Chapter 2
The Price of Uncertainty in Security Games

Jens Grossklags, Benjamin Johnson, and Nicolas Christin

Abstract In the realm of information security, lack of information about other users' incentives in a network can lead to inefficient security choices and reductions in individuals' payoffs. We propose, contrast and compare three metrics for measuring the *price of uncertainty* due to the departure from the payoff-optimal security outcomes under complete information. Per the analogy with other efficiency metrics, such as the price of anarchy, we define the price of uncertainty as the maximum discrepancy in expected payoff in a complete information environment versus the payoff in an incomplete information environment. We consider *difference*, *payoff-ratio*, and *cost-ratio* metrics as canonical nontrivial measurements of the price of uncertainty. We conduct an algebraic, numerical, and graphical analysis of these metrics applied to different well-studied security scenarios proposed in prior work (i.e., best shot, weakest-link, and total effort). In these scenarios, we study how a fully rational expert agent could utilize the metrics to decide whether to gather information about the economic incentives of multiple nearsighted and naïve agents. We find substantial differences between the various metrics and evaluate the appropriateness for security choices in networked systems.

Jens Grossklags
Princeton University, Center for Information Technology Policy, Sherrerd Hall, Princeton, NJ 08544, e-mail: jensg@princeton.edu

Benjamin Johnson
Carnegie Mellon University, CyLab, 4720 Forbes Ave, Pittsburgh, PA 15213, e-mail: johnsonb@andrew.cmu.edu

Nicolas Christin
Carnegie Mellon University, CyLab, 4720 Forbes Ave, Pittsburgh, PA 15213, e-mail: nicolasc@andrew.cmu.edu

T. Moore et al. (eds.), *Economics of Information Security and Privacy,*
DOI 10.1007/978-1-4419-6967-5_2, © Springer Science+Business Media, LLC 2010

2.1 Introduction

The importance of (the lack of) information about security threats, response mechanisms, and associated expected losses and cost has long been identified in the computer science, risk management and economics communities. Granick, for example, argues that weaknesses in our understanding of the measurability of losses serve as an impediment in sentencing cybercrime offenders [14]. Swire adds that deterring fraudsters and criminals online is hampered if we cannot correctly aggregate their offenses across different jurisdictions [37].

The question arises how much defenders can gain by investing in techniques or other efforts to improve information availability for decision-making? Swire's analysis foreshadows significant costs to create an information exchange for law enforcement that could support evidence gathering. Similarly, private organizations struggle with how to accumulate data about security risks and incidents in their respective industries. Past work has, for example, considered the role of intermediaries such as Information Sharing & Analysis Centers to create incentives for exchanging and disclosing data between companies. Researchers investigated under which conditions organizations are willing to contribute to an information pool about security breaches and investments when (negative) competitive effects may result from this cooperation [10, 13]. In different contexts disclosure is not always voluntary and companies may question how much profit they squander when undesirable information is released. For example, other economics research explores the impact of (mandated) breach disclosures [5] or publication of software vulnerabilities [38] on the financial market value of corporations. While other work shows that the information gathering or disclosure effect is not always unambiguously positive or negative [7].

This trade-off between cost and benefits of information gathering, sharing or disclosure reappears in many contexts. From a viewpoint of individual rationality it is decided based on the difference of how much the individual can learn in comparison to the advantage gained by attackers or competitors [36].

Our contribution is to propose and evaluate a set of generic metrics that are applicable to different security decision-making situations to help with this trade-off calculation. In particular, we are interested in quantifying the payoff differential that results from the changes in security choices given different information available. In economic terms we thereby refer to the differences in payoff that results from changes in the underlying *information structure* of the scenario that makes explicit the nature of the utility of information to agents [27].

Specifically, we introduce the *price of uncertainty* metric that quantifies the maximum discrepancy in the total expected payoff between exactly two information conditions.[1] Our terminology is made per analogy with Koutsoupias and Papadim-

[1] After our initial proposal of the price of uncertainty [19], Balcan *et al.* published a research study in which they defined the price of uncertainty as the degree that small fluctuations in costs impact the result of natural best-response and improved-response dynamics [3].

itriou’s “price of anarchy” [24]. We consider *difference*, *payoff-ratio*, and *cost-ratio* metrics as canonical nontrivial measurements of the price of uncertainty.

Since the possibilities for the economic formalization of information are vast, we illustrate our approach on an example model for security choices. Specifically, we introduce uncertainty by assuming that each agent faces a randomly drawn probability of being subject to a direct attack. We study how the decisions and payoffs of an individual agent differ if all draws are common knowledge, compared to a scenario where this information is only privately known [18].

We conduct this analysis within the framework of security games [15, 16] to understand the behavior of the price of uncertainty across different canonical interdependency cases: best shot, weakest-link and total effort [39]. We further consider a recent extension of our work in which we distinguish between the roles of a fully rational expert agent and naïve end users [17]. The inexperienced users conduct a simple self-centered cost-benefit analysis, and neglect interdependencies. We analyze the price of uncertainty from the perspective of the expert agent that fully comprehends the benefits of information in the context of the interrelationship with the naïve users [18]. This allows us to make a general observation. The value of information for the expert agent is always weakly positive [27] since naïve users do not strategize based on additional information.

In this model, the price of uncertainty can depend on several different parameters: the cost of security measures, the magnitude of potential losses, the initial security budget or endowment, and the number of other naïve agents. We study the impact of these parameters algebraically, numerically and graphically.

We show that the difference metric of the price of uncertainty increases linearly in losses, L, and decreases super-linearly in the number of agents, N. That is, only in the presence of extremely large losses would a decision-maker strictly prefer to explore the threat probabilities of other agents at a reasonable cost. The payoff-ratio metric is strictly decreasing in N and independent of the magnitude of potential losses, L. Finally, our cost-ratio metric serves as an example for misleading advice because it overemphasizes the need for action in the presence of relatively small costs.

By evaluating the price of uncertainty for a range of parameters in different security scenarios, we can determine which configurations can accommodate limited information environments (i.e., when being less informed does not significantly jeopardize an expert user’s payoff). We also provide a framework for future work in the area of analysis of the value of security-relevant information. For example, we believe that the game-theoretic analysis in specialized scenarios, e.g., intrusion detection games [28], and security patrol versus robber avoidance scenarios [32] can benefit from a substantiation of the significance of informational assumptions by studying the price of uncertainty.

In Section 2.2, we summarize the security games framework we developed in prior work, and detail our assumptions about agent behaviors and information conditions. We present the different metrics for the price of uncertainty and describe our analysis methodology in Section 2.3. We conduct our analysis and discuss the results in Section 2.4. Finally, we close with a discussion and concluding remarks in Section 8.8.

2.2 Decision Theoretic Model

Our current analysis of the *price of uncertainty* is based on the security games framework [15, 16] and our consecutive work that extends this model to an economy consisting of an expert user and several unsophisticated users that follow a simple but reasonable rule-of-thumb strategy [17, 18]. The latter investigation is a decision-theoretic approach [6, 12]. In the following, we present the key aspects of our model.

2.2.1 Basic Model

Self-protection and self-insurance. In practice, the action portfolio of a defender may include different options to prevent successful compromises and to limit losses that result from a breach. In Grossklags *et al.* [15] we provide a model that allows a decoupling of investments in the context of computer security. On the one hand, the perimeter can be strengthened with a higher self-protection investment (e.g., implementing or updating a firewall). On the other hand, the amount of losses can be reduced by introducing self-insurance technologies and practices (e.g., backup provisions). Formally, player i decides whether to invest in protection ($e_i = 1$) or not ($e_i = 0$). Similarly, each player can adopt a self-insurance technology ($s_i = 1$) or not ($s_i = 0$). In other words, e_i and s_i are two discrete decision variables.

