Secure or Insure: An Economic Analysis of Security Interdependencies and Investment Types

by

Jens Grossklags

Diplom (Humboldt University of Berlin) 2001
M.I.M.S. (University of California, Berkeley) 2004
M.S. (University of California, Berkeley) 2008

A dissertation submitted in partial satisfaction of the
requirements for the degree of
Doctor of Philosophy

in

Information Management and Systems

in the

GRADUATE DIVISION
of the
UNIVERSITY OF CALIFORNIA, BERKELEY

Committee in charge:
Professor John Chuang, Chair
Professor Teck-Hua Ho
Professor Deirdre K. Mulligan
Professor Hal R. Varian

Fall 2009

Secure or Insure: An Economic Analysis of Security Interdependencies and Investment Types

Jens Grossklags

Abstract

Secure or Insure: An Economic Analysis of Security Interdependencies and Investment Types

by

Jens Grossklags

Doctor of Philosophy in Information Management and Systems

University of California, Berkeley

Professor John Chuang, Chair

Computer users express a strong desire to prevent attacks, and to reduce the losses from computer and information security breaches. However, despite the widespread availability of various technologies, actual investments in security remain highly variable across the Internet population. As a result, attacks such as distributed denial-of-service and spam distribution continue to spread unabated.

Users may struggle to respond vigorously because the effectiveness of security decisions is subject to strong interdependencies in a network, and different types of threats. In this dissertation, we address this complexity by analyzing investment decision-making in a unified framework of established games (i.e., weakest-link, best shot, and total effort) and novel games (e.g., weakest-target).

We examine how incentives shift between investment opportunities in a cooperative good (protection) and a private good (self-insurance), subject to factors such as network size, type of attack, loss probability, loss magnitude, and cost of technology. We capture security weaknesses due to monocultures by analyzing decision-making for an economy of homogeneous, selfish and fully rational agents under complete information. We compare our analysis to the case of heterogeneous users modeling efforts for security diversity. The findings highlight circumstances where poorly aligned incentives lead to security failures, and how interventions may be helpful.

Extending our analysis and relaxing assumptions on individuals' rationality, we consider the case of a single rational expert agent in an economy of nearsighted agents that under-appreciate the effect of security interdependencies. We further measure the value of information availability in the security context. Specifically, we introduce the *price of uncertainty* metric that quantifies the maximum discrepancy between the total expected payoffs for different information conditions. By evaluating the metric in different interdependency scenarios, we can determine which configurations can better accommodate limited information environments.

For my mother, father, and sister.

Contents

List of Figures v

List of Tables vii

1 Introduction 1
- 1.1 Technical security vulnerabilities . . . 3
- 1.2 Failures in infrastructure security investments . . . 5
- 1.3 Understanding individual security decisions . . . 6
- 1.4 Developing an economic framework for user security decisions . . . 8
- 1.5 Summary of contributions . . . 15
- 1.6 Roadmap . . . 19

2 Security monocultures: Homogeneous agents 21
- 2.1 Background . . . 22
  - 2.1.1 Security monocultures . . . 22
  - 2.1.2 The social organization of security: Interdependencies . . . 23
  - 2.1.3 Individual choice: Security as a hybrid good . . . 25
- 2.2 Basic model of security games . . . 27
- 2.3 Tightly coupled security interdependencies . . . 30
  - 2.3.1 Total effort security game . . . 30
  - 2.3.2 Weakest-link security game . . . 31
  - 2.3.3 Best shot security game . . . 33
- 2.4 Loosely coupled security interdependencies . . . 34
  - 2.4.1 Weakest-target security game (without mitigation) . . . 34
  - 2.4.2 Weakest-target security game (with mitigation) . . . 35
- 2.5 Nash equilibrium analysis . . . 36
  - 2.5.1 Total effort security game . . . 38
  - 2.5.2 Weakest-link security game . . . 42
  - 2.5.3 Best shot security game . . . 45
  - 2.5.4 Weakest-target security game (without mitigation) . . . 47

2.5.5 Weakest-target security game (with mitigation) 51
2.6 Identification of social optima 52
2.6.1 Total effort security game 53
2.6.2 Weakest-link security game 54
2.6.3 Best shot security game 55
2.6.4 Weakest-target security game (without mitigation) 57
2.6.5 Weakest-target security game (with mitigation) 59
2.7 Practical implications 60
2.8 Summary 63

3 Security diversity: Heterogeneous agents 65
3.1 Background: Heterogeneity in system security 66
3.2 Modification of the basic model 69
3.3 Nash equilibrium analysis 70
3.3.1 Total effort security game 70
3.3.2 Weakest-link security game 75
3.3.3 Best shot security game 77
3.3.4 Weakest-target security games 80
3.4 Intervention mechanisms 82
3.5 Summary 87

4 Bounded rationality and limited information 89
4.1 Background 92
4.1.1 Bounded rationality 93
4.1.2 Limited information 94
4.1.3 Heterogeneous agents 95
4.2 Decision-theoretic model 95
4.2.1 Modifications to the basic model 96
4.2.2 Player behavior 97
4.2.3 Information conditions 99
4.3 Analysis methodology 99
4.4 Results 106
4.4.1 Strategies and payoffs 106
4.4.2 Value of information 113
4.5 Summary 114

5 The price of uncertainty 120
5.1 Decision-theoretic model 124
5.2 Price of uncertainty metrics 125
5.2.1 Three metrics for the price of uncertainty 126
5.2.2 Discussion of the definitions 126
5.3 Analysis 128

5.3.1 Best shot game 128
5.3.2 Weakest-link game 141
5.3.3 Total effort game 150
5.4 Summary 155

6 Conclusions **159**
6.1 Contributions 159
6.2 Open questions 161

Bibliography **163**

A Derivations and tables for complete/incomplete information security game **196**
A.1 Derivations for weakest-link game 196
A.2 Derivations for best shot game 201
A.3 Derivations for total effort game 204
A.4 Tabulated results 214

List of Figures

1.1 Overview of security games. 11
1.2 Security games: Perfect defense and perfect attack. 12

2.1 Impact of network size, N, and loss magnitude, L, on threshold value for total effort self-insurance strategy. 41
2.2 Impact of loss magnitude, L, on threshold value for weakest-link protection strategy. 44

3.1 Reaction functions for a two-player total effort game. 70
3.2 Reaction functions for a two-player weakest-link game. 76
3.3 Reaction functions for a two-player best shot game. 78

4.1 Strategy boundaries in the incomplete information scenario for the sophisticated player. 111
4.2 Total expected payoffs for the strategic player under different information conditions, compared with that of the naïve agents. 118
4.3 The value of information for the three games. 119

5.1 Best shot – Difference metric: Maximizing b for $BPoU_1(L, N)$. 131
5.2 Best shot – Difference metric: $BPoU_1(L, N)$. 133
5.3 Best shot – Payoff-ratio metric: Maximizing b for $BPoU_2(L, N)$. 136
5.4 Best shot – Payoff-ratio metric: $BPoU_2(L, N)$. 137
5.5 Best shot – Cost-ratio metric: Maximizing b for $BPoU_3(L, N)$. 139
5.6 Best shot – Cost-ratio metric: $BPoU_3(L, N)$. 140
5.7 Weakest-Link – Difference metric: Maximizing b, c for $WPoU_1(L, N)$. . . 144
5.8 Weakest-Link – Difference metric: $WPoU_1(L, N)$. 144
5.9 Weakest-Link – Payoff-ratio metric: Maximizing b, c for $WPoU_2(L, N)$. . 146
5.10 Weakest-Link – Payoff-ratio metric: $WPoU_2(L, N)$. 147
5.11 Weakest-Link – Cost-ratio metric: Maximizing b, c for $WPoU_3(L, N)$. . . 148
5.12 Weakest-Link – Cost-ratio metric: $WPoU_3(L, N)$. 149
5.13 Total effort – Difference metric: $TPoU_1(L, N)$. 152

5.14 Total effort – Payoff-ratio metric: $TPoU_2(L, N)$. 153
5.15 Total effort – Cost-ratio metric: $TPoU_3(L, N)$. 154

List of Tables

A.1 Weakest-link security game: Payoffs for different strategies under different information conditions . . . 215
A.2 Weakest-link security game: Conditions to select protection, self-insurance or passivity strategies . . . 216
A.3 Weakest-link security game: Probabilities to select protection, self-insurance or passivity strategies . . . 217
A.4 Weakest-link security game: Total expected game payoffs, conditioned on other players . . . 218
A.5 Weakest-link security game: Total expected game payoffs, not conditioned on other players . . . 219
A.6 Best shot security game: Payoffs for different strategies under different information conditions . . . 220
A.7 Best shot security game: Conditions to select protection, self-insurance or passivity strategies . . . 221
A.8 Best shot security game: Probabilities to select protection, self-insurance or passivity strategies . . . 222
A.9 Best shot security game: Total expected game payoffs, conditioned on other players . . . 223
A.10 Best shot security game: Total expected game payoffs, not conditioned on other players . . . 224
A.11 Total effort security game: Payoffs for different strategies under different information conditions . . . 225
A.12 Total effort security game: Conditions to select protection, self-insurance or passivity strategies . . . 226
A.13 Total effort security game: Probabilities to select protection, self-insurance or passivity strategies . . . 227
A.14 Total Effort security game: Total expected game payoffs, conditioned on other players . . . 228
A.15 Total effort security game: Total expected game payoffs, not conditioned on other players . . . 229

Acknowledgments

This dissertation would not have been feasible without my advisors, colleagues, family and friends who have in various ways provided support, counseling and creative spirit. A short note in this section cannot convey the deep gratitude I have for all of you.

