incorruptible is costly, if not impossible. There exists a large body of literature on public choice theory, which studies imperfections due to state interventions and adverse incentives in government decision-making [15, 16]. Note that we hesitate to argue that regulation of computer security is generally a bad idea or inferior to market approaches. We rather consider it as an option which needs to be studied, though it is beyond the scope of this paper.

The second possible response to a market failure is establishing new markets with mechanisms that eventually feedback and thus mitigate the problems at their source. If the markets are designed properly, then market prices serve as valid indicators for underlying security properties and thus make security measurable. Moreover, markets can well differentiate between good and bad security. For instance, cyber-insurance contracts could contain deductions for customers where good security technology and practices are in place. Conversely, users who do not invest appropriately in security pay a higher premium, which corresponds to the objective of taxing bad security.

This is the theoretical justification for vulnerability markets. In the following section we present concrete concepts for vulnerability markets before we discuss how suitable each concept is to counter the market failure.

[1] Another often-cited topic is the discussion on software liability [11, 12], which we omit for the sake of brevity.

Table 1. Alternative names for vulnerability markets in the literature

Proposed term	Equivalents in the literature
Bug challenges	*vulnerability markets* in [13]
	bug auctions in [17, 18]
	bug bounties on some blogs
Vulnerability brokers	*vulnerability markets* in [19]
	also *vulnerability sharing circles*
Exploit derivatives	related to *security tokens* in [20], but not the same
	prediction markets in more general contexts
Cyber-insurance	(not ambiguous)

3 Classifying Vulnerability Markets

This section contains a typology of possible market concepts for security-related information. Note that our terminology is deliberately not consistent with all prior art, because some terms have been used ambiguously in the past. Therefore, we collected alternative names for each concept together with the corresponding references in Table 1.

3.1 Bug challenges

Bug challenges are the oldest concept to "prove" the security strength of a product, or to guarantee invisible properties of traded goods in general. In the simplest scenario, the vendor allots a monetary reward for vulnerability reports related to his product. Then the amount of the reward is a lower bound to the security strength of the product: it can be safely used to handle and secure assets totalling up to this amount because a rational adversary would prefer to report possible vulnerabilities and cash the reward over attacking the system and capitalising the information gained. Stuart Schechter coined the term *market price of vulnerability* (MPV) for a metric derived from this model [13]. Examples for simple bug challenges in the real world include the Mozilla Security Bug Bounty Program[2], the RSA factoring contests, and the Argus Security Challenges[3].

One of the main issues in bug challenges is the difficulty to find an appropriate level of reward. Therefore, several extensions to fixed-sum bug challenges have been proposed in the literature. For example, the reward could be initialised at a very low level and then gradually grow over time. The most widely known

[2] `http://www.mozilla.org/security/bug-bounty.html`

[3] `http://www.wired.com/news/technology/0,1282,43234,00.html`; its aftermath demonstrates the need for a trusted third party to settle the deals: `http://www.net-security.org/news.php?id=1522`.

example of this type is Donald E. Knuth's reward of initially 1.28 USD for each bug in his TEX typesetting system. His reward grows exponentially with the number of years the program is in use. To limit the expenses, the vulnerability buyer may decide to reset the reward after each vulnerability report [13].

This scheme allows for a certain dynamic in price-setting, which is similar to market mechanisms designed as auctions [17]. This is the reason why bug challenges are sometimes referred to as "bug auctions", which *should not be mistaken* as offering vulnerability reports on auction platforms such as eBay.[4] For a precise terminology, we propose to distinguish between *buyer-administered* bug auctions and *seller-administered* bug auctions.

Even with this extension, the price quote is not always a reliable indicator for the true security of a product. Consider the case where two vulnerabilities are discovered at the same time. A rational agent would sell the first one and then wait with the second release until the reward has climbed back to a worthwhile amount. In the meantime, the mechanism fails completely in aggregating information about the security of a product, and prudent users should stop using it until the reward signals again a desirable level of security.