Discrete choice decision-making captures many practical security problems. Examples include purchase and adoption investments as well as updating and patching of protection and self-insurance technologies [2, 25, 29, 30]. We have further conducted a sensitivity analysis with respect to the discrete choice assumption and find that, for the study in the present paper, the only differences between the discrete and continuous cases (where e_i and s_i are continuous variables over the interval $[0, 1]$ as opposed to be mere binary variables) arise when there is strict equality between some of the terms in our case-specifying inequality conditions (see derivations in [18]). We believe that focusing on these boundary cases is of limited practical applicability, and could even be misleading. For comparison, we refer to our prior work where we considered the continuous case in a full information environment [15].

We further denote by $b \geq 0$ and $c \geq 0$ the cost of protection and self-insurance, respectively, which are homogeneous for the agent population. So, player i pays be_i for protection and cs_i for self-insurance.

Interdependency. Decisions by one defender frequently influence the incentives for security investments by her peers [39]. For example, the lack of protection efforts by a subset of agents will often allow an attacker to also compromise resources of other agents if a common perimeter is breached. We denote H as a "contribution" function that characterizes the effect of e_i on agent's utility U_i, subject to the protection levels chosen (contributed) by *all* other players. We require that H be defined for all values over $[0,1]^N$. We distinguish three canonical cases that we discussed in-depth in prior work [15]:

- Best shot: $H = \max(e_i, e_{-i})$.
- Weakest-link: $H = \min(e_i, e_{-i})$.
- Total effort: $H = \frac{1}{N}\sum_k e_k$.

where, following common notation, e_{-i} denotes the set of protection levels chosen by players other than i.

Attack probabilities, network size and endowment. Each of $N \in \mathbb{N}$ agents receives an endowment M. If she is attacked and compromised successfully she faces a maximum loss of L. Her expected loss $p_i L$ is mitigated by a scaling factor p_i randomly drawn from a uniform distribution on $[0,1]$.[2] Instead of interpreting the parameter p_i as the probability of a successful attack; we consider the *expected loss*, $p_i L$, as the primary heterogeneous parameter under consideration. The same familiar notation with p_i considered as a heterogeneous mitigating factor as opposed to an attack probability facilitates this perspective.

The choice to consider a heterogeneous expected loss is motivated by practical considerations, as different targets may have highly variable liabilities, due to their economic, political, or reputational agenda. The choice of a uniform distribution on mitigating factors ensures the analysis remains tractable, while already providing numerous insights. We conjecture that different distributions (e.g., power law) may also be appropriate in practice.

2.2.2 Player Behavior

At the core of our analysis is the observation that expert and non-expert users differ in their understanding of the complexity of networked systems. Indeed, consumers' knowledge about risks and means of protection with respect to privacy and security can be quite varied [1], and field surveys separate between high and low expertise users [34].

Sophisticated (expert) user. Advanced users can rely on their superior technical and structural understanding of computer security threats and defense mechanisms, to analyze and respond to changes in the environment [8]. In the present context, expert users, for example, have less difficulty to conclude that the goal to avoid

[2] Technically, our analysis does not require complete knowledge of the distribution on the various p_i. The distribution informs the probability that a given number of p_j are above the rule-of-thumb threshold; but to conduct our analysis, it suffices to know only these threshold probabilities, and not the full distribution.

censorship points is a best shot scenario, whereas the protection of a corporate network frequently suggests a weakest-link optimization problem [15]. Accordingly, a sophisticated user correctly understands her utility to be dependent on the interdependencies that exist in the network:

$$U_i = M - p_i L(1 - s_i)(1 - H(e_i, e_{-i})) - be_i - cs_i \; .$$

Naïve (non-expert) user. Average users underappreciate the interdependency of network security goals and threats [1, 34]. We model the *perceived* utility of each naïve agent to only depend on the direct security threat and the individual investment in self-protection and self-insurance. The investment levels of other players are *not* considered in the naïve user's decision making, despite the existence of interdependencies. We define the perceived utility for a specific naïve agent j as:

$$PU_j = M - p_j L(1 - s_j)(1 - e_j) - be_j - cs_j \; .$$

Clearly, perceived and realized utility actually differ: by failing to incorporate the interdependencies of all agents' investment levels in their analysis, naïve users may achieve sub-optimal payoffs far below their anticipated expected payoffs. This paper does not aim to resolve this conflict, and, in fact, there is little evidence that users will learn the complexity of network security over time [34]. We argue that non-expert users would repeatedly act in an inconsistent fashion. This hypothesis is supported by findings in behavioral economics that consumers repeatedly deviate from rationality, however, in the same predictable ways [23].

2.2.3 Information Conditions

Our analysis is focused on the decision making of the expert user subject to the bounded rational behaviors of the naïve network participants. That is, more precisely, the expert agent maximizes their expected utility subject to the available information about other agents' drawn threat probabilities and their resulting actions. Two different information conditions may be available to the expert agent:

Complete information: Actual draws of attack probabilities p_j for all $j \neq i$, and her own drawn probability of being attacked p_i.

Incomplete information: Known probability distribution of the unsophisticated users' attack threat, and her own drawn probability of being attacked p_i.

Therefore, the expert agent can accurately infer what each agent's investment levels are in the complete information scenario. Under incomplete information the sophisticated user has to develop an expectation about the actions of the naïve users.

2.2.4 Remarks on Basic Results

We have conducted the basic analysis of this scenario in [18]. Below we are making several general observations to guide the reader through the results in this paper.

Every security scenario (i.e., best-shot, weakest-link and total effort) involves simple cost-benefit analyses for both sophisticated and naïve agents [11]. Agents remain passive when the cost of self-protection and self-insurance exceeds the expected loss. Further, they differentiate between the two types of security actions based on their relative cost. This behavior describes what we would usually consider as basic risk-taking that is part of everyday life: It is not always worth protecting against known risks.

One important feature of our model is the availability of self-insurance. If the cost of self-insurance c is less than the cost of protection b, the decision scenario significantly simplifies for all games and both information conditions. This is because once self-insurance is applied, the risk and interdependency among the players is removed. The interesting cases for all three games arise when $b \leq c$ and protection is a potentially cost-effective option. Within this realm insurance has a more subtle effect on the payoffs.

Tables 2.1, 2.2 and 2.3 contain the total expected payoff for decisions made by the sophisticated agent, but also for the naïve agents. We have already highlighted that for $c < b$ all agents follow the same simple decision rule to decide between passivity and self-insurance. Therefore, payoffs in this region are identical for all agent types in the case of homogeneous security costs. But, there are payoff differences among all three information conditions for some parts of the parameter range when $b \leq c$.

It is intuitive that the naïve agents suffer in the weakest-link game since they do not appreciate the difficulty to achieve system-wide protection. Similarly, in the best shot game too many unsophisticated agents will invest in protection lowering the average payoff. In the total effort game, sophisticated agents realize that their contribution is only valued in relation to the network size. In comparison, naïve agents invest more often. Further, the payoff profile of the unsophisticated agents remains flat for $b < c$. This reflects the fact that the naïve agent ignores the insurance option whenever protection is cheaper.