I want to thank my family and close friends for their their moral, and emotional support and limitless patience during my studies at Berkeley. Without you it would not have been possible.

I am highly appreciative of the advice and mentoring I received during the completion of my dissertation from the members of my qualifying and thesis committees, Michael Buckland, John Chuang, Teck-Hua Ho, Deirdre Mulligan and Hal Varian. I learned so much from you.

My advisor and dissertation chair, John Chuang, has been my primary source of support, technical expertise and guidance during the completion of my graduate degrees at UC Berkeley. His evenhanded approach, abilities to distill theory and technical concepts, and his generosity have been instrumental on my path from student to researcher. I am tremendously grateful for everything, John.

I am thankful for countless insights I garnered from my co-authors during our constructive arguments and the many hours completing working paper drafts. In particular, together with John Chuang, Nicolas Christin, and Benjamin Johnson, I worked on research that found its place in this dissertation. I also want to thank Alessandro Acquisti, Bettina Berendt, John Canny, Rachna Dhamija, Neal Fultz, Abhishek Ghose, Nathan Good, Oliver

Günther, Chris Hoofnagle, Joe Konstan, Carsten Schmidt and Sarah Spiekermann for our exciting and inspiring joint research projects and their advice. Thank you so much.

My thanks go to my fellow doctoral and masters students for many serious and funny conversations and forming such a wonderful community at the School of Information.

The following paragraphs acknowledge the work of individuals specific to each published chapter in this thesis.

- Chapter 2 is based on joint work with Nicolas Christin and John Chuang. The earlier work has been peer-reviewed by an academic program committee and published in proceedings of a conference organized by the International World Wide Web Conferences Steering Committee (IW3C2): Secure or insure? A game-theoretic analysis of information security games, in: *Proceedings of the 17th International World Wide Web Conference (WWW2008)*, ©2008 IW3C2 [94]. `http://doi.acm.org/10.1145/1367497.1367526`

- Chapter 3 includes co-authored work with Nicolas Christin and John Chuang. In particular, parts of this chapter are based on an earlier work that was competitively reviewed and published in proceedings of the Association for Computing Machinery (ACM): Security and insurance management in networks with heterogeneous agents, in: *Proceedings of the 9th ACM Conference on Electronic Commerce*, ©ACM, 2008 [95]. `http://doi.acm.org/10.1145/1386790.1386818`

  A summary paper of Chapters 2 and 3 was peer-reviewed and discussed at the *Seventh*

Workshop on the Economics of Information Security (WEIS), Dartmouth College, Hanover, New Hampshire, 2008.

- Chapter 4 draws from a research project with Benjamin Johnson and Nicolas Christin. This chapter is a substantially different and more comprehensive version of a prior work that was carefully peer-reviewed and appeared in proceedings of the Institute of Electrical and Electronics Engineers (IEEE): Uncertainty in the weakest-link security game, in: *Proceedings of the International Conference on Game Theory for Networks (GameNets '09)*, ©IEEE, 2009 [96]. Digital Object Identifier 10.1109/GAMENETS.2009.5137460

- Chapter 5 is the result of a combined research effort with Benjamin Johnson and Nicolas Christin. Our results have been thoroughly peer-reviewed and discussed at the *Eighth Workshop on the Economics of Information Security (WEIS)*, University College, London, England, 2009 [97].

I appreciate the insightful feedback received through the review process, and by the participants of the academic conferences. I also want to thank the organizations that supported my doctoral research.

Chapter 1

Introduction

"*One of the lamentable principles of human productivity is that it is easier to destroy than to create.*" Thomas C. Schelling (Arms and Influence [191])

The globally interconnected network serves as the basic infrastructure for countless aspects of the information society and substantially supports the worldwide economy and civil organization. The Internet has opened new and attractive channels to create, publicize and market products, to communicate with friends and colleagues, and to access information from spatially distributed resources. Though it has grown significantly, the network's architecture still reflects the cooperative spirit of its original designers [185]. Unfortunately, today's network users are no longer held together by that same sense of camaraderie and common purpose.

In fact, the expansion of the Internet has attracted individuals, groups and even orga-

nizations sponsored by nation states with often destructive and intrusive motivations [58]. During the early days of the Internet, the majority of these malefactors[1] were motivated by peer recognition, or curiosity, and were often undecided regarding the ethical legitimacy of their behavior, but actual damages were limited [89, 90, 119, 202].[2] However, with the increasing centrality of the Internet for the public and private sectors, Internet miscreants learned to conceive and implement techniques with the objective to exploit network stakeholders for their financial gain on a large-scale basis [1, 73, 123, 130, 159, 204].

The associated security attacks are common, widespread and highly damaging [166]. The "I Love You" virus [145], Code Red [158], Slammer [157], Storm [111], and Downadup worms [163] to cite some of the most famous cases, have infected hundreds of thousands of machines and installed code to steal personal and financial information, or to misappropriate resources for distributed denial of service attacks. Participants in underground markets are actively trading goods as diverse as individuals' banking credentials, or processing time on compromised resources [73].[3] Companies suffer from directed security attacks exposing their business secrets and customer data [180]. All together, these activities cost users and businesses billions of dollars in damages, while governments fear for public safety and security balance [19, 52, 208].

[1] We are avoiding the term "hacker" due to its semantic ambiguity including legitimate and illegitimate computer wizardry [110, 141].

[2] Already in the mid-nineties, a large share of businesses reported security breaches. For example, a CSI/FBI survey showed that 40 percent of the surveyed sites suffered at least one unauthorized access [51].

[3] The existence of active marketplaces evidences a high degree of specialization. However, some researchers also note significant inefficiencies due to fraud between criminals [102].

1.1 Technical security vulnerabilities

The surge of problems can be partly attributed to the fact that the global network infrastructure "is neither secure enough nor resilient enough" given current and future needs [100]. For example, security researchers have observed a multitude of problems associated with the basic routing infrastructure. Chakrabati and Manimaran distinguish between attacks on the Domain Name System (DNS)[4], routing table poisoning, packet mistreatment, and denial-of-service efforts [43]. Butler *et al.* provide an in-detail analysis of security threats challenging the dominant protocol for interdomain routing (i.e., the Border Gateway Protocol (BGP)) [39]. The limited technical guarantees provided by the protocol can cause instability and outages that, although frequently limited in impact and scope, may result in crippling and widespread harm (see, for example, research on prefix hijacking [30]).

Further, communications and information security protocols and algorithms are continuously probed. For example, Borisov *et al.* report ways to break the still widely used Wired Equivalent Privacy (WEP) protocol for wireless communications [31]. Given its popularity, it is surprising that the weakness of the protocol stems from failures in the application of cryptographic primitives rather than, for instance, side-channel attacks on the underlying hardware, implementation or configuration errors. In general, most security protocols can be challenged with brute force attacks, i.e., systematically probing with a large number

[4]The DNS is a global, distributed, and hierarchical directory with the primary purpose of translating domain names to numeric IP addresses.

of attempts. However, these strategies are usually computationally expensive. For example, in theory it would require 2^{80} attempts to find a message that would hash to the same value given the algorithm for the SHA-1 cryptographic hash function.[5] However, groups of researchers developed sophisticated collision strategies to reduce the required number of attempts to 2^{69} computations and later to 2^{63} [214], and eventually to 2^{52} [151].

Vulnerabilities in desktop and networking software are at the heart of many harmful security incidents. Weaknesses can be introduced at any stage during the software product life cycle, including specification, design, implementation, deployment and maintenance [136, 170]. The prevention of vulnerabilities is complicated because they can manifest themselves in a myriad of ways, such as in the form of an unchecked buffer or a race condition. Further, the requirements for modern software systems frequently lead to large increases in code volume and, therefore, code complexity.[6]

Researchers and industry have responded to these key weaknesses by developing numerous security technologies to alleviate many of the aforementioned problems [9].[7] However, many security compromises could potentially be prevented with more diligent adoption of improved security software, patches and protocols.

[5]Cryptographic hash functions should create a unique relationship between an input message and a hash value. They are used in many security applications, e.g., the protection of the integrity of a data repository [172].

[6]The measurement of software complexity is not straightforward, and many mathematical formulations have been proposed (e.g., [68,149]) and evaluated (e.g., [18,218]. See also the literature on software reliability [24,33].

[7]Several other technical challenges exist such as overcoming insecurities associated with hardware devices. Examples include challenges to card readers for online banking [61], appliances using radio frequency identification [120], smart cards [193], automatic teller machines [109], electronic voting machines [67] etc.

1.2 Failures in infrastructure security investments

All stakeholders of Internet communications experience the negative consequences of security incidents. For instance, service providers suffer from the additional cost of abuse notifications and management, and software vendors and content providers endure brand tarnishing, patching and notification costs [85, 166, 182]. Unfortunately, the deployment of promising countermeasures is frustratingly slow and hampered by cost considerations, misaligned incentives and coordination problems [10].

A primary reason why security technologies are not adopted or upgraded is the *direct economic cost*. For example, the introduction of protocols using cryptographic primitives may necessitate additional hardware investments. A frequently discussed case is the lack of widespread deployment of Internet Protocol Security (IPsec) used to authenticate and encrypt Internet Protocol (IP) packets. Miltchev *et al.* conduct a protocol benchmark analysis of IPsec. They acknowledge that the benefits of IPsec are immediately obvious. But the advantages must be weighted with the need for technology purchases, e.g., hardware cryptographic accelerators [155].