As to the operational aspects, it is still questionable whether the rewards can ever be high enough to secure the accumulated assets at risk for software with large installation bases in critical environments, such as finance, health care, or governmental use. Even when taking into account that the actual amount can be smaller than the assets at stake by assuming a risk-averse adversary (the reward is certain whereas making a fortune as black-hat is risky), the so-reduced sum still requires a financial commitment of vulnerability-buyers which exceeds the tangible assets of many software vendors, let alone the case of open source software or depreciated systems, where the vendor ceased to exists.

3.2 Vulnerability brokers

Vulnerability brokers are often referred to as "vulnerability sharing circles". These clubs are built around independent organisations, mostly private companies, who offer money for new vulnerability reports. They then circulate the acquired information within a closed group of subscribers to their security alert service. The customer bases are said to consist of both vendors, who thus learn about bugs to fix, and corporate users, who want to protect their systems even before a patch becomes available. In the standard model, only honest users are assumed to join the club, though it might be very difficult to enforce this policy in practice.[5] With annual subscription fees of more than ten times the reward for a vulnerability report, the business model seems so profitable that there are multiple players in the market: iDefense was first with its "Vulnerability Contributor Program", TippingPoint/3COM followed with a "Zero-day Initiative"

[4] `http://it.slashdot.org/article.pl?sid=05/12/12/1215220`

[5] It is remarkable that Nizovtsev and Thursby in [3] model the proportion of 'black hats' within vulnerability sharing circles equal to the proportion in the population. They justify this decision with frequent reports of insider attacks.

and Digital Armaments also offers money or barter deals for vulnerability information. This kind of competition increased the (publicly communicated) reward sums to 4-digit dollar amounts per bug and led to sophisticated bonus schemes, which resemble customer loyalty plans. This business model has been criticised as blackmail, because vendors and users are forced to subscribe to all services in order to avoid missing important information, even when the frequency of actually relevant reports is very low.

A technically similar but socially more acceptable service is offered by CERT (Computer Emergency Response Team). It also acts as a vulnerability broker, albeit on a non-profit basis. It does not pay any reward for reporting vulnerability information and disseminates that information for free. A recent paper compares the social welfare of vulnerability markets (more precisely: commercial vulnerability brokers) with the CERT approach [19]. The authors conclude that a single CERT acting as a social planner always performs better than commercial brokers.[6] Being exposed to competition with commercial brokers, however, the authors suggest for a CERT-type model to offer monetary rewards as well. This, in turn, means that it must be subsidised from public money (which reduced overall welfare) and it remains unclear how to assure that the social planner works efficiently and turns away from hidden action.

3.3 Exploit derivatives

Exploit derivatives apply the idea of binary options, as known in the theory of financial markets, to computer security events. Instead of trading sensitive vulnerability information directly—with all its negative consequences from trading information goods—, a market is constructed for contracts with pay-out functions *derived* from security events [18].

Consider a pair of contracts $(C, \bar{C})$, where C pays a fixed amount of money, say 100 EUR, if there exists a remote root exploit against some specified server software X on platform Y at date D in the future. The inverse contract, $\bar{C}$ pays out the same face value if there is *no* remote root exploit submitted to a market authority—not a trusted third party in a strict sense—before date D. It is evident that the value of the bundle $(C, \bar{C})$ is 100 EUR at any time and that selling and buying it is risk-free.[7] Therefore, one or many market makers can issue as many bundles as demanded by the market participants. Now assume that there is an exchange platform, where the contracts C and $\bar{C}$ can be traded individually at prices determined by matching bid and ask orders. Then the ratio of the market price of C and its face value approximately indicates the probability of software X being compromised before date D.

The accuracy of the price information depends on the liquidity of the market, hence for accuracy we need a high number of participants and low transaction

[6] Note that the authors come from Carnegie Mellon University, which hosts the headquarters of CERT/CC.

[7] Ignoring interest rate yield of alternative investment, which can be easily compensated for, but is omitted here for the sake of brevity.