We can observe that the sophisticated agents will suffer from their misallocation of resources in the weakest-link game when information is incomplete. In the best shot game this impact is limited, but there is a residual risk that no naïve agent willingly protects due to an unlikely set of draws. In such cases the fully informed expert could have chosen to take it upon herself to secure the network. In the total effort game we observe a limited payoff discrepancy for expert users as a result of limited information.

2.2.5 Outlook on Further Analyses

Above we have provided a short summary of the key results that help to distinguish the three canonical scenarios and the decision-making of the expert and naïve agents (as detailed in [18]). From this point on we venture into new territory.

We start with the total payoff results in Tables 2.1, 2.2, and 2.3 and derive metrics to compare the impact of the important decision making parameters on the payoffs achievable in the two different information conditions. Thereby, we focus on the choices and payoffs garnered by the expert agent.

2.3 Price of Uncertainty Metrics

2.3.1 The Price of Uncertainty

In previous work we discussed two information conditions (complete information and incomplete information) for an expert player in three canonical security games. In this context, the price of uncertainty measures the disadvantage of the expert player when she has incomplete information, compared to when she has complete information. Depending on the form this measure takes, the price of uncertainty potentially depends on five different parameters:

1. Cost of protection b,
2. Cost of insurance c,
3. Magnitude of potential losses L,
4. Initial endowment M, and
5. Number of other players N.

Because the analysis of five-variable functions is somewhat cumbersome, a central objective in our metric-creation exercise is to reduce the number of parameters in a manner such that something both relevant and interesting can be said. Therefore, we focus on how the price of uncertainty depends on the magnitude of potential losses L and the number of other players N. To eliminate M we choose a canonical value of either 0 or L, and to eliminate b and c we chose the values that cause the price of uncertainty to have the greatest significance. This choice depends on the metric.

2.3.2 Three Metrics for the Price of Uncertainty

For each of the security games (i.e., best shot, weakest link, and total effort), we define and analyze three metrics for the price of uncertainty:

1. The difference metric $PoU_1(L,N)$, defined by

$$\max_{b,c\in[0,L]} [\text{Expected Payoff Complete}(b,c,L,L,N) -$$
$$\text{Expected Payoff Incomplete}(b,c,L,L,N)]$$

2. The payoff-ratio metric $PoU_2(L,N)$ defined by

$$\max_{b,c\in[0,L]} \left[\frac{\text{Expected Payoff Complete}(b,c,L,L,N)}{\text{Expected Payoff Incomplete}(b,c,L,L,N)}\right]$$

3. The cost-ratio metric $PoU_3(L,N)$ defined by

$$\min_{b,c\in[0,L]} \left[\frac{\text{Expected Payoff Complete}(b,c,L,0,N)}{\text{Expected Payoff Incomplete}(b,c,L,0,N)}\right]$$

2.3.3 Discussion of the Definitions

2.3.3.1 The Difference Metric

The difference metric is our most straightforward metric. It says the price of uncertainty is the worst case difference in payoff between complete and incomplete information, where the maximum is taken over all possible prices for protection and self-insurance. In this metric, a completely insignificant price of uncertainty yields an output of zero, and the metric's output increases directly as the price of uncertainty becomes more significant.

2.3.3.2 The Payoff-Ratio Metric

The payoff-ratio metric is motivated by the game-theoretic notion of the "price of anarchy", which is defined as a payoff-ratio of a game's socially optimal equilibrium to its worst case Nash equilibrium [24]. By analogy, we define the price of uncertainty as the worst case ratio between the payoffs for the expert with complete information to the expert with incomplete information, with the worst case taken over all possible prices of protection and self-insurance. One advantage of using a ratio-style metric of this type is that its output is currency-independent. In other words, while our difference metric might depend on say dollars or euros, this ratio metric is just a pure number. In the payoff-ratio metric, a completely insignificant price of uncertainty yields an output of 1, and the metric's output *increases* as the price of uncertainty becomes more significant.

2.3.3.3 The Cost-Ratio Metric

The cost-ratio metric is similar to the payoff-ratio metric, but with a different canonical choice of 0 for the initial endowment M. This metric directly measures the ratio of costs induced by the expert's choices. These costs are reflected in formulas involving b, c, L, and N. Mathematically, the cost-ratio allows for a simpler algebraic analysis due to an abundance of term cancellations. A minor disadvantage of this metric's formulation is that it has a somewhat nonstandard orientation, in the sense that it decreases as the price of uncertainty becomes more significant. There are two justifications for this choice. First, we wanted to cast this metric as being a simpler analogue to the payoff-ratio metric. Second, we wanted to avoid values at infinity, which would have resulted had we used this metric's multiplicative inverse. In our cost-ratio metric, a completely insignificant price of uncertainty yields an output of 1, and the metric's output *decreases* toward zero as the price of uncertainty becomes more significant.

2.4 Analysis

In this section, we analyze the price of uncertainty as defined by each of the three metrics in each of the canonical security games. In each case the analysis proceeds as follows. First, considering the magnitude of potential loss L and the number of other players N as fixed parameters, we determine the protection cost b and self-insurance cost c which cause the metric under consideration to yield its most significant value. This process defines a function of two parameters L and N, which we then analyze as a measure of the price of uncertainty. In some scenarios we are able to produce clean algebraic results with tight asymptotic bounds. For others we must rely almost completely on computer-aided numerical analysis and graphs. Each subsection contains graphs of all relevant metrics and maximizing parameters, and concludes with some important observations.

2.4.1 Best Shot Game

In the best shot game (introduced in [15]), the security of the network is determined by the protection level of the individual with the highest investment. The relevant expected payoffs for an expert player in the best shot game are shown in Table 2.1. These results will be used throughout this section. Complete derivations for these payoffs can be found in [18].

Table 2.1 Best shot security game: Total expected game payoffs.

Case Name	Case Condition	Information Type	Total Expected Payoff
BC1	$c < b$	Complete	$M - c + \frac{c^2}{2L}$
BC2	$b \leq c$	Complete	$M - b\left(1 - \frac{b}{2L}\right)\left(\frac{b}{L}\right)^{N-1}$
BI1	$c < b$	Incomplete	$M - c + \frac{c^2}{2L}$
BI2	$b \leq c$	Incomplete	$M - \frac{L}{2}\left(\frac{b}{L}\right)^{N-1}$
BN1	$c < b$	Naive	$M - c + \frac{c^2}{2}$
BN2	$b \leq c$	Naive	$M - b + \frac{b^2}{2L}$

2.4.1.1 The Best Shot Difference Metric: $BPoU_1(L,N)$

In this section, we analyze the price of uncertainty metric $BPoU_1(L,N)$ defined as:

$$\max_{b,c \in [0,L]} [\text{Best Shot Expected Payoff Complete}(b,c,L,M,N) - \text{Best Shot Expected Payoff Incomplete}(b,c,L,M,N)] \tag{2.1}$$

In the best shot game, the complete and incomplete payoffs are the same when $c < b$; hence to compute the maximum payoff difference we may assume that $b \leq c$. Observe that in this case, the payoffs do not depend on c at all. This will help to simplify our analysis, and in fact allows us to compute $BPoU_1(L,N)$ in a purely algebraic manner.