Further significant hurdles for deployment arise due to the various *interdependencies* and the associated positive and negative externalities between the different stakeholders of Internet communications [49]. A key entity are Internet Service Providers who are generally (technically) capable of undertaking changes from the physical infrastructure level up to the application layer, but only *within their domains*. The global network infrastructure

consists of more than 20000 globally accessible autonomous systems [59].[8] And, typically, a service provider does not have purview and control over an entire end-to-end path [65]. Accordingly, the benefit that providers can derive from a deployment of new technology may depend on the number of other entities taking the same measure. Depending on the type of improvement, a unilateral enhancement might yield the desired benefit. At the other extreme, complete cooperation of all network operators may be required.[9] The secure shell (SSH) protocol is an example for a communication primitive where even a small group of adopters - like a single organization - can reap immediate benefits [181]. In contrast, the Domain Name System Security Extensions (DNSSEC) necessitate widespread adoption to improve network security [167].

1.3 Understanding individual security decisions

Consumers and small businesses are caught between a rock and a hard place. Some reports declare they are a significant factor for the prevailing security problems [77].[10] However, as a group they bear the brunt of the attacks and cannot rely on protection from

[8] An autonomous system is a group of computer networks using the same routing policy that are, therefore, commonly under the same administrative authority [28]. From 1997 to 2005, researchers observed an increase of the number of globally routable autonomous system identifiers from less than 2000 to more than 20000 [59]. Dimitropoulos *et al.* provide a classification based on size and ownership: large ISPs (44), small ISPs (5599), customer autonomous systems (11729), universities (877), Internet exchange points (33), and Network Information Centers (332) [59].

[9] Network virtualization has been proposed to allow providers to run customized networks in a parallel fashion over a shared infrastructure (e.g., [65]). However, Laskowski and Chuang point at the lack of adoption incentives for the virtualization technologies themselves [137].

[10] For example, in the context of identity theft one can find commentary stating that "consumers are to blame for many identity theft incidents, because they fall for phishing attacks, they fail to secure personal information, or they allow family members or friends to steal their identity" (collected by Hoofnagle [113]).

their service and content providers [11]. Further, consumers are severely limited in finding redress with the help of law enforcement or government agencies (e.g., the Federal Trade Commission).[11]

When asked in surveys, network users say they are interested in preventing attacks and mitigating the damages from computer and information security breaches [3, 36, 121, 200]. And consumers and small businesses are certainly taking some precautionary measures. But the evidence is mixed. A 2008 home user study provided adoption data for important security software. The report evidenced that some security software solutions are quite common, e.g., anti-virus software (95% installation rate), anti-spyware applications (82%). Other basic safety solutions are notably absent, i.e., firewall technologies (58%), anti-spam (42%) and anti-phishing (50%) [162].[12] Another survey presents discouraging data about the usage of information security practices. For example, at least 67 percent of the surveyed home users never used email encryption, 82 percent never utilized credit alerts, and 83 percent never removed their private telephone numbers from public directories [3].[13]

Further, the speed of attack innovation will render even up-to-date prevention measures sometimes powerless, and users are therefore urged to invest in mitigation and backup

[11]Consider the following non-exhaustive list of obstacles. First, while the number of identity theft incidents grows faster than other types of theft, the clearance rate is falling behind [6]. Second, transborder communication and cooperation difficulties (both national and international) impede law enforcement effectiveness [207]. Third, the Federal Trade Commission is tasked to prevent fraudulent, deceptive, and unfair business practices in the marketplace and to provide information to help consumers, however, it usually does not resolve individual consumer complaints [66]. Fourth, consumers will not always be able to prove minimum damages as required by legal statutes. For example, the Computer Fraud and Abuse Act requires a "loss to one or more others of a value aggregating $1,000 or more during any one year period (18 USC 1030)."

[12]For the report 400 personal computers located in the United States were scanned [162]. See also earlier reports conducted by the National Cyber Security Alliance, e.g., [12, 13, 148].

[13]The survey study was conducted online and included 119 responses [3].

technologies. However, a 2001 survey found that about 25 percent lost data to security incidents as well as hardware and software faults. Further, only 41 percent personally conducted data backups and 69 percent did not recently create a copy of their data [36].[14] In 2009, an international survey found that 66 percent have lost files (with 42 percent within the last 12 months). The survey also noted a low usage rate of backup technologies of less than 50 percent [121].[15]

The empirical evidence shows that users and small businesses are concerned about security, but they follow highly different security strategies.[16] In the following, we ask whether these observations can be explained with economic considerations. Alternatively, do users invest too little, or too much, or are we witnessing the results of naïve consumer decision-making [21]?

1.4 Developing an economic framework for user security decisions

In this dissertation, we take a theoretical approach to analyze individuals' security incentives. We are thereby building on prior work from risk management, public goods economics, and the relatively novel research area on the economics of computer and network security. We develop a mathematical framework that addresses important aspects of

[14] About 1000 computer users were surveyed for the report [36].
[15] 4257 consumers from 129 countries were included in the survey study [121].
[16] Further support is provided by ethnographic studies in workplace environments [60].

security investment decisions. In the following, we are highlighting several focus areas of our models that we will address more thoroughly in the remainder of this work. We begin with the modeling of interdependencies, and the consideration of preventive and mitigating security investment types. We continue with a discussion of rationality assumptions that shape individuals' reactions to threats. Finally, we discuss different assumptions about the amount of information users have at their disposal about security threats and consequences.

Interdependencies: While many models prescribe behavior in individual choice situations, the focus of our work is to model and study strategic interaction in networked systems to understand the impact of individual choices within a larger group. Such interactions usually involve common as well as conflicting interests [190].[17]

A further concern is the large and growing number of security threats, vulnerabilities and implications, which pose significant challenges for end users. We argue that by considering a finite number of *canonical cases* we can highlight important differences in the incentive structure for security investments. The analysis explains security failures in practical settings and helps to evaluate intervention mechanisms and countermeasures that respect economic considerations.

- First, in the *perimeter defense* scenario a breach at a single point will leave the whole of a network unprotected and open to harm. This scenario is a major cause of concern for businesses that have to trust all employees to adhere to security guidelines and practices [81]. Similarly, this weakness affects any group of users that relies

[17]This mutual dependence as well as opposition guarantees for a more realistic scenario for analysis. Pure conflict, in which the interests of all agents are completely opposed, is a special case [190].

on the confidentiality of a shared secret (e.g., access passwords, membership in an organization, customer data, business secrets) [53, 153].

- Second, in the *cumulative defense* scenario the probability of harm and/or the magnitude of potential damages incrementally increases with the number of users not taking security measures. For example, the volume of spam amplifies with the share of consumers who open or conduct purchases based on unsolicited communications [144]. In a like manner, the threat of distributed denial of service attacks becomes more significant with the number of compromised end-user resources [178].

- Third, with the *last stand defense* we address situations in which a protective effort's success critically depends on the survival of at least a single part of the network (e.g., a user's computer with a valuable document). Primary examples are the threat of censorship whether attempted by technical or legal means [8, 64], and the protection of decentralized storage systems in peer-to-peer networks [183] and wireless sensor networks [80].

- Fourth, the effectiveness of a *comparative defense* relies on the proposition that other entities are less well protected against security threats or more attractive to an attacker.[18] Many cybercrime activities rely on low costs, for example, when sending unsolicited bulk email or reselling stolen credentials at underground marketplaces [50, 70].

[18] We consider attackers that search for the least protected defenders. More generally, malefactors might react to preventive investments by considering, for example, alternate times, places, methods, or completely different offenses [105].

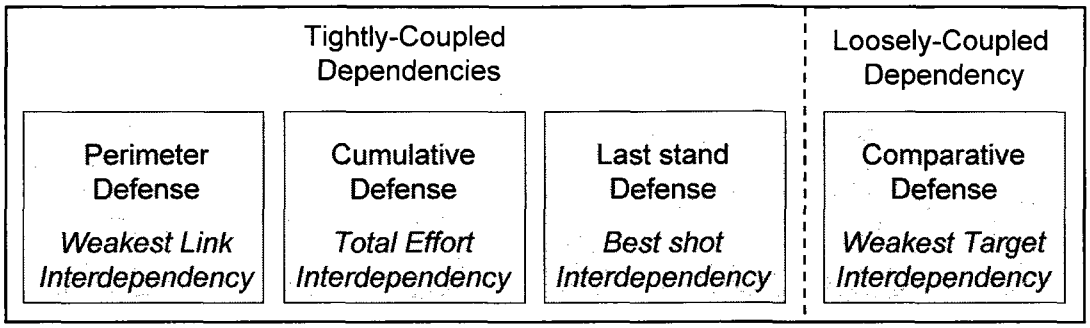

Figure 1.1: **Overview of security games.**

To analyze these canonical cases, we build upon public goods literature [107, 212] by mapping to the economic games depicted in Figure 1.1. We consider the classical weakest-link, total effort and best shot games and analyze them in a security context.[19] We complement these three games with a novel model, called the *weakest-target game* to capture the comparative defense scenario.

On a high level, we make a distinction between tightly-coupled and loosely-coupled interdependencies [75, 169]. In a tightly-coupled network, all defenders will face a loss if the condition of a security breach is fulfilled. This description applies to the public goods interdependencies. In a loosely-coupled network consequences may differ for network participants. Particularly lucrative or unprotected targets receive preference from an attacker while other defenders are not attacked and remain unharmed. We capture this type of interdependency with the weakest-target game.

A point of contention for practitioners is the feasibility of an infallible defense in the

[19] We will discuss our approach more closely in Chapter 2.