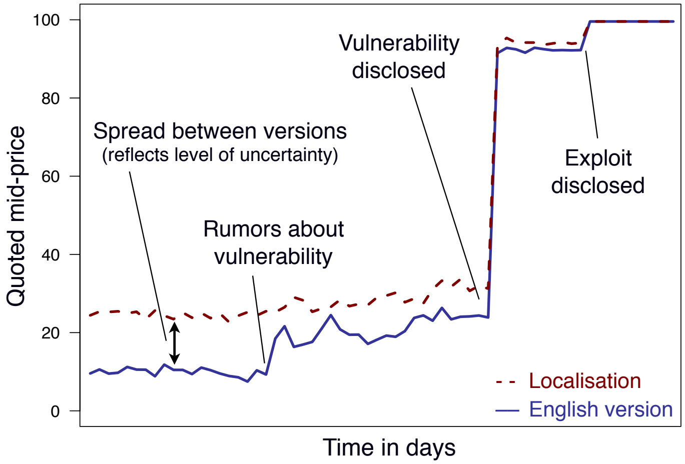

Fig. 1. Relation of events and price quotes of hypothetical exploit derivatives

costs. This market type, however, has the potential to attract far more groups of participants than bug challenges or vulnerability brokers. Software users would demand contracts C in order to hedge the risks they are exposed to due to their computer systems in place. The same applies for cyber-insurance companies underwriting their customers' cyber-risks. Conversely, investors would buy contracts $\bar{C}$ to diversify their portfolios. Software vendors could demand both types of contracts: contracts $\bar{C}$ that pay if their software remains secure as a means to signal to their customers that they trust their own system; or contracts C_{comp} that pay if their competitors' software gets compromised. One could even think of software vendors using exploit derivatives as part of their compensation schemes to give developers an incentive to secure programming.

Finally, security experts (a.k.a. “vulnerability hunters”) could use the market to capitalise efforts in security analyses. If, after a code review, they consider a software as secure, they could buy contracts $\bar{C}$ at a higher rate than the market price. Otherwise they buy contracts C and afterwards follow their preferred vulnerability disclosure strategy. As interaction on the market influences the price, the quotes are constantly updated and can be used as reliable indicators for security strength. Note that this concept does not require the co-operation of the vendor, and the number of different contracts referring to different pieces of software, versions, localisations, etc., is solely limited by demand.

Figure 1 displays a hypothetical price development for an exploit derivative over time. The price quotes reflect changes in the expected security level of the underlying software. Combining information from more than one contract allows for even more interesting metrics. Differences between related contracts (“spreads” in financial terms) can be directly attributed to variations in security or public scrutiny between well-defined technical differences. In the figure, this is illustrated as differences in the perceived security between two localisations of the same software. In addition, joint probabilities of failure can be computed from pairs of contracts to measure the total security of cascaded protection mechanisms.

Like other market types, exploit derivatives require a trusted third party to test candidate exploits at the end of each contract and announce the result. However, if the TTP is required to publish the exploit candidates together with the announcement, it becomes verifiable and cannot cheat. The job can also be distributed to a number of TTPs. Therefore, the assumptions about the TTP are much gentler in this scenario than in other market types.

The concept of exploit derivatives is a modification of seminal work by Kanta Matsuura [20]. He studied the use of option pricing models to assess the risk of cryptographically secured digital objects being compromised. Exploit derivatives, unlike Matsuura's *security tokens*, start with modelling the risk of *components* or mechanisms being compromised rather assessing the risk of loosing the value of the *content* processed in the system. This generalisation, however, does not limit the range of applications at all, because given a set of critical components it is possible to choose a portfolio that exactly matches the risk profile of the defined system. The total value of its content can be matched in a second step by a linear adjustment of the investment volume up to the desired level. In this framework one can even think of cyber-insurance companies being merely intermediaries to whom users and firms outsource their exploit derivatives portfolio management. This guides us to the remaining type of vulnerability markets.