We find that any b maximizing this equation satisfies

$$b = L \cdot \left(\frac{N-1}{N+1}\right),$$

and that consequently,

$$BPoU_1(L,N) = 2L \cdot \frac{(N-1)^{N-1}}{(N+1)^{N+1}}. \tag{2.2}$$

To give an asymptotic analysis, we begin by noting that $\lim_{n\to\infty}\left(\frac{N-1}{N+1}\right)^{N-1} = \frac{1}{e^2}$. Rewriting the expression above as $2L\left(\frac{N-1}{N+1}\right)^{N-1} \cdot \frac{1}{(N+1)^2}$, we see that the first part approaches $\frac{2L}{e^2}$ as N gets large, and that the second part decreases to zero quadratically in $\frac{1}{N}$. Hence this metric for the price of uncertainty increases linearly in L for fixed N and decreases quadratically to zero in $\frac{1}{N}$ for fixed L. Figure 2.1(a) shows a graph of the maximizing b for $BPoU_1$ as a function of N and L; while Figure 2.1(b) shows a graph of the metric $BPoU_1$ as a function of N and L. A complete algebraic derivation is also available in the workshop version of this paper [19].

Observations. The interpretation of our numerical results for this metric is that the price of uncertainty increases with the potential losses, but as the number of players

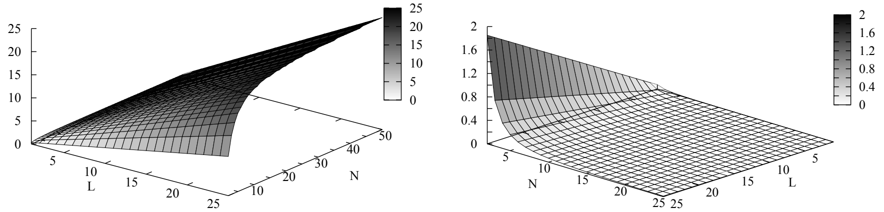

(a) The maximizing b for $BPoU_1(L,N)$

(b) $BPoU_1(L,N)$

Fig. 2.1 Best shot – Difference metric: $BPoU_1(L,N)$. The metric grows linearly in the potential loss L for a fixed network size N, and decreases inverse-quadratically in the network size N for a fixed loss L.

increases, the price of uncertainty diminishes (unless the losses are quite high) and approaches the square of the number of players.

2.4.1.2 The Best Shot Payoff-Ratio Metric $BPoU_2(L,N)$

In this section, we analyze the price of uncertainty metric $BPoU_2(L,N)$, defined as

$$\max_{b,c\in[0,L]}\left[\frac{\text{Best Shot Expected Payoff Complete}(b,c,L,L,N)}{\text{Best Shot Expected Payoff Incomplete}(b,c,L,L,N)}\right]. \tag{2.3}$$

After substituting $B=\frac{b}{L}$ we may derive

$$BPoU_2(L,N)=\max_{B\in[0,1]} 1+\frac{\frac{1}{2}B^{N-1}(1-B)^2}{1-\frac{1}{2}B^{N-1}},$$

and the maximizing B for this equation occurs when

$$0=\frac{1-B}{2}B^{N-2}\left(B^N-B(N+1)+N-1\right).$$

Observing that $B^N-B(N+1)+N-1$ is positive at $B=0$ and negative at $B=1$, and making additional arguments, it can be shown that this equation has exactly one solution in $(0,1)$. Due to well-known algebraic results, this solution cannot be expressed algebraically for $N\geq 5$, but we can plot the solution graphically. Figure 2.2 Grossklags/plots a graph of the maximizing $b=LB$ as a function of N and L. Figure 2.2(b) Grossklags/plots $BPoU_2$ as a function of N. As can be seen from the graph (or from our derivation), this metric does not depend on L, and it approaches 1 as N increases.

Observations. Since 1 represents the smallest price possible in this metric, the interpretation would be that the price of uncertainty is independent of the magnitude of potential losses and becomes insignificant as the number of players increases.

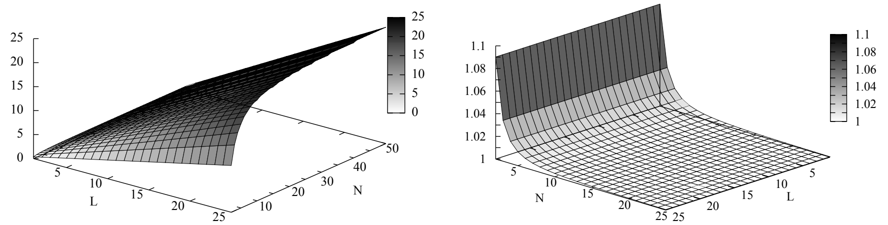

(a) The maximizing b for $BPoU_2(L,N)$

(b) $BPoU_2(L,N)$

Fig. 2.2 Best shot – Payoff-ratio metric: $BPoU_2(L,N)$. The metric is independent of L.

2.4.1.3 The Best Shot Cost-Ratio Metric $PoU_3(B,L,N)$

In this section, we analyze the price of uncertainty metric $BPoU_3(L,N)$, defined as

$$\min_{b,c\in[0,L]} \left[\frac{\text{Best Shot Expected Payoff Complete}(b,c,L,0,N)}{\text{Best Shot Expected Payoff Incomplete}(b,c,L,0,N)} \right]. \tag{2.4}$$

This metric is expressed in terms of our payoff functions, but by starting with an initial endowment of zero, it becomes a ratio of costs. If the cost of limited information is great compared to the cost of complete information, this ratio will tend toward zero. On the other hand, if the costs are similar, then the ratio will tend toward one. We select the minimizing b and c for this ratio so as to obtain the most significant price of uncertainty under the metric. Using this strategy, we obtain

$$BPoU_3(L,N) = \min_{b\in[0,L]} \frac{2b}{L}\left(1 - \frac{b}{2L}\right).$$

Clearly the minimum value (of zero) for this expression (assuming $0 \leq b \leq L$) is achieved by taking $b = 0$. This cost-ratio metric always measures the price of uncertainty at its greatest possible value, independent of N or L.

Observations. The most direct interpretation for this result would be that the price of uncertainty is very significant, regardless of the number of players or the potential losses. An alternative, and arguably better explanation is that this particular metric is not a very useful provider of information for the best shot game.

2.4.2 Weakest Link Game

In the weakest link game (introduced in [15]), the security of the network is determined by the protection level of the individual with the lowest protection investment. The relevant expected payoffs for the weakest link game are shown in Table 2.2. These results will be used throughout this section. Complete derivations for these payoffs can be found in [18].

Table 2.2 Weakest link security game: Total expected game payoffs.

Case Name	Case Condition	Information Type	Total Expected Payoff
WC1	$c < b$	Complete	$M - c + \frac{c^2}{2L}$
WC2	$b \le c$	Complete	$M - c$ $+ \frac{c^2}{2L} + (c-b)\left(1 - \frac{c+b}{2L}\right)\left(1 - \frac{b}{L}\right)^{N-1}$
WI1	$c < b$	Incomplete	$M - c + \frac{c^2}{2L}$
WI2	$b \le c \le \frac{b}{\left(1-\frac{b}{L}\right)^{N-1}}$	Incomplete	$M - c + \frac{c^2}{2L}$
WI3	$\frac{b}{\left(1-\frac{b}{L}\right)^{N-1}} < c < b + L\left(1 - \left(1 - \frac{b}{L}\right)^{N-1}\right)$	Incomplete	$M - c$ $+ \frac{b^2}{2L\left(1-\frac{b}{L}\right)^{N-1}} + \frac{(c-b)^2}{2L\left(1-\left(1-\frac{b}{L}\right)^{N-1}\right)}$
WI4	$\frac{b}{\left(1-\frac{b}{L}\right)^{N-1}} < b + L\left(1 - \left(1 - \frac{b}{L}\right)^{N-1}\right) \le c$	Incomplete	$M - b$ $- \frac{L}{2}\left(1 - \left(1 - \frac{b}{L}\right)^{N-1}\right) + \frac{b^2}{2L\left(1-\frac{b}{L}\right)^{N-1}}$
WN1	$c < b$	Naive	$M - c + \frac{c^2}{2}$
WN2	$b \le c$	Naive	$M - b$ $+ \frac{b^2}{2L} - \frac{L}{2}\left(1 - \frac{b^2}{L^2}\right)\left(1 - \left(1 - \frac{b}{L}\right)^{N-1}\right)$