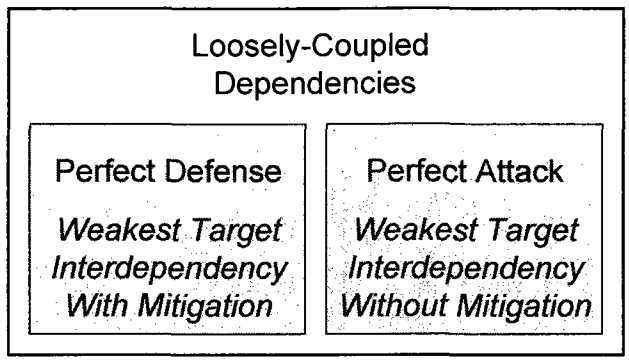

Figure 1.2: **Security games: Perfect defense and perfect attack.**

context of computer security [186]. We contribute to this debate by proposing two variations of the weakest-target game (see Figure 1.2). On the one hand, in the case with mitigation defenders can invest in safe protection with the guarantee of evading security compromises. On the other hand, the weakest-target game without mitigation allows an attacker to overcome the defense of even well-protected entities.[20]

Diversity of security actions: Past research on the economics of computer security focused on security investments as a problem with a single variable (i.e., amount of money spent on preventive security). In our work, we consider two key components to be part of a comprehensive security strategy. First, individuals can invest in *self-protection.* Protection efforts include the patching of system vulnerabilities, investments in firewall and intrusion detection systems, and the adoption of scanning software against viruses, spyware, spam, and other forms of malicious code. We also consider security-conscious behaviors such as the immediate deletion of suspect email in this category. Second, individuals can select *self-*

[20]This distinction could also be made for the public goods games. We defer this extension to future work.

insurance to reduce the damages from security breaches. The most important mitigating measure are comprehensive backups.[21]

Rationality assumptions: The complexity of network security poses immense requirements on the rationality of decision-makers. In order to implement optimal security investments, individuals need to comprehend the impact of interdependencies and the trade-off between different security actions. Further, agents must collect essential data, be able to compute strategies, and practically execute their plans without mistakes [3, 194]. These challenges cause computer users to apply qualitative evaluations and aspirational solutions to security problems. This leaves the potential for weaknesses in their defenses [83, 84].[22] In prior work, we addressed individuals' innate *bounded rationality* by developing a model of near-rational decision-making in networked systems [47].[23]

In this dissertation, we are considering fully rational as well as bounded rational decision-makers. In the latter case, we are interested in the structural understanding that average computer users have of system interdependencies. We anticipate the vast majority of users to be *non-expert*, and to apply approximate decision-rules that fail to accurately appreciate the impact of their security decisions on others. In particular, we assume non-expert users to conduct a simple self-centered cost-benefit analysis, and to neglect interdependencies. Such users would secure their system only if the vulnerabilities being exploited can cause

[21] In the market insurance context, the importance of self-protection and self-insurance was first recognized by [62].

[22] We have also reviewed research at the intersection of computer science and behavioral experimentation [93].

[23] In particular, we applied the notion of the ϵ-equilibrium concept [5, 177]. Each individual is satisfied to get within a certain bound of the optimal payoff that would result from her best response to others' strategies [47].

significant harm or a direct annoyance to them (e.g., their machines become completely unusable), but would not act when they cannot perceive or understand the effects of their insecure behavior (e.g., when their machine is used as a relay to send moderate amounts of spam to third parties). In contrast, an expert user fully comprehends to which extent her and others' security choices affect the network as a whole, and responds rationally.

Information: Lack of information can significantly affect security-decision making and economic outcomes. In this dissertation, we are considering static games with *complete and incomplete information.* In particular, we address how users' security choices are mediated by the information available on the severity of the threats the network faces. At first, we develop models under the assumption that agents are fully aware of their own and all others' parameters (including security costs, and expected damages), strategies and resulting payoffs. Then, we relax this informational assumption. In practice, different targets, even if they are part of a same network, are not all equally attractive to an attacker: a computer containing payroll information is, for instance, considerably more valuable than an legacy system that holds historical data. In our incomplete information model, a user is aware of the potential damage that would result from a security breach (i.e., she is holding private information). However, individuals do not know the precise harm that successful attacks will inflict on their peers, and can only form an expectation on their respective willingness to invest in protection.

1.5 Summary of contributions

We propose and analyze models on network and computer security decision-making. We aim to better understand investment decisions for protection and mitigation in several canonical security scenarios. Further, we expect to gain insights regarding security failures which we observe in practical settings. We discuss different intervention techniques and propose economic metrics to guide system design decisions.

In our analysis, we take a systematic, step-by-step approach. 1) We start by discussing our basic model and conduct an economic analysis for symmetric static games with fully rational users in possession of full information. 2) We evaluate optimal security strategies from the perspective of a social planner that has complete control over users' security investments. 3) By considering heterogeneity in the user population we are able to draw comparisons between uniform and diverse networks. 4) We relax assumptions about user rationality and information conditions to address concerns about practical challenges of security decision-making. 5) We develop metrics to evaluate the economic importance of complete information in comparison to restricted information about the impact of security attacks. We structure and summarize our contributions in the following section.

The consideration of different security scenarios: As a first contribution, we present and discuss canonical types of security interdependencies. We focus in our interpretation on security challenges faced by end users and small businesses. We mathematically embed the interdependencies in five static economic games with complete information and multiple agents. We distinguish between scenarios in which agents share the consequences

of security successes and failures (i.e., tightly-coupled games), and loosely-coupled games allowing users to differentiate themselves from their peers.

Our analysis of tightly-coupled interdependencies relies on prior work in public goods economics. We study the total effort, weakest-link, and best shot game in the security context. Further, we develop two variants of the novel weakest-target game to address loosely-coupled dependencies. In this game, attackers will successfully compromise users with the lowest security settings.[24]

The set of security games allows us to contrast and compare economic incentives across important security scenarios. In practice, a particular type of threat will be dominant for a group of defenders. For example, a group of dissidents or political activists wants to prevent censorship of their documents (i.e., the best shot game) [78]. Similarly, savvy consumers want to avoid being targeted by financial fraud or network threats by investing in less frequently targeted (and potentially) safer technologies (i.e., the weakest-target game).

The temptation to deviate from protection in the presence of self-insurance options: We provide a model that allows a decoupling of investments in the context of computer security. On the one hand, the defense 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).

We study the strategic interactions of users by applying the Nash equilibrium solution concept. We examine how incentives shift between the two types of investment opportuni-

[24]In the weakest-target game with mitigation there is a safe investment level.

ties, subject to factors such as network size, type of attack, loss probability, loss magnitude, and cost of technology.

We characterize for each game equilibria for protection, self-insurance and passivity. We find that for the total effort game and the weakest-link game multiple equilibria exist for the same parameter values. This multiplicity invokes complicated coordination problems. As a result, the availability of self-insurance technologies may weaken the willingness to protect when agents fear lack of others' cooperation.

The role of a social planner in security games: Our analytic results for fully decentralized decision-making highlight the benefits of an intermediary with the capability to enforce superior protection and self-insurance strategies. For example, in the total effort game, individuals will frequently select protection efforts that are below the social optimum (e.g., they may fail to coordinate on a protection strategy). We determine and discuss the strategies that maximize overall utility for all users in a network.

In games with tightly-coupled interdependencies, we observe that a central planner may increase the average protection level of the network. However, we found that the common wisdom that having a central planner who decides upon security implementation always yields higher protection contributions by individual players does not hold (i.e., in the weakest-target game).

The trade-off between monocultures and diversity: Some security researchers warn that our vulnerability to security threats is exacerbated by dominant products in the software and hardware segments. As a remedy the injection of heterogeneity has been proposed. We

evaluate the economic impact of the introduction of diversity by studying the incentives of a heterogeneous population in networks.

We find that the robustness of protection outcomes in the presence of user heterogeneity depends on the type of interdependency. For example, in the weakest-link game the injection of heterogeneity may likely lead some individuals to prefer self-insurance or passivity in comparison to protection investments. As a result, the willingness to protect will unravel in the population. In contrast, in the best shot game heterogeneity can serve as a coordination tool.

A model with bounded rational agents and incomplete information: In this thesis, we develop a decision-theoretic model with a mixed population of expert and inexperienced users. Savvy users understand the implications of interdependencies. However, non-expert users are nearsighted and only perceive direct security threats on their resources, and neglect the impact of their decisions on other agents.

We study the strategic optimization behavior of such a rational expert user in an economy of naïve users. We also address how the security choices by users are mediated by the information available on the severity of the threats the network faces.

Naïve users experience a payoff reduction as a result of their limited understanding of correlated threats, whereas experts may suffer from the impact of limited information.

The development of metrics for measuring the value of information: We further ask how much defenders can gain by investing in techniques or other efforts to improve information availability about attack threats and other users' incentives. Our contribution is to pro-

pose and evaluate a set of generic metrics that are applicable to different security decision-making situations to help with this trade-off calculation.

Specifically, we introduce the *price of uncertainty* metric that quantifies the maximum discrepancy in the total expected payoff between different information conditions. We consider difference, payoff-ratio, and cost-ratio sub-metrics as canonical nontrivial measurements of the price of uncertainty.

By evaluating the metrics for a range of parameters in the different security games, we can determine which configurations can accommodate limited information environments (i.e., when being less informed does not significantly jeopardize a rational user's payoff). We expect these results to be of relevance for systems designers and the economic metrics community.

1.6 Roadmap

In the following, we present the structure of this dissertation:

- In Chapter 2 we provide an overview of related work including economic literature on security interdependencies and on the distinction between prevention and mitigation. Then we introduce the framework for security games with multiple agents. We analytically study individually rational decision-making of homogeneous agents in tightly and loosely coupled games. We derive Nash equilibrium strategies and determine social optima.

- Chapter 3 includes our discussion of heterogeneity in the context of security. We study how equilibrium predictions are impacted when agents face different losses, and dissimilar costs of security.