3.4 Cyber-insurance

Cyber-insurance is among the oldest proposals for market mechanisms to overcome the security market failure (see [21–24, 11]). The idea that cures the market failure goes as follows: end users demand insurance against financial losses from information security breaches and insurance companies sell this kind of coverage after a security audit. The premium is assumed to be adjusted by the individual risk, which depends on the IT systems in use and the security practices in place. Therefore, it would be costly to buy insurance coverage for insecure software. This gives users an incentive to invest in security technology. One would even raise the willingness to pay more for secure products if—in the long run—the total cost of ownership including insurance premiums is below the expenses for a less secure product.

In theory, on a long-term average the premiums converge to the actual security risk (plus a constant overhead) because competition sets an upper and profitability a lower bound. Premiums are never completely ill-aligned (like in bug challenges after a reset of the reward). In contrast to bug challenges and exploit derivatives, the premiums are adjusted to each individual insured's risk profile and not on the expected security strength of standard components. This tailored nature is advantageous for the application as a metric, because an organisation or a system is measured on the whole and there is no need for sophisticated and error-prone aggregation to high-level indicators.

However, despite the presence of potent insurance companies, cyber-insurance business remains on a comparatively low volume. One of the reasons could be that insurance companies hesitate in underwriting cyber-risks, because the losses from information security risks are highly correlated globally—think about viruses and worms, and the lack of diversity in installed platforms. This concentration of risk is contrary to the insurance principle of portfolio balancing and requires additional safety premiums that render cyber-insurance policies economically uninteresting [25]. Apart from the fear of "cyber-hurricanes", there are other operational obstacles, such as the difficulty to substantiate claims, the intangible nature of cyber-assets, and unclear legal grounds.

4 Comparison of Market Types

The typology presented in the previous section demonstrates that there is not one "vulnerability market" but rather a family of different concepts. It also becomes evident that the different mechanisms are hardly comparable per se. Nevertheless we try to tackle the research question which market type serves best to counter the security market failure by defining a set of criteria that allow for a more objective comparison. For an ideal vulnerability market, with respect to its ability to counter the security market failure, we have identified three functions, which are elaborated in detail below.

4.1 Information function

The information function refers to the possibility to use market prices as forward-looking indicators for security properties. This function is important to counter the lemon effect because it makes security measurable. It can be divided in sub-dimensions, such as the accuracy of price information, its timeliness and availability to the public.

Some empirical studies show that even existing stock markets do accumulate information related to security events [26–29]. However, stock markets aggregate a large set of different information so that only very extreme security events can be identified in the stochastic movements of market prices. Consequently, an ideal vulnerability market should isolate security-relevant information from other sources of noise and distortion.

4.2 Incentive function

The incentive function addresses the monetary compensation for security research and development. It motivates firms and individuals to participate in the exchange of vulnerability information. Possible incentives from vulnerability markets include incentives for individual bug hunters as well as incentives for developers.

In the absence of operable vulnerability markets, individual bug hunters are motivated by altruism and the prospect of reputation, and—perhaps—by monetary compensation on the black market. Vulnerability markets add monetary rewards on top of the gain in reputation and, depending on the price level, may convince bug hunters to turn away from selling on the black market.

The sole motivation for software developers to invest in security is trust of satisfied customers, which can be capitalised in the long run only. Vulnerability markets add short-term profits and competitive advantage on top of the long-term benefits. With hyperbolic discounting of (uncertain) future revenues [30] and a general tendency to short-term oriented management decisions, vulnerability markets thus add a strong incentive to give security a higher priority.

4.3 Risk-balancing function

The risk balancing function means that the vulnerability market provides instruments to hedge against large information security risks. This is important to mitigate the financial impact of (occasional) security breaches, which may help firms to survive attacks rather than filing for bankruptcy with all its adverse social and economic consequences.

It is also the risk balancing function which contributes to the objective of taxing bad security, both directly and indirectly. The direct effect comes from the fact that instruments covering extreme events are less costly if the extreme events become less likely. The probability of failure, in turn, is related to the level of security (in terms of resistance against attacks) and exposure (in terms of likelihood of being targeted by an attack). As exposure is said to depend largely on how widely a system is deployed, diversity gets rewarded as well. Since diversity is a desirable security property on an aggregated level [31, 32, 25], the risk balancing function taxes bad security also indirectly.