In the weakest link game, the complete and incomplete payoffs are the same when $c < b$, but for $b \le c$ there is a wide variety of cases to consider, and without some direction it is not obvious which direction we should take in our analysis. Unlike the best shot game in which most of our equational analysis involved a single variable b in a relatively simple expression, a soft algebraic analysis of the weakest link game is much more difficult to conduct. Our strategy is to use numerical approximations and graphs to determine which cases to consider, and consequently which equations to work with. Thus, most of our algebraic work for this game takes the form of supporting, verifying, and clarifying the numerical analysis.

2.4.2.1 The Weakest Link Difference Metric: $WPoU_1(L,N)$

In this section, we analyze the price of uncertainty metric $WPoU_1(L,N)$ defined as:

$$\max_{b,c \in [0,L]} [\text{Weakest Link Expected Payoff Complete}(b,c,L,L,N) - \\ \text{Weakest Link Expected Payoff Incomplete}(b,c,L,L,N)]. \quad (2.5)$$

Our numerical analysis of this difference metric indicates that all the highest values lie in the weakest link game's case WI3, in which we have $\frac{b}{\left(1-\frac{b}{L}\right)^{N-1}} < c < b + L\left(1 - \left(1 - \frac{b}{L}\right)^{N-1}\right)$. Using this, we may derive an expression for this metric involving equations; however the minimizing values of b and c that yield the final solution are roots of polynomial equations whose degree depends on N. Here we will dispense with the partial derivations and refer the reader to the graphs. Figure 2.3 gives the maximizing b and c (respectively) as functions of L and N. Then, Figure 2.4 gives the weakest link difference metric $WPoU_1$ as a function of L and N.

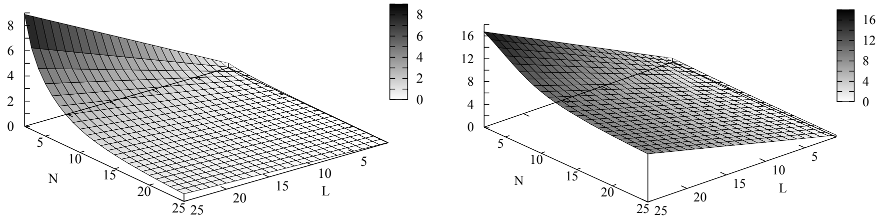

(a) Maximizing b

(b) Maximizing c

Fig. 2.3 Weakest Link – Difference metric: The maximizing b and c (respectively) for $WPoU_1(L,N)$.

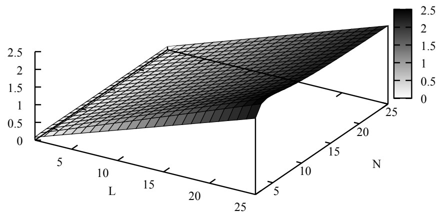

Fig. 2.4 Weakest Link – Difference metric: $WPoU_1(L,N)$. The metric grows linearly in the losses L and remains relatively constant for fixed L regardless of the network size N.

Observe that the maximizing b decreases to 0 as a function of N but increases linearly in L. The maximizing c also decreases in N and increases linearly in L. The difference metric itself increases linearly in L, but remains relatively-constant as N grows. This phenomenon can be explained by the following observation. The maximizing b for this metric satisfies the relation $\frac{b}{L} \in O\left(\frac{1}{N}\right)$, whence the expression $\left(1 - \frac{b}{L}\right)^{N-1}$ approaches a constant as N increases. All terms in $WPoU_1(L,N)$ involving N have this form; thus as N grows the function value does not change. The graph shows additionally that the convergence to a constant value is quite fast in N.

Observations. The interpretation for these numerical results is that the price of uncertainty in the weakest link game is highest when protection is cheap and self-insurance is competitively priced. The price of uncertainty increases directly with the potential loss, and is unaffected by the number of other players.

2.4.2.2 The Weakest Link Payoff-Ratio Metric $WPoU_2(L,N)$

In this section, we analyze the price of uncertainty metric $WPoU_2(L,N)$, defined as

$$\max_{b,c\in[0,L]}\left[\frac{\text{Weakest Link Expected Payoff Complete}(b,c,L,L,N)}{\text{Weakest Link Expected Payoff Incomplete}(b,c,L,L,N)}\right]. \tag{2.6}$$

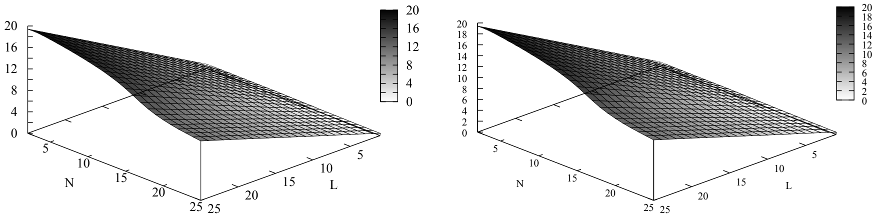

(a) Maximizing b (b) Maximizing c

Fig. 2.5 Weakest Link – Payoff-ratio metric: The maximizing b and c (respectively) for $WPoU_2(L,N)$.

We begin by considering the graphs in Figure 2.5, which give as functions of L and N the b and c (respectively) which maximize the price of uncertainty under this metric. We see that the maximizing b increases linearly with L, but decreases to zero super-linearly in $\frac{1}{N}$. The maximizing c also increases linearly with L, and decreases with N. For the weakest link payoff-ratio metric, we observe that the metric has no dependence on L, and that there is a local maximum very close to $N = 4$, and that after $N = 4$ the ratio decreases toward zero as N increases.

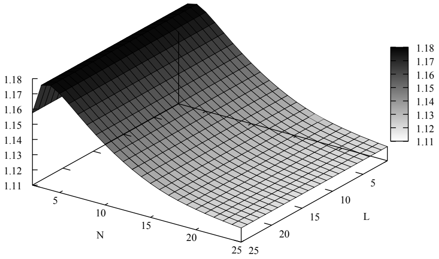

Fig. 2.6 Weakest Link – Payoff-ratio metric: $WPoU_2(L,N)$. Numeric simulations confirm the metric is independent of L.

The graph for the payoff-ratio metric is given in Figure 2.6. We see from the figure that the metric does not depend on L. We can also derive this observation by considering the equations as we did in the best shot case, specifically noting that it is without loss of generality to consider a maximum over $\frac{b}{L}$ and $\frac{c}{L}$ in place of b and

c, respectively. Because the metric only depends on $\frac{b}{L}$ and $\frac{c}{L}$ with the conditions $0 \leq b, c \leq L$, it follows that $L = 1$ without loss of generality, and hence the metric does not depend on L.

Observations. We observe that in the weakest link payoff-ratio metric, the price of uncertainty is highest when there are exactly 4 players, and it decreases toward its minimum possible value as the number of players increases.