- We modify our framework to account for bounded rationality of agents, and different information conditions in Chapter 4. Agents differ in how they understand the interrelatedness of security decisions. Information availability is restricted about the expected damages that other agents face.

- In Chapter 5 we develop metrics to quantify the value of complete information in comparison to incomplete information (using the model from Chapter 4). We study this price of uncertainty analytically and graphically.

- We summarize contributions and discuss implications of our results in Chapter 6. We conclude with a discussion of opportunities for future work.

Chapter 2

Security monocultures: Homogeneous agents

In this chapter, building upon the literature on public goods [107, 212], we consider the classical best shot, total effort, and weakest-link games, and analyze them in a security context. We complement these three games with a novel model, the weakest-target game, which allows us to describe a whole class of attacks ranging from insider threats to very aggressive worms. Furthermore, while most research on the economics of security focuses on security investments as a problem with a single variable (e.g., amount of money spent on security), we propose to decouple protection investments (e.g., setting up a firewall) from self-insurance coverage (e.g., archiving data as backup). This separation allows us to explain a number of inefficiencies in the observed user behaviors.

This chapter is only a first step toward a more comprehensive modeling of user atti-

tudes toward security issues. We analyze security decision-making for a number of homogeneous, selfish and fully rational agents under complete information. We will relax these assumptions in the later chapters of this dissertation.

The rest of this chapter is organized as follows. We discuss the background of our model and related work in Section 2.1. We introduce our game-theoretic models in Section 2.2. Then we present an analysis of the Nash equilibria (Section 2.5) and social optima (Section 2.6) for all games. We discuss and summarize our results in Sections 2.7 and 2.8, respectively.

2.1 Background

2.1.1 Security monocultures

Some security researchers warn that our vulnerability to malware (malicious software) attacks is exacerbated by monoculture: the predominance of a single operating system or software application. The debate was sparked by a Computer and Communications Industry Association report [79], which argued that Microsoft's dominance in the market for operating systems represents a grave threat to Internet safety and to national security. Reports by other industry research groups echo these concerns by recommending that companies reduce their reliance on a single operating system in order to avoid damage caused by malware attacks [213].

In particular, users of highly popular software products may suffer from a *monoculture*

penalty. The more common a product is, the more attractive it appears to malefactors, and the more likely this product is targeted by malware attacks. Users are also more often subject to propagated threats that utilize the similarities between individual computing devices to more easily spread across the network.

In contrast, monoculture has also direct security benefits. Uniform systems are easier to update and patch. Less complexity usually equals better understanding, and may therefore increase the incentives to contribute to system-wide security [23].

Our basic model captures decision-making in a homogeneous environment to address the impact of monocultures for network security in the presence of different types of dependencies. Agents share the same cost for security, as well, as identical losses. In the following, we describe the security games framework in more detail.

2.1.2 The social organization of security: Interdependencies

The prevalence of widely spread, propagated and correlated threats such as distributed denial of service attacks (DDoS), worms and spam has brought attention to interdependencies existing in computer networks. For an attacker this might create strong economies but sometimes also diseconomies of scale. For example, a single breach of a corporate perimeter may allow an attacker to harvest resources from all machines located within its borders. In other scenarios an attacker may have to shut down every single computer or network connection to achieve an attack goal and thereby incur large costs potentially proportional to network size. More generally, there is an interaction between the structure of

the defenders' network, the attack goal and threat model.

To better understand the implications of this mutual dependence, Varian [212] conducts an analysis of system reliability within a public goods game-theoretical framework. He discusses the best effort, weakest-link and total effort games, as originally analyzed by Hirshleifer [107, 108]. The main difference from classical public goods theory is that within the framework of computer reliability "considerations of costs, benefits, and probability of failure become paramount, with income effects being a secondary concern." [212] Varian focuses on two-player games with heterogeneous effort costs and benefits from reliability.[1] He also adds an inquiry into the role of taxes and fines, and differences between simultaneous and sequential moves.

We distinguish between tightly and loosely-coupled networks [169]. In a tightly coupled network all defenders will face a loss if the condition of a security breach is fulfilled. This may be a suitable description, for example, of a network perimeter breach that causes the spread of malicious code to all machines, but also applies to independently acting defenders that try to preserve a common secret or resist censorship. In a loosely coupled network consequences may differ for network participants. For example, an attacker might be interested to gain control over a limited set of compromised machines ("zombies" or "bots") and to organize them into a logical network ("botnet") with the goal of executing a DDoS attack against third parties or sending unsolicited information to and from the bots (i.e., popup advertisements and spam). At other times, an attacker might target a specific

[1] A distinction between reliability and security, in terms of consequences, may exist [112]. In this study, we do not follow this distinction and consider reliability as a key component of security.

set of users (e.g., wealthy users in spearphishing scams). Other users would stay unharmed and are never targeted.

There are several research papers that capture some degree of security interdependency between agents. Clark and Konrad analyze a variation of the weakest-link game in which one attacker is satisfied by breaching at least one point in the defense that is managed by a single player (who acts as a social planner for multiple defense points) [48]. Kunreuther and Heal derive rational strategies when defenders can protect themselves against directed attacks, but are helpless against propagated threats from their own peers [132].[2] August and Tunca study patching behavior in a continuum of users when lack of security updates causes negative network externalities [16]. Further, several research papers explore the optimal strategies of defenders and attackers in graph-theoretic network security games [15, 139, 147, 160].

Other works are concerned with how defenders' investment choices influence the number or identities of the specifically targeted individuals based on economic considerations (if the attacker has the capability to select) [22, 55, 189].

2.1.3 Individual choice: Security as a hybrid good

Our work generalizes the research by Varian [212] and others in several aspects. First, instead of considering security decisions to be determined by a single "security" variable, we identify two key components of a security strategy: self-protection (e.g., patching sys-

[2]This work has been applied to the airline security context [101, 125].

tem vulnerabilities) and self-insurance (e.g., having good backups). More precisely, we allow agents to self-protect and/or self-insure their resources in N-player games. We also contrast the three canonical games discussed by Varian with two more complex "weakest-target" games that represent a more complicated incentive structure, which we believe applies to a whole class of security issues.

Outside the network security context, the dual role of self-protection and self-insurance was first recognized by [62]. To provide a more precise definition, self-protection stands for the ability to reduce the probability of a loss – for example, by installing a firewall application which limits the amount of traffic allowed to communicate with one's network. Self-insurance, on the other hand, denotes a reduction in the magnitude of a loss, e.g., by performing regular backups on existing data. Some technologies and practices such as disconnecting a computer from a network do both. Ehrlich and Becker [62] focus in their analysis on the comparison of self-protection and self-insurance to market insurance. They find that, for rare loss events, there is less incentive to self-insure losses than to use market insurance. This is due to their assumption, that the price of self-insurance is independent of the probability of the loss. An additional result is that the demand for self-insurance grows with the base loss of a security threat. As an outcome of their work, they characterize self-insurance and market insurance as substitutes, and self-protection and market insurance as complements. Our analysis complements the work in [62] by extending the concepts of self-protection and self-insurance to the public goods and security context.

Several other researchers have included in their analyses the trade-off between differ-

ent security measures [32, 44]. Konrad and Skaperdas consider self-insurance and self-protection decisions by individuals with rank-dependent expected utility preferences [128]. Briys *et al.* evaluate the benefits of self-protection and self-insurance when the associated technologies are not always reliable [35]. Another model addresses how a firm can optimally respond to hazards in the production process by investing in mitigating or preventive technologies [106].

Our research complements work on market insurance for security and privacy [7, 219]. Cyberinsurance can fulfill several critical functions. For example, audit requirements for cyberinsurance can motivate investments in security, and might contribute to a better understanding of the economic value of the protected resources [127]. Several researchers have investigated the impact of correlation of risks and interdependency of agents in networks on the viability of insurance [26, 29].

2.2 Basic model of security games

We define a security game as a game-theoretic model that captures essential characteristics of decision making to protect and self-insure resources within a network. Varian [212] observed that frequently the success of security (or reliability) decision making depends on a joint protection level determined by all participants of a network. The computation of the protection level will often take the form of a public goods contribution function with nonrival and nonexcludable benefits or consequences. A main observation is that dependent on

the contribution function individuals may be able to freeride on others' efforts. However, individuals may also suffer from inadequate protection efforts by other members if those have a decisive impact on the overall protection level.

Following Varian's exposition, we analyze three canonical contribution functions that determine a global protection level. Different from Varian's work however, here network members have a second action available: They can decide to self-insure themselves from harm. The success of insurance decisions is completely independent of protection choices made by the individual and others. Consequently, the games we consider share qualities of private (on the self-insurance side) and public (on the protection side) goods. We further add to the research literature by studying two additional games with a more complex determination of protection levels.

Security games share the following key assumptions: (i) all entities in the network share a single purely public protection output, (ii) a single individual decides on protection efforts for each entity (so we do not assume a second layer of organizational decision making), (iii) protection costs per unit are identical for each entity, and (iv) all decisions are made simultaneously. These assumptions are commonly made also in models on decision making of partners in military alliances [187]. We add to these main assumptions that individuals are able to self-insure resources at a homogeneous cost with self-insurance being a purely private good.