4.4 Market efficiency

Orthogonal to the functions, market efficiency is a criterion which expresses the absence of additional burden in realising the functions. Therefore, efficiency should be regarded as a property of the market, which subsumes the following aspects:

- low transaction costs (it is inexpensive to participate in the market)
- liquidity (high number of participants and possible trade counterparts)
- accountability (low counterparty risk)
- transparency (fair rules, public price quotes)

Not all of these properties are necessary to make vulnerability markets operable, but any of them increases the potential of a vulnerability market to actually counter the security market failure. There exist also a number of dependencies between these sub-dimensions. For example, low transaction costs allow more people to participate in the market and thus automatically improve the liquidity; accountability reduces the transaction costs because the average loss due to unsettled positions decreases, asf.

Table 2. Comparison of Vulnerability Markets

	Criterion			
			Risk-	
Market type	Infomation	Incentives	balancing	Efficiency
Bug challenges	−	+	−−	−
Vulnerability brokers	−−	±	−−	−−
Exploit derivatives	++	+	+	+
Cyber-insurance	+	++	++	−

Symbols ranging from −− (poor) to ++ (excellent)

4.5 A provisional assessment of market types

Putting the three functions and the efficiency property together, gives us a framework for a structured comparison of the market types discussed in Section 3. A summary of the correspondence of each market type to the criteria is given in Table 2. Note that the evaluation is based on a qualitative assessment and should be regarded as a starting point for exchanges of view rather than as outright evidence. Some arguments backing the relative assessment of different market types are given below.

The incentive function is fulfilled by all market types, though to varying degree. The ambivalent evaluation for vulnerability brokers is due to the questionable incentives created for adversaries to join the circle in order to obtain sensitive vulnerability information before the general public [3]. Conversely, we consider cyber-insurance as particularly good at the incentive function because the incentives to give security a higher priority are not limited to bug hunters and developers, but also affect the end user. This fosters security awareness on a large basis.

As to the information function, bug challenges fail to provide accurate indicators when vulnerabilities are reported frequently. Vulnerability brokers do not reveal timely information to the public at all. Even worse, the usual practice of requiring vulnerability discoverers to sign non-disclosure agreements hinders the vital exchange of security-relevant information. We consider exploit derivatives as superior to cyber-insurance, because insurance contracts are re-negotiated less frequently, which negatively affects the timeliness of a price indicator. And it is questionable whether price information on actual cyber-insurance contracts—not merely unspecified offers—will ever be made available to the public on a large and regular basis. This together with the presumably high transaction costs of insurance contracts justifies a slightly negative assessment of cyber-insurance with respect to efficiency.

Bug challenges and vulnerability brokers provide no risk-balancing instruments at all. Exploit derivatives are somewhat worse than cyber-insurance because it is more difficult to manage optimal portfolios for individual risk profiles when the pay-outs are defined by global events rather than by a firm's individual losses.

Overall, it appears that exploit derivatives and cyber-insurance are both acceptable concepts for vulnerability markets, and it is a matter of fact that both can complement one another.

5 Concluding Remarks

This paper contributes to the literature on vulnerability disclosure policy and economics of information security by differentiating classes for vulnerability market concepts. Moreover, criteria for a better comparability of market types with regard to their potential as tools to moderate the flow of security-relevant information have been proposed. An application of this framework to four market types resulted in a qualitative assessment, which may serve as a first guideline for practitioners in the security industry as well as for policy makers on topics related to information security. Primarily, however, it is intended to be a starting point for academic discussion on the basis of further refined analyses and more rigorous models.

As to future research, there remains to be written chapters on possible conflicts of interest, and on the consequences for disclosure policies. The entire comparison could be repeated on the basis of formal models for each of the market types. Although it might be tricky to model all properties, it will help to understand the exact conditions under which each market type performs optimal.