2.4.2.3 The Weakest Link Cost-Ratio Metric $WPoU_3(L,N)$

In this section, we analyze the price of uncertainty metric $WPoU_3(L,N)$, defined as

$$\min_{b,c\in[0,L]} \left[\frac{\text{Weakest Link Expected Payoff Complete}(b,c,L,0,N)}{\text{Weakest Link Expected Payoff Incomplete}(b,c,L,0,N)} \right]. \tag{2.7}$$

Plotting as functions of L and N the b and c (respectively) which maximize the price of uncertainty under this metric (not shown for space purposes) shows that the maximum value for b is always achieved when b (and consequently $\frac{b}{L}$) is close to zero. The maximizing c is attained when $\frac{c}{L}$ is scaled with $\frac{b}{L}$ appropriately.

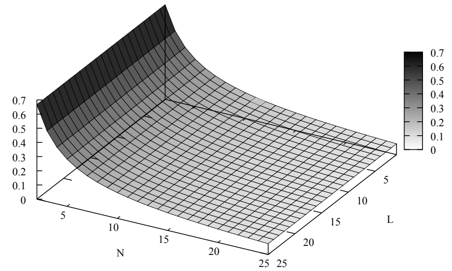

Fig. 2.7 Weakest Link – Cost-ratio metric: $WPoU_3(L,N)$.

The graph for the payoff ratio metric is given in Figure 2.7. As with the payoff-ratio metric considered above, this ratio-based metric does not depend on L. The plot gives nonzero values for all N but decreases to zero as N increases. Recall that zero in this metric represents the most significant price of uncertainty.

Observations. The results for this metric can be interpreted as saying that the price of uncertainty becomes more significant as the number of players increases. This interpretation contradicts our observations in the difference and payoff-ratio metrics for this game, and serves as a prime example to illustrate that the choice of metric makes a significant difference in the interpretation. Our explanation of the discrepancy is that this cost-ratio metric focuses on comparing costs which are insignificantly small in both the complete and incomplete information environments, but whose limiting ratio indicates a significant discrepancy. Based on this observation, a blunt assessment is that the cost-ratio metric for the weakest link game does not measure what we most generally think of as important.

2.4.3 Total Effort Game

In the total effort game (introduced in [15]), the security of the network is determined by the average protection level of all individual players in the network. The relevant expected payoffs for the total effort game are shown in Table 2.3. These results will be used throughout this section. Complete derivations for these payoffs can be found in [18].

Table 2.3 Total effort security game: Total expected payoffs.

Case Name	Case Condition	Information Type	Total Expected Payoff
TC1	$c < b$	Complete	$M - c + \frac{c^2}{2L}$
TC2	$bN \leq L$ and $b \leq c$	Complete	$\sum_{k=0}^{\lfloor N-\frac{c}{b}\rfloor} Pr[k]\cdot\left(M - c + \frac{c^2}{2L\left(1-\frac{k}{N}\right)}\right)$ $+\sum_{k=\lfloor N-\frac{c}{b}+1\rfloor}^{\lfloor N-1-\frac{N}{L}(c-b)\rfloor} Pr[k]\cdot\left(M - c + \frac{b^2N}{2L} + \frac{(c-b)^2}{2L\left(1-\frac{k+1}{N}\right)}\right)$ $+\sum_{k=\lfloor N-\frac{N}{L}(c-b)\rfloor}^{N-1} Pr[k]\cdot\left(M - b - \frac{L}{2}\left(1-\frac{k+1}{N}\right) + \frac{b^2N}{2L}\right)$
TC2	$L < bN$ and $b \leq c$	Complete	$\sum_{k=0}^{\lfloor N-\frac{cN}{L}\rfloor} Pr[k]\cdot\left(M - c + \frac{c^2}{2L\left(1-\frac{k}{N}\right)}\right)$ $+\sum_{k=\lfloor N-\frac{cN}{L}+1\rfloor}^{N-1} Pr[k]\cdot\left(M - \frac{L}{2N}(N-k)\right)$
TI1	$c < b$	Incomplete	$M - c + \frac{c^2}{L}$
TI2	$bN \leq L$ and $b \leq c \leq b + \frac{b^2}{L}(N-1)$	Incomplete	$M - c + \frac{c^2}{2\left(b+\frac{L-b}{N}\right)}$
TI3	$bN \leq L$ and $b + \frac{b^2}{L}(N-1) < c < 2b - \frac{b}{N}$	Incomplete	$M - c + \frac{b^2N}{2L} + \frac{(c-b)^2}{2\left(b-\frac{b}{N}\right)}$
TI4	$bN \leq L$ and $2b - \frac{b}{N} \leq c$	Incomplete	$M - b - \frac{1}{2}\left(b - \frac{b}{N}\right) + \frac{b^2N}{2L}$
TI5	$L < bN$ and $b \leq c < b + \frac{L-b}{N}$	Incomplete	$M - c + \frac{c^2}{2\left(b+\frac{L-b}{N}\right)}$
TI6	$L < bN$ and $b + \frac{L-b}{N} \leq c$	Incomplete	$M - \frac{1}{2}\left(b + \frac{L-b}{N}\right)$
TN1	$c < b$	Naive	$M - c + \frac{c^2}{2}$
TN2	$b \leq c$	Naive	$M - b - \frac{1}{2}\left(b - \frac{b}{N}\right) + \frac{b^2}{L}\left(1 - \frac{1}{2N}\right)$

2.4.3.1 The Total Effort Difference Metric: $TPoU_1(L,N)$

In this section, we analyze the price of uncertainty metric $TPoU_1(L,N)$ defined as:

$$\max_{b,c\in[0,L]} [\text{Total Effort Expected Payoff Complete}(b,c,L,M,N) - \\ \text{Total Effort Expected Payoff Incomplete}(b,c,L,M,N)]. \tag{2.8}$$

As with the weakest link game, there are a number of cases to consider when beginning to analyze the price of uncertainty metrics. Numerical evidence suggests that the maximizing b and c for this game are in the total effort game's case TI3, in which we have $bN \leq L$ and $b + \frac{b^2}{L}(N-1) < c < 2b - \frac{b}{N}$. Using the payoff equations from this case, we can make some progress toward an algebraic solution, deriving the following condition for (b,c) to maximize the payoff difference:

$$c = \frac{\sum_{k=\lfloor N-\frac{c}{b}+1\rfloor}^{\lfloor N-1-\frac{N}{L}(c-b)\rfloor} \left(\frac{Pr[k]}{L\left(1-\frac{k+1}{N}\right)}\right) - \sum_{k=\lfloor N-\frac{N}{L}(c-b)\rfloor}^{N-1} Pr[k] - \frac{b}{\left(b-\frac{b}{N}\right)}}{\sum_{k=0}^{\lfloor N-\frac{c}{b}\rfloor} \left(\frac{Pr[k]}{L\left(1-\frac{k}{N}\right)}\right) + \sum_{k=\lfloor N-\frac{c}{b}+1\rfloor}^{\lfloor N-1-\frac{N}{L}(c-b)\rfloor} \left(\frac{Pr[k]}{L\left(1-\frac{k+1}{N}\right)}\right) - \frac{1}{b-\frac{b}{N}}}.$$

This equation meets the frontiers of our algebraic simplification skills and motivates our haste in proceeding to the numerical analysis. Figure 2.8 Grossklags/plots the price of uncertainty as a function of N and L. We observe that the price of uncertainty in this metric increases linearly in L and decreases to zero with N significantly more quickly than $\frac{1}{N}$.