Formally, the basic model from which we develop the security games has the following payoff structure. Each of $N \in \mathbb{N}$ players receives an endowment M. If she is attacked and

compromised successfully she faces a loss L. Attacks arrive with an exogenous probability of p ($0 \leq p \leq 1$). Players have two security actions at their disposition. Player i chooses an insurance level $0 \leq s_i \leq 1$ and a protection level $0 \leq e_i \leq 1$. Finally, $b \geq 0$ and $c \geq 0$ denote the unit cost of protection and insurance, respectively. The generic utility function has the following structure:

$$U_i = M - pL(1 - s_i)(1 - H(e_i, e_{-i})) - be_i - cs_i \,, \tag{2.1}$$

where following usual game-theoretic notation, e_{-i} denotes the set of protection levels chosen by players other than i. H is a "contribution" function that characterizes the effect of e_i on U_i, subject to the protection levels chosen (contributed) by *all* other players. We require that H be defined for all values over $(0, 1)^N$. However, we do not place, for now, any further restrictions on the contribution function (e.g., continuity). From Eqn. (2.1), the magnitude of a loss depends on three factors: i) whether an attack takes place (p), ii) whether the individual invested in self-insurance ($1 - s_i$), and iii) the magnitude of the joint protection level ($1 - H(e_i, e_{-i})$). Self-insurance always lowers the loss that an individual incurs when compromised by an attack. Protection probabilistically determines whether an attack is successful. Eqn. (2.1) therefore yields an expected utility.

We introduce five games in the following discussion. In selecting and modeling these games we paid attention to comparability of our security games to prior research (e.g., [107, 187, 212]). The first three specifications for H represent important baseline cases recognized in the public goods literature. To allow us to cover most security dilemmas, we

add two novel games, for which we could not find a formal representation in the literature. All games are easy to interpret within and outside the online security context.

2.3 Tightly coupled security interdependencies

In a tightly coupled network all defenders will face a loss if the condition of a security breach is fulfilled [75]. In the following, we discuss canonical examples of such interdependencies and their interpretation in the security context.

2.3.1 Total effort security game

In the total effort security game the global protection level of the network depends on the sum of contributions normalized over the number of all participants. That is, we define $H(e_i, e_{-i}) = \frac{1}{N}\sum_i e_i$, so that Eqn. (2.1) becomes

$$U_i = M - pL(1 - s_i)(1 - \frac{1}{N}\sum_k e_k) - be_i - cs_i \,. \tag{2.2}$$

Economists identified the sum of efforts (or total effort) contribution function long before the remaining cases included in this section [107]. We consider a slight variation of this game to normalize it to the desired parameter range. A typical parable for the total sum function is that the effectiveness of a dam or city wall depends on its strength that is contributed to by all players. In terms of security the average contributions matters if an attacker wants to successfully conquer the majority of machines in a network one-by-one. For instance, consider a building plan for a new technology that is spread across a com-

pany's network and which is considerably more valuable to an attacker, if obtained in its entirety.

As another example, maybe more related to Internet security, consider parallelized file transfers, as in the BitTorrent peer-to-peer service. It may be the case that an attacker wants to slow down transfer of a given piece of information; but the transfer speed itself is a function of the aggregate effort of the machines participating in the transfer. Note that, the attacker in that case is merely trying to slow down a transfer, and is not concerned with completely removing the piece of information from the network: censorship actually results in a different, "best shot" game, as we discuss later.

2.3.2 Weakest-link security game

The overall protection level depends on the minimum contribution offered over all entities. That is, we have $H(e_i, e_{-i}) = \min(e_i, e_{-i})$, and Eqn. (2.1) takes the form:

$$U_i = M - pL(1 - s_i)(1 - \min(e_i, e_{-i})) - be_i - cs_i \,. \tag{2.3}$$

This game describes the situation where a levee or city wall that is too low at any point leads to a negative payoff to all players in the event of a flood or attack. The weakest-link game is easily the most recognized public goods problem in computer security by business professionals and researchers alike.[3] In the weakest-link externality an attacker is able (after approaching her target) to identify the least protected point in the system

[3]See, for example, see a recent interview with a security company CEO. New York Times (September 12, 2007), "Who needs hackers," available at `http://www.nytimes.com/2007/09/12/technology/techspecial/12threat.html`. Stating that: "As computer networks are cobbled together [...] the Law of the Weakest Link *always* seems to prevail."

of interconnected resources in which the target is embedded. Depending on the type and security actions of defenders the weaknesses of a system can be costly to circumvent, and of surprising variety.

On the one hand, technology and code quality are often the culprits of (un)predictable weaknesses in the chain of defense. The increasing complexity of software products (for example, because of code bloat and feature creep) leaves little doubt that most publicly available software products include several significant security vulnerabilities [79]. But even sophisticated and thoroughly tested security software and protocols (e.g., certain hard disk encryption packages) can sometimes be broken with non-standard attacks [98], or large-scale brute-force efforts [126]. Legal, regulatory and law enforcement requirements can also put limits on security effectiveness (e.g., through mandatory escrow of encryption keys or inclusion of back doors in hardware and software technologies) [25].

On the other hand, many observers argue that the "human factor is truly security's weakest link" [156]. First, insiders may maliciously interfere with data and network security to the disadvantage of other individuals [179]. Second, an abundance of incidents involving lost and stolen property (e.g., laptops and storage devices), as well as individuals' susceptibility to deception and social engineering are evidence of breaches characterizing weakest-link vulnerabilities. Third, users may out of convenience, cognitive limitations or economic considerations engage in insecure practices. The most common example is the prevalent use of weak passwords in organizations reported in many empirical and behav-

ioral studies [37,221].[4] Password misuse can sometimes be remedied, but it "only requires one indiscretion to destroy a secret" [53] such as the identities of members of a darknet (for filesharing purposes).

2.3.3 Best shot security game

In this game, the overall protection level depends on the maximum contribution offered over all entities. Hence, we have $H(e_i, e_{-i}) = \max(e_i, e_{-i})$, so that Eqn. (2.1) becomes

$$U_i = M - pL(1 - s_i)(1 - \max(e_i, e_{-i})) - be_i - cs_i \,. \tag{2.4}$$

As an example of a best shot game, consider a set of walls of which the highest sets the effectiveness benchmark. Among information systems, networks with built-in redundancy, such as peer-to-peer, sensor networks, or even Internet backbone routes, share resilience qualities with the best shot security game; for instance, to completely take down communications between two (presumably highly connected) backbone nodes on the Internet, one has to shut down all possible routes between these two nodes. Censorship-resistant networks are another example of best shot games. A piece of information will remain available to the public domain as long as a single node serving that piece of information can remain unharmed [57].

[4] For example, the analysis of a leaked password data set from phpbb.com revealed that 16% of passwords matched a person's first name, 14% of passwords were patterns on the keyboard, 4% were variations of the word "password" etc. Analysis available as online article "PHPBB Password Analysis" at: `http://www.darkreading.com/blog/archives/2009/02/phpbb_password.html`.

2.4 Loosely coupled security interdependencies

In a loosely coupled network consequences may differ for network participants [75]. For example, an attacker might target wealthy or inexperienced users. Other individuals would stay unharmed and are never targeted. In the following, we present two variations of the weakest-target game and propose security interpretations.

2.4.1 Weakest-target security game (without mitigation)

Here, an attacker will *always* be able to compromise the entity (or entities) with the lowest protection level, but will leave other entities unharmed. This game derives from the security game presented in [47]. Formally, we can describe the game as follows:

$$H(e_i, e_{-i}) = \begin{cases} 0 & \text{if } e_i = \min(e_i, e_{-i}), \\ 1 & \text{otherwise,} \end{cases} \tag{2.5}$$

which leads to

$$U_i = \begin{cases} M - pL(1 - s_i) - be_i - cs_i & \text{if } e_i = \min(e_i, e_{-i}), \\ M - be_i - cs_i & \text{otherwise.} \end{cases} \tag{2.6}$$

The weakest-target game markedly differs from the weakest-link. There is still a decisive security level that sets the benchmark for all individuals. It is determined by the individual(s) with the lowest chosen effort level. However, in this game all entities with a protection effort strictly larger than the minimum will remain unharmed.

In information security, this game captures the situation in which an attacker is interested in securing access to an arbitrary set of entities with the lowest possible effort. Ac-

cordingly, she will select the machines with the lowest security level. An attacker might be interested in such a strategy if the return on attack effort is relatively low, for example, if the attacker uses a compromised machine to distribute spam. Such a strategy is also relevant to an attacker with limited skills, a case getting more and more frequent with the availability of automated attack toolboxes [210]; or, when the attacker's goal is to commandeer the largest number of machines using the smallest investment possible [73]. Likewise, this game can be useful in modeling insider attacks – a disgruntled employee may for instance very easily determine how to maximize the amount of damage to her corporate network while minimizing her effort.

2.4.2 Weakest-target security game (with mitigation)

This game is a variation on the above weakest-target game. The difference is that, the probability that the attack on the weakest protected player(s) is successful is now dependent on the security level $\min e_i$ chosen. That is,

$$H(e_i, e_{-i}) = \begin{cases} 1 - e_i & \text{if } e_i = \min(e_i, e_{-i}), \\ 1 & \text{otherwise,} \end{cases} \tag{2.7}$$

so that

$$U_i = \begin{cases} M - pL(1 - s_i)(1 - e_i) - be_i - cs_i & \text{if } e_i = \min(e_i, e_{-i}), \\ M - be_i - cs_i & \text{otherwise.} \end{cases} \tag{2.8}$$

This game represents a nuanced version of the weakest-target game. Here, an an attacker is not necessarily assured of success. In fact, if all individuals invest in full protection, not a single machine will be compromised. This variation allows us to capture scenarios where,

for instance, an attacker targets a specific vulnerability, for which an easily deployable countermeasure exists.

2.5 Nash equilibrium analysis

We consider strategic interactions (called *games*) of the following simple form: the individual decision-makers (also called *players*) of a game simultaneously choose actions that are derived from their available strategies. The players will receive payoffs that depend on the combination of the actions chosen by each player.