There is also room for more general critiques on the market approach. One might question whether vulnerability hunting actually leads to more secure products because the supply of vulnerabilities is deemed to be unlimited [4]. So why bother putting market incentives in place for something allegedly useless? (See [33] for a discussion and evidence *for* vulnerability hunting.) Moreover, it is well-known that markets tend to err in the short term—but it is still very difficult to outpace existing markets in the long run. Therefore, we need to assess the harm a "vulnerability market bubble" potentially causes, and weight it against the welfare gains from better information, more secure products, and the possibility to hedge information security risks.

Finally, it is important to ask the questions whether a closer link between information security and financial markets is desirable at all from a stability point of view. A higher interdependency between two previously separate systems implies also a larger sensitivity to mutual shocks, even if the now combined system is less likely to face extreme outcomes because of better risk sharing. Whatever mechanisms get implemented in practice, an individual virus author's potential to halt computers in offices all over the world (which already translates to enormous financial losses) must not get leveraged to cause global asset price deterioration.

Acknowledgements

The author gratefully acknowledges the valuable comments he received from Thorsten Holz, Gaurav Kataria, and Andy Ozment.

References

1. Arora, A., Telang, R., Xu, H.: Optimal policy for software vulnerability disclosure. In: *Workshop on the Economics of Information Security (WEIS)*, University of Minnesota, Minneapolis, MN (2004) `http://www.dtc.umn.edu/weis2004/xu.pdf`.
2. Arora, A., Krishnan, R., Telang, R., Yang, Y.: An empirical analysis of vendor response to software vulnerability disclosure. In: *Workshop on Information Systems and Economics (WISE)*, University of California, Irvine, CA (2005)
3. Nizovtsev, D., Thursby, M.: Economic incentives to disclose software vulnerabilities. In: *Workshop on the Economics of Information Security (WEIS)*, Harvard University, Cambridge, MA (2005) `http://infosecon.net/workshop/pdf/20.pdf`.
4. Rescorla, E.: Is finding security holes a good idea? In: *Workshop of Economics and Information Security (WEIS)*, University of Minnesota, Minneapolis, MN (2004) `http://www.dtc.umn.edu/weis2004/rescorla.pdf`.
5. Anderson, R.J.: Why information security is hard – An economic perspective (2001) `http://www.cl.cam.ac.uk/~rja14/econsec.html`.
6. Akerlof, G.A.: The market for 'lemons': Quality, uncertainty and the market mechanism. *Quarterly Journal of Economics* **84** (1970) 488–500
7. Shapiro, C., Varian, H.R.: *Information Rules. A Strategic Guide to the Network Economy.* Harvard Business School Press (1998)
8. Hardin, G.: The tragedy of the commons. *Science* **162** (1968) 1243–1248
9. Freiling, F., Holz, T., Wicherski, G.: Botnet tracking: Exploring a root-cause methodology to prevent distributed denial-of-service attacks. In S. de Capitani di Vimercati et al., ed.: *Proc. of ESORICS.* LNCS 3679, Berlin Heidelberg, Springer Verlag (2005) 319–335
10. Varian, H.R.: System reliability and free riding. In: *Workshop on Economics and Information Security (WEIS)*, Berkeley, CA (2002) `http://www.sims.berkeley.edu/resources/affiliates/workshops/econsecurity/`.
11. Varian, H.R.: Managing online security risks. *New York Times* (2000) `http://www.nytimes.com/library/financial/columns/060100econ-scene.html`.
12. Ryan, D.J., Heckmann, C.: Two views on security software liability. *IEEE Security & Privacy* **1** (2003) 70–75
13. Schechter, S.E.: *Computer Security Strength & Risk: A Quantitative Approach.* PhD thesis, Harvard University, Cambridge, MA (2004)
14. Camp, J.L., Wolfram, C.: Pricing security. In: *Proc. of the CERT Information Survivability Workshop*, Boston, MA (2000) 31–39 `http://www.cert.org/research/isw/isw2000/papers/54.pdf`.
15. Downs, A.: *An Economic Theory of Democracy.* Harper and Brothers, New York (1957)
16. Stigler, G.J.: *The Citizen and the State: Essays on Regulation.* University Press, Chicago (1975)
17. Ozment, A.: Bug auctions: Vulnerability markets reconsidered. In: *Workshop of Economics and Information Security (WEIS)*, University of Minnesota, Minneapolis, MN (2004) `http://www.dtc.umn.edu/weis2004/ozment.pdf`.