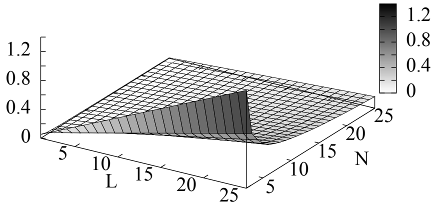

Fig. 2.8 Total Effort – Difference metric: $TPoU_1(L,N)$.

Observations. The interpretation of our numerical results for this metric is that the price of uncertainty increases with the potential losses, but as the number of players increases, the price of uncertainty diminishes quickly.

2.4.3.2 The Total Effort Payoff-Ratio Metric: $TPoU_2(L,N)$

In this section, we analyze the price of uncertainty metric $TPoU_2(L,N)$ defined as:

$$\max_{b,c\in[0,L]} \left[\frac{\text{Total Effort Expected Payoff Complete}(b,c,L,L,N)}{\text{Total Effort Expected Payoff Incomplete}(b,c,L,L,N)}\right]. \tag{2.9}$$

For the remaining total effort metrics, our analysis relies exclusively on numerical approximations. Figure 2.9(b) Grossklags/plots the total effort game's payoff-ratio price of uncertainty as a function of N. The figure shows that the price of uncertainty does not depend on L and that it decreases toward 1 as N increases.

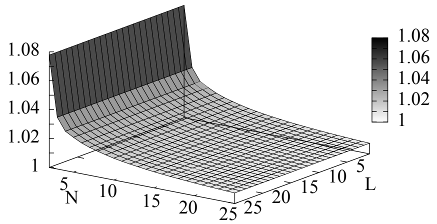

Fig. 2.9 Total Effort – Payoff-ratio metric: $TPoU_2(L,N)$.

Observations. In the total effort game, the payoff-ratio metric depends only on the number of players, and it diminishes to its least significant possible value as the number of players increases.

2.4.3.3 The Total Effort Cost-Ratio Metric: $TPoU_3(L,N)$

In this section, we analyze the price of uncertainty metric $TPoU_3(L,N)$ defined as:

$$\max_{b,c\in[0,L]}\left[\frac{\text{Total Effort Expected Payoff Complete}(b,c,L,0,N)}{\text{Total Effort Expected Payoff Incomplete}(b,c,L,0,N)}\right]. \qquad (2.10)$$

Figure 2.10(c) Grossklags/plots the total effort game's cost-ratio price of uncertainty as a function of N. As can be seen from the graph, the price of uncertainty does not depend on L, and decreases as N increases.

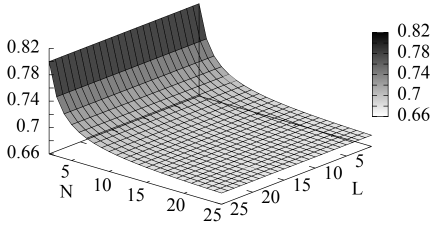

Fig. 2.10 Total Effort – Cost-ratio metric: $TPoU_3(L,N)$.

Observations. Using the cost-ratio metric for the total effort game, the price of uncertainty becomes more significant with an increase in the number of players. Once again this goes against the analogous conclusions drawn with the other two metrics. We surmise that this happens because the cost-ratio metric focuses on the cases where the costs for the complete and incomplete information scenarios are quite small, while the ratio indicates a significant distinction.

2.5 Conclusions

Users frequently fail to deploy, or upgrade security technologies, or to carefully preserve and backup their valuable data [22, 31], which leads to considerable monetary losses to both individuals and corporations every year. This state of affairs can be partly attributed to economic considerations.

Significant challenges for average users arise when they have to determine optimal security strategies in the presence of interdependencies between security choices of other agents [15, 25]. Struggling with this task we anticipate the vast majority of users to be naïve, and to apply approximate decision-rules that fail to accurately appreciate the impact of their decisions on others [1].

In this paper we continue our investigation into the incentives of an individual expert user that rationally responds to the security choices of unsophisticated end-users under different informational assumptions [18]. In particular, we study how the expert evaluates the importance of improving the information available for her decision-making. We propose three variations of the *price of uncertainty* metric that may serve as a decision help for the expert user. We distinguish between a difference, a payoff-ratio, and a cost-ratio metric.

Our work complements the rich area of security metrics that are commonly technical, financial [21] or market-based [4]. However, the price of uncertainty is motivated by game-theory and, more specifically, by Koutsoupias and Papadimitriou's metric to evaluate worst-case equilibria [24], and adds to the rich literature on information sharing, (mandatory) disclosure, and notice and consent that we reviewed in the introductory section.

Our research yields a number of somewhat counter-intuitive results:

- Using cost-ratio metrics can be misleading, as two negligible costs in front of a large endowment may still produce a large ratio when divided by each other. While mathematically trivial, such a pitfall is relatively easy to get into. We showed that, unfortunately, for *all* games we studied, cost-ratios are *never* an appropriate metric. The cynic in ourselves could actually point out that their main use would be for marketing purposes. Beware of snake oil!
- Aside from the cost-ratio metric, the other metrics show a relatively low price of uncertainty across all the scenarios we considered, and this is especially true with a large number of players. The difference metric shows some signs of a penalty for lack of information, but if we consider the absolute payoff values (reported in Tables 2.1, 2.2, and 2.3) we find the price of uncertainty in the difference metric is at most 20% of the magnitude of the potential loss. Accordingly, we can summarize that in scenarios with many players the lack of information does not penalize an expert too much. On the other hand, the lack of knowledge (about interdependencies) that makes a user naïve, as opposed

to expert, results in significant payoff degradation regardless of the number of players [18].

- Assuming fixed possible losses, the more players are in a network, the less information matters. This is actually good news, as full information typically gets increasingly difficult to gather as the number of players grows large.
- In contrast to our arguments in favor of difference-based metrics behavioral research has shown that individuals are frequently influenced by ratio-difference evaluations [33]. However, this makes consumers more vulnerable to (numerical) framing differences that change perceptions about the benefits of additional information. For example, experimental research has reported robust evidence for consumers' preferences for benefits that are presented as large ratios in comparison to small ratios [26]. In the security context, marketers could easily switch the framing from a security to a reliability measure and thereby vary the size of the benefit ratio (e.g., from 3% vs. 5% failure to 97% vs. 95% reliability). As a result, individuals may exaggerage the importance of changes when risks or benefits are small [20, 35].
- We have also shown that the payoff-ratio and the cost-ratio metrics are independent of the size of the losses, L. Human-subject experiments suggest, however, that decision-makers may falsely utilize ratio considerations in the presence of (apparently) irrelevant information. For example, psychologists have found that investments in measures leading to savings of a fixed number of lives were preferred if the total number of individuals at risk was decreased [9]. Unfortunately, such a bias would lead to even less optimal decisions when considering the difference metric since the loss, L, is shown to be positively and linearly related to the price of uncertainty.

Of course, we should not forget that we consider a rather specialized environment, where only one single expert is alone in a population of naïve users. However stringent this assumption may sound, one should note that in reality, the number of expert users is dwarfed by the number of "lambda" users, that may not have the expertise, or inclination, to act very strategically.

Regardless of these limitations, we hope that our work will be a useful starting point for a serious discussion of information metrics applied to interdependent security scenarios. As we have shown here, picking the right metric is not a straightforward choice, and several pitfalls exist.