More precisely, consider a set $N = \{2, ..., n\}$ of players. Denote as S_i the set of *pure* (i.e., deterministic) strategies available to player i, and denote as s_i an arbitrary member of i's strategy set. A probability distribution over pure strategies is called a *mixed* strategy σ_i. Accordingly, the set of mixed strategies for each player, Σ_i, contains the set of pure strategies, S_i, as degenerate cases. Each player's randomization is statistically independent of those of the other players. Then, u_i represents player i's payoff (or *utility*) function: $u_i(\sigma_i, \sigma_{-i})$ is the payoff to player i given her strategy (σ_i) and the other players' strategies (summarized as σ_{-i}). An n-player game can then be described as $G = \{N; \Sigma_i, \Sigma_{-i}; u_i, u_{-i}\}$.

Players are in a Nash equilibrium if a change in strategies by any one of them would lead that player to obtain a lower utility than if she remained with her current strategy [161]. Formally, we can define a Nash equilibrium as follows:

Definition of Nash Equilibrium: A vector of mixed strategies $\sigma^* = (\sigma_1^*, ..., \sigma_n^*) \in \Sigma$ comprises a mixed-strategy Nash equilibrium of a game G if, for all $i \in N$ and for all $\sigma_i^{'} \in \Sigma_i$, $u_i(\sigma_i^{'}, \sigma_{-i}^*) - u_i(\sigma_i^*, \sigma_{-i}^*) \leq 0$.

A pure-strategy Nash equilibrium is a vector of pure strategies, $s^* \in S$, that satisfies the equivalent condition.

The main advantage of the concept of Nash equilibrium resides in its simplicity. However, because Nash equilibria rely on very stringent assumptions on the capabilities and objectives of each player, they can predict counter-intuitive or unrealistic outcomes. Thus, the economics community has provided an increasing number of refinements to strengthen the concept of Nash equilibrium (e.g., perfect vs. proper equilibria). Similarly, some have investigated how to weaken the rational choice assumptions on which the Nash equilibrium concept is built: a rational player is expected to demonstrate error-free decision-making, to have perfect foresight of the game and to be unbounded in her computational abilities.[5] Intuitively, players such as network users (which are not necessarily perfectly rational) or automated agents (which can be faulty, due to software bugs or misconfiguration, or have limited computational resources) will likely deviate from these rigid assumptions [40, 74, 93].

We next determine the equilibrium outcomes where each individual chooses protection effort and self-insurance investments unilaterally, in an effort to maximize her own utility. We focus our analysis on symmetric equilibrium strategies since the implementation of asymmetric investment decisions is a difficult coordination problem in networks without

[5] See, for example, the research by [5, 82, 152, 177].

information exchange between agents. In Section 2.6, we then compare these results to the protection efforts and self-insurance levels chosen if coordinated by a social planner.

2.5.1 Total effort security game

Let us focus on player i, and consider e_k for $k \neq i$ as exogenous. Then, U_i is a function of two variables, e_i and s_i. From Eqn. (2.2), U_i is twice differentiable in e_i and s_i, with $\partial^2 U_i/\partial s_i^2 = 0$ and $\partial^2 U_i/\partial e_i^2 = 0$. Hence, according to the second derivative test, only $(e_i, s_i) \in \{(0,0),(0,1),(1,0),(1,1)\}$ can be an extremum – that is, possible Nash equilibria are limited to these four values (or to strategies yielding a payoff constant regardless of e_i and/or s_i). As long as at least one of b or c is strictly positive, $(e_i, s_i) = (1,1)$ is always dominated by either $(e_i, s_i) = (1,0)$ or $(e_i, s_i) = (0,1)$ and cannot define a Nash equilibrium. Let us analyze the three other cases:

- $(e_i, s_i) = (0,0)$. Replacing in Eqn. (2.2), we get a payoff for passivity:

$$U_i = M - pL\left(1 - \frac{1}{N}\sum_{k\neq i} e_k\right) . \tag{2.9}$$

- $(e_i, s_i) = (0,1)$. Replacing in Eqn. (2.2), we get a payoff for full self-insurance:

$$U_i = M - c\,. \tag{2.10}$$

- $(e_i, s_i) = (1,0)$. Replacing in Eqn. (2.2), we get a payoff for full protection:

$$U_i = M - pL\left(1 - \frac{1}{N} - \frac{1}{N}\sum_{k\neq i} e_k\right) - b\,. \tag{2.11}$$

Result: *After investigating Eqs. (2.9–2.11) we can identify three Nash equilibrium strategies.*

- *Passivity eq.:* If $pL < bN$ and $pL < c$, then $(0, 0)$ (passivity) is a symmetric Nash equilibrium.

- *Full self-insurance eq.:* If $c < pL$ and $c < b + pL\left(1 - \frac{1}{N}\right)$, then $(0, 1)$ (full self-insurance) is a symmetric Nash equilibrium.

- *Full protection eq.:* If $bN < pL$ and $b < c$, then $(1, 0)$ (full protection) is a symmetric Nash equilibrium.

These algebraic inequality conditions on parameters are both necessary and sufficient for the specified Nash equilibrium to be strict.

To derive these conditions, note that for any one of the three viable symmetric strategies to be a (strict) Nash equilibrium, it must be the case that for each player participating in such a strategy, it is disadvantageous for that player to switch to an alternative strategy. We obtain the conditions above by writing down the payoff for a given player i, using the definition of a (strict) Nash equilibrium, and simplifying. Below we will carry out the full argument for characterizing the passivity strategies that are Nash equilibrium. The other conditions are derived in a similar manner.

In the case of the symmetric passivity strategy, the payoff for player i is $M - pL$, and we seek to derive conditions under which any alternative strategy of the form $(e_i, s_i) \neq (0, 0)$ yields an inferior payoff. We begin by considering what happens when player i

adjusts her strategy by increasing only her level of self-protection. Since none of the other players are protecting, when player i chooses a non-zero protection level e_i, her payoff is $M - pL(1 - \frac{e_i}{N}) - be_i = M - pL + e_i\left(\frac{pL}{N} - b\right)$. Thus, it is a strict disadvantage for player i to increase only her protection level if and only if $\frac{pL}{N} - b < 0$, or equivalently, $pL < bN$.[6] Similarly, if player i increases her self-insurance level to some non-zero quantity s_i, her payoff becomes $M - pL(1 - s_i) - cs_i = M - pL + s_i(pL - c)$; so it is a disadvantage for player i to increase her level of self-insurance if and only if $pL - c < 0$, or equivalently, $pL < c$. Since player i could conceivably increase her payoff by adjusting her levels of self-protection or self-insurance individually, the preceding argument shows that the conditions $pL < bN$ and $pL < c$ are necessary for the symmetric passivity strategy to be a strict Nash equilibrium. To complete the argument for sufficiency of these conditions, we now show that under the specified conditions $pL < bN$ and $pL < c$, any strategy of the form $(e_i, s_i) \neq (0, 0)$ results in a strictly inferior payoff for player i. For this we have:

$$
\begin{aligned}
U_i &= M - pL(1 - \frac{e_i}{N})(1 - s_i) - be_i - cs_i && \text{(since all other players are passive)} \\
&= M - pL + e_i\left(\frac{pL}{N}(1 - s_i) - b\right) + s_i(pL - c) \\
&< M - pL + e_i\left(\frac{pL}{N}(1 - s_i) - b\right) && \text{(since } pL < c \text{ and } 0 \leq s_i\text{)} \\
&\leq M - pL + e_i\left(\frac{pL}{N} - b\right) && \text{(since } 0 \leq e_i\frac{pL}{N} \text{ and } 0 \leq s_i\text{)} \\
&< M - pL && \text{(since } 0 \leq e_i \text{ and } bN < pL\text{)}
\end{aligned}
$$

[6]Note that if $bN = pL$, then each player would be ambivalent about switching strategies. This strict equality condition can lead to various weak equilibrium strategies, but we consider the derivation of these conditions to be cumbersome and uninteresting and we do not address it in this treatment.

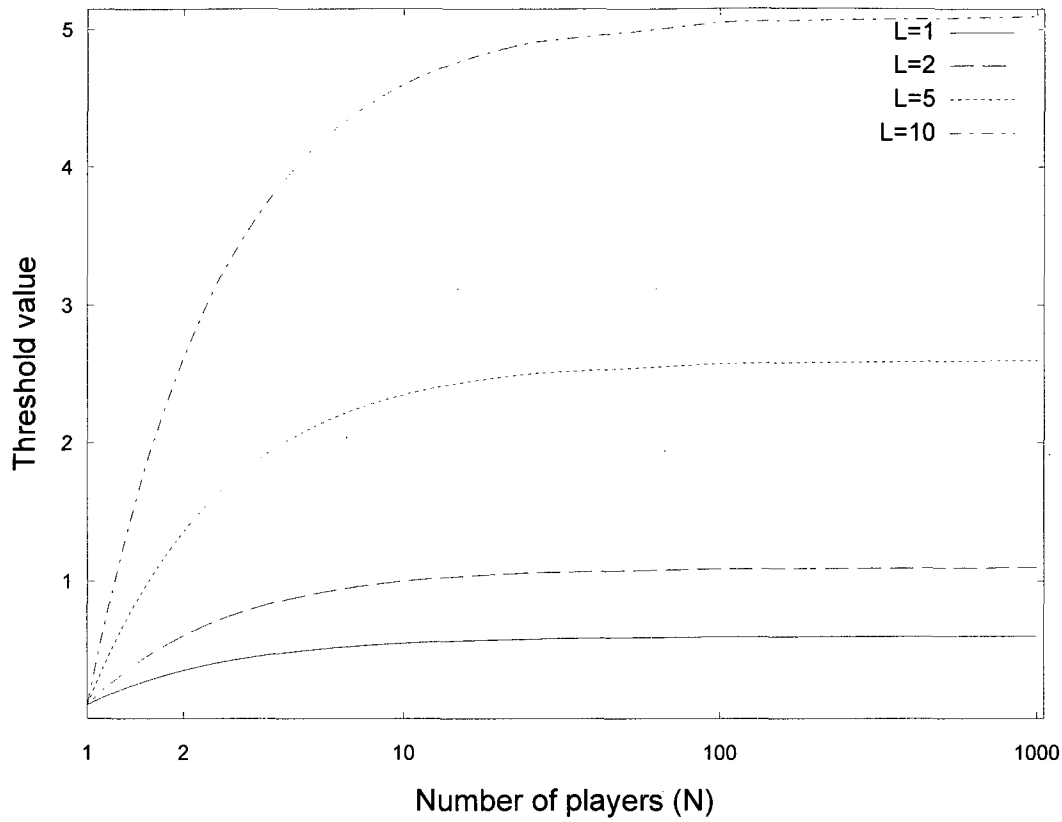

Figure 2.1: **Impact of network size, N, and loss magnitude, L, on threshold value for total effort self-insurance strategy.** (Protection cost b=0.1, Attack probability p=0.5)

This completes the derivation of necessary and sufficient parameter conditions for passivity to be a strict Nash equilibrium. As remarked previously, the derivations for other conditions are similar. Since there are eight additional conditions to address, we omit the remaining derivations for the purpose of saving the reader's time.