18. Böhme, R.: Vulnerability markets – What is the economic value of a zero-day exploit? In: *Proc. of 22C3: Private Investigations*, Berlin, Germany (2005) https://events.ccc.de/congress/2005/fahrplan/attachments/542-Boehme2005_22C3_VulnerabilityMarkets.pdf.
19. Kannan, K., Telang, R.: An economic analysis of markets for software vulnerabilities. In: *Workshop of Economics and Information Security (WEIS)*, University of Minnesota, Minneapolis, MN (2004) http://www.dtc.umn.edu/weis2004/kannan-telang.pdf.
20. Matsuura, K.: Security tokens and their derivatives. Technical report, Centre for Communications Systems Research (CCSR), University of Cambridge, UK (2001)
21. Gordon, L.A., Loeb, M.P., Sohail, T.: A framework for using insurance for cyber-risk management. *Communications of the ACM* **46** (2003) 81–85
22. Kesan, J.P., Majuca, R.P., Yurcik, W.J.: The economic case for cyberinsurance. In: *Workshop on the Economics of Information Security (WEIS)*, Harvard University, Cambridge, MA (2005) http://infosecon.net/workshop/pdf/42.pdf.
23. Schneier, B.: Hacking the business climate for network security. *IEEE Computer* (2004) 87–89
24. Yurcik, W., Doss, D.: Cyberinsurance: A market solution to the internet security market failure. In: *Workshop on Economics and Information Security (WEIS)*, Berkeley, CA (2002) http://www.sims.berkeley.edu/resources/affiliates/workshops/econsecurity/.
25. Böhme, R.: Cyber-insurance revisited. In: *Workshop on the Economics of Information Security (WEIS)*, Harvard University, Cambridge, MA (2005) http://infosecon.net/workshop/pdf/15.pdf.
26. Ettredge, M., Richardson, V.J.: Assessing the risk in e-commerce. In Sprague, R.H., ed.: *Proc. of the 35th Hawaii International Conference on System Sciences*, Los Alamitos, CA, IEEE Press (2002)
27. Campbell, K., Gordon, L.A., Loeb, M.P., Zhou, L.: The economic cost of publicly announced information security breaches: Empirical evidence from the stock market. *Journal of Computer Security* **11** (2003) 431–448
28. Cavusoglu, H., Mishra, B., Raghunathan, S.: The effect of internet security breach announcements on market value: Capital market reactions for breached firms and internet security developers. *International Journal of Electronic Commerce* **9** (2004) 69–104
29. Telang, R., Wattal, S.: Impact of software vulnerability announcements on the market value of software vendors – An empirical investigation. In: *Workshop on the Economics of Information Security (WEIS)*, Harvard University, Cambridge, MA (2005) http://infosecon.net/workshop/pdf/telang_wattal.pdf.
30. Kahneman, D., Tversky, A.: *Choices, Values, and Frames.* Cambridge University Press (2000)
31. Geer et al., D.: CyberInsecurity – The cost of monopoly (2003) http://www.ccianet.org/papers/cyberinsecurity.pdf.
32. Chen, P.Y., Kataria, G., Krishnan, R.: Software diversity for information security. In: *Workshop on the Economics of Information Security (WEIS)*, Harvard University, Cambridge, MA (2005) http://infosecon.net/workshop/pdf/47.pdf.
33. Ozment, A.: The likelihood of vulnerability rediscovery and the social utility of vulnerability hunting. In: *Workshop on the Economics of Information Security (WEIS)*, Harvard University, Cambridge, MA (2005) http://infosecon.net/workshop/pdf/10.pdf.