Acknowledgements We thank John Chuang, the anonymous reviewers and the participants of the Workshop on the Economics of Information Security for their helpful comments to an earlier version of this paper. All remaining errors are our own. This work is supported in part by the National Science Foundation under ITR award ANI-0331659 (100x100), with a University of California MICRO project grant in collaboration with DoCoMo USA Labs, and by CyLab at Carnegie Mellon under grant DAAD19-02-1-0389 from the Army Research Office. The views and conclusions contained here are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either express or implied, of ARO, CMU, or the U.S. Government or any of its agencies. We further want to thank Christos Ioannidis, Tyler Moore and David Pym for assembling this edited volume.

References

1. Acquisti, A., Grossklags, J.: Privacy and rationality in individual decision making. IEEE Security & Privacy **3**(1), 26–33 (2005)
2. August, T., Tunca, T.: Network software security and user incentives. Management Science **52**(11), 1703–1720 (2006)
3. Balcan, M., Blum, A., Mansour, Y.: The price of uncertainty. In: Proceedings of the ACM Conference on Electronic Commerce (EC), pp. 285–294. ACM Press, New York (2009)
4. Böhme, R., Nowey, T.: Economic security metrics. In: I. Eusgeld, F. Freiling, R. Reussner (eds.) Dependability Metrics, *LNCS*, vol. 4909, pp. 176–187. Springer, Berlin Heidelberg (2008)
5. Campbell, K., Gordon, L., Loeb, M., L. Zhou, L.: The economic cost of publicly announced information security breaches: Empirical evidence from the stock market. Journal of Computer Security **11**(3), 431–448 (2003)
6. Cavusoglu, H., Raghunathan, S., Yue, W.: Decision-theoretic and game-theoretic approaches to IT security investment. Journal of Management Information Systems **25**(2), 281–304 (2008)
7. Choi, J., Fershtman, C., Gandal, N.: Network security: Vulnerabilities and disclosure policy. Journal of Industrial Economics (forthcoming)
8. Dörner, D.: The Logic Of Failure: Recognizing And Avoiding Error In Complex Situations. Metropolitan Books (1996)
9. Fetherstonhaugh, D., Slovic, P., Johnson, S., Friedrich, J.: Insensitivity to the value of human life: A study of psychophysical numbing. Journal of Risk & Uncertainty **14**(3), 283–300 (1997)
10. Gal-Or, E., A. Ghose, A.: The economic incentives for sharing security information. Information Systems Research, **16**(2), 186–208 (2005)
11. Gordon, L., Loeb, M.: Managing Cyber-Security Resources: A Cost-Benefit Analysis. McGraw-Hill (2006)
12. Gordon, L.A., Loeb, M.: The economics of information security investment. ACM Transactions on Information and System Security **5**(4), 438–457 (2002)
13. Gordon, L.A., Loeb, M., Lucyshyn, W.: Sharing information on computer systems security: An economic analysis. Journal of Accounting and Public Policy, **22**(6), 461–485 (2003)
14. Granick, J.: Faking it: Calculating loss in computer crime sentencing. I/S: A Journal of Law and Policy for the Information Society **2**(2), 207–228 (2006)
15. Grossklags, J., Christin, N., Chuang, J.: Secure or insure? A game-theoretic analysis of information security games. In: Proceedings of the 17th International World Wide Web Conference (WWW), pp. 209–218. (2008)
16. Grossklags, J., Christin, N., Chuang, J.: Security and insurance management in networks with heterogeneous agents. In: Proceedings of the ACM Conference on Electronic Commerce (EC), pp. 160–169. ACM Press, New York (2008)
17. Grossklags, J., Johnson, B.: Uncertainty in the Weakest-link security game. In: Proceedings of GameNets, pp. 673-682. (2009)
18. Grossklags, J., Johnson, B., Christin, N.: When information improves information security. Tech. Rep. CMU-CyLab-09-004 (2009)
19. Grossklags, J., Johnson, B., Christin, N.: The price of uncertainty in security games. In: Proceedings of the 8th Workshop on the Economics of Information Security (WEIS). London, UK (2009)
20. Hershey, J., Baron, J.: Clinical reasoning and cognitive processes. Medical Decision Making **7**(4), 203–211 (1987)
21. Jaquith, A.: Security Metrics: Replacing Fear, Uncertainty, and Doubt. Pearson Education (2007)
22. Kabooza: Global backup survey: About backup habits, risk factors, worries and data loss of home PCs (2009). `http://www.kabooza.com/globalsurvey.html`
23. Kahneman, D., Tversky, A.: Choices, Values and Frames. Cambridge University Press (2000)

24. Koutsoupias, E., Papadimitriou, C.: Worst-case equilibria. In: Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science (STOC), pp. 404–413. ACM Press, New York (1999)
25. Kunreuther, H., Heal, G.: Interdependent security. Journal of Risk & Uncertainty **26**(2–3), 231–249 (2003)
26. Kwong, J., Wong, K.: The role of ratio differences in the framing of numerical information. International Journal of Research in Marketing **23**(4), 385–394 (2006)
27. Laffont, J.: The Economics of Uncertainty and Information. MIT Press (1989)
28. Liu, Y., Comaniciu, C., Man, H.: A Bayesian game approach for intrusion detection in wireless ad hoc networks. In: Proceedings of the Workshop on Game Theory for Communications and Networks, article no. 4. ACM Press, New York (2006)
29. Meier, D., Oswald, Y., Schmid, S., Wattenhofer, R.: On the windfall of friendship: Inoculation strategies on social networks. In: Proceedings of the ACM Conference on Electronic Commerce (EC), pp. 294–301. ACM Press, New York (2008)
30. Moscibroda, T., Schmid, S., Wattenhofer, R.: When selfish meets evil: Byzantine players in a virus inoculation game. In: Proceedings of the ACM Symposium on Principles of Distributed Computing (PODC), pp. 35–44. ACM Press, New York (2006)
31. NCSA/Symantec: Home user study (2008). `http://staysafeonline.org/`
32. Paruchuri, P., Pearce, J., Marecki, J., Tambe, M., Ordonez, F., Kraus, S.: Playing games for security: An efficient exact algorithm for solving Bayesian Stackelberg games. In: Proceedings of AAMAS, pp. 895–902. IFAAMAS, Richland, South Carolina (2008)
33. Quattrone, G., Tversky, A.: Contrasting rational and psychological analyses of political choice. The American Political Science Review **82**(3), 719–736 (1988)
34. Stanton, J., Stam, K., Mastrangelo, P., Jolton, J.: Analysis of end user security behaviors. Computers & Security **2**(24), 124–133 (2005)
35. Stone, E., Yates, F., Parker, A.: Risk communication: Absolute versus relative expressions of low-probability risks. Organizational Behavior & Human Decision Processes **3**(60), 387–408 (1994)
36. Swire, P.: A model for when disclosure helps security: What is different about computer and network security? Journal on Telecommunications and High Technology Law **3**(1), 163–208 (2004)
37. Swire, P.: No cop on the beat: Underenforcement in e-commerce and cybercrime. Journal on Telecommunications and High Technology Law **7**(1), 107–126 (2009)
38. Telang, R., Wattal, S.: An empirical analysis of the impact of software vulnerability announcements on firm stock price. IEEE Transactions on Software Engineering **33**(8), 544–557 (2007)
39. Varian, H.R.: System reliability and free riding. In: L.J. Camp and S. Lewis (eds.) Economics of Information Security, pp. 1–15. Kluwer Academic Publishers (2004)