Notice that there are several case conditions in which it is possible to have multiple types of equilibria within the same case. In particular, full protection and full self-insurance can both be Nash equilibria for the same set of parameters.

Increasing number of players N: As the number of players increases, protection equilibria become more and more unlikely to occur. Indeed, in a total effort scenario, benefits yielded by a player's investment in security have to be shared with all of the other participants, making it an increasingly uninteresting strategy for the player as the network grows.

Further, with increasing N the self-insurance equilibrium becomes more attractive (see threshold level $b+pL\left(1-\frac{1}{N}\right)$ plotted in Figure 2.1). That is, in larger networks protection has to compete with self-insurance as a viable option for individual security investments.

2.5.2 Weakest-link security game

Let $e_0 = \min_i(e_i)$. From Eqn. (2.3), we have $U_i = M - pL(1-s_i)(1-e_0) - be_i - cs_i$, so that $\frac{\partial U_i}{\partial s_i} = pL(1-e_0) - c$, and, for all i,

$$U_i \leq M - pL(1-s_i)(1-e_0) - be_0 - cs_i \ ,$$

which is reached for $e_i = e_0$. So, in a Nash equilibrium, everybody picks the same $e_i = e_0$. It follows that Nash equilibria are of the form $(e_0, 0)$ or $(0, 1)$.

- Selecting $(e_i, s_i) = (0, 0)$ yields a payoff for passivity:

$$U_i = M - pL.$$

- Selecting $(e_i, s_i) = (0, 1)$ yields a payoff for full self-insurance:

$$U_i = M - c.$$

- Selecting $(e_i, s_i) = (e_0, 0)$ yields a payoff for self-protection at uniform level e_0:

$$U_i = M - pL(1 - e_0) - be_0.$$

Result: *In the weakest-link security game, we can identify three types of Nash equilibrium strategies. However, there exist multiple pure protection equilibria.*

- *Passivity eq.:* If $pL < c$, then $(0, 0)$ (passivity) is a symmetric Nash equilibrium.
- *Full self-insurance eq.:* If $c < pL$ then $(0, 1)$ (full self-insurance) is a symmetric Nash equilibrium.
- *Multiple protection equilibria:* If $b < pL$ and $b < c$, then $(\hat{e}_0, 0)$ (protection at level $\hat{e}_0$) is a symmetric Nash equilibrium for any $\hat{e}_0$ between $\frac{pL-c}{pL-b}$ and 1.

In the weakest-link security game, several Nash equilibria can co-exist for the same parameter settings. We find that protection and self-insurance equilibria compete with $c < pL$. Further, passivity and protection are both Nash equilibria if $b < pL < c$.

Interestingly, an increase in the loss magnitude, L, will reduce the feasible space of concurrent Nash strategies for self-protection by reducing the perceived difference between the cost of protection and self-insurance (i.e., the existence of feasible joint protection levels, $\hat{e}_0$, between $\frac{pL-c}{pL-b}$ and 1). This effect reduces strategic uncertainty, when the loss is large compared with security costs. See Figure 2.2.

Increasing number of players N: The weakest-link security game, much like the tacit coordination game of [211] has highly volatile protection equilibria when the number of

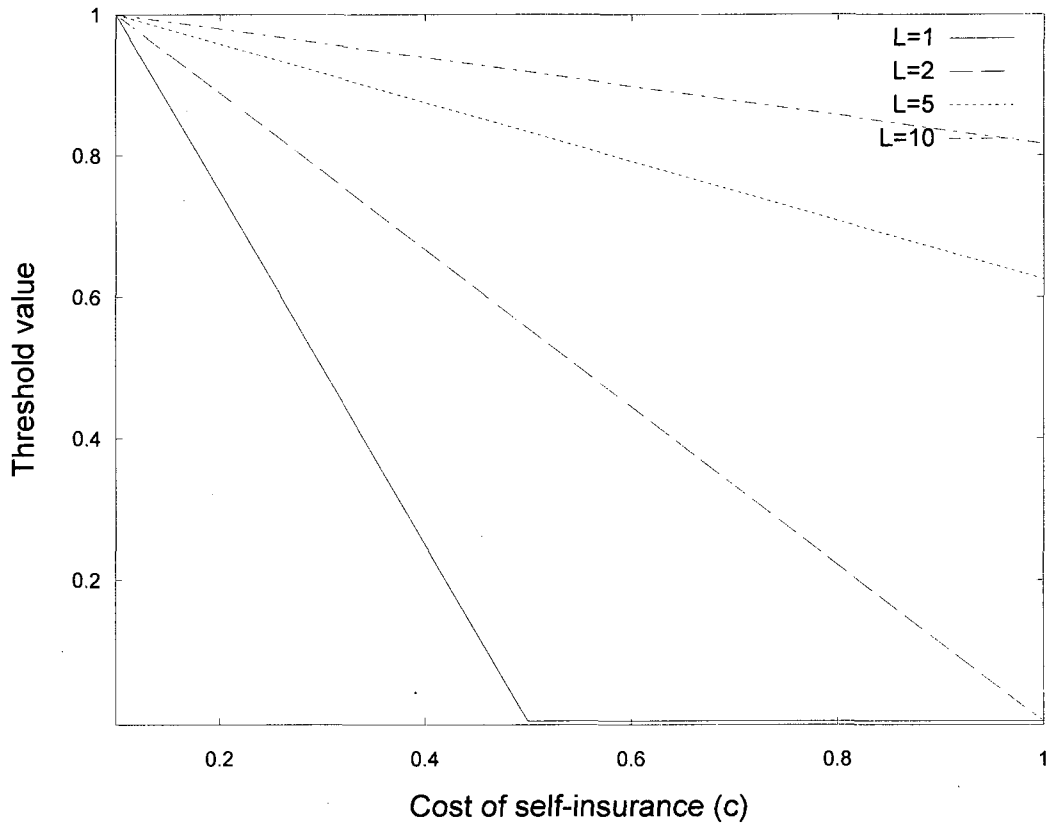

Figure 2.2: **Impact of loss magnitude, L, on threshold value for weakest-link protection strategy.** (Protection cost b=0.1, Attack probability p=0.5)

players increases. In fact, any protection equilibrium has to contend with the strategic certainty of a self-insurance equilibrium. To view this, consider the cumulative distribution function $F(e_i)$ over the protection strategies e_i of a given player i. From what precedes, with pure strategies, in the Pareto-optimum, $F(1) = 1$ and $F(e_i) = 0$ for $e_i < 1$. Assuming all N players use the same c.d.f. F, then the c.d.f. of $e_0 = \min_i\{e_i\}$ is given by $F_{\min}(e_0) = 1-(1-F(e_0))^N$ [211]. So, $F_{\min}(1) = 1$ and $F_{\min}(e_0) = 0$ for $e_0 < 1$ as well. Now, assume there is an arbitrarily small probability $\varepsilon > 0$ that one player will defect, that is $F(0) = \varepsilon$. Then, $F_{\min}(0)$ converges quickly to 1 as N grows large. That is, it only takes the slightest rumor that one player may defect for the whole game to collapse to the $(e_i, s_i) = (0, 1)$ equilibrium.

If players are infallible and will not defect then the actual number of players has no impact on the expected outcome [107].

2.5.3 Best shot security game

Let $e^* = \max_i(e_i)$. Eqn. (2.4) gives

$$U_i = M - pL(1 - s_i)(1 - e^*) - be_i - cs_i \ .$$

An argument analogous to those above shows that the only possible equilibrium strategies are ones in which e_i and s_i take binary values. Clearly, $(e_i, s_i) = (1, 1)$ is suboptimal, so that three strategies may yield the highest payoff to user i.

- Selecting $(e_i, s_i) = (0, 0)$ yields a payoff for passivity:

$$U_i = M - pL(1 - e^*).$$

- Selecting $(e_i, s_i) = (0, 1)$ yields a payoff for full self-insurance:

$$U_i = M - c.$$

- Selecting $(e_i, s_i) = (1, 0)$ yields a payoff for full self-protection:

$$U_i = M - b.$$

Result: *From the above relationships, we can identify the following pure Nash equilibrium strategies.*

- *Passivity eq.:* If $pL < b$ and $pL < c$, then $(0, 0)$ (passivity) is a symmetric Nash equilibrium.

- *Full self-insurance eq.:* If $c < b$ and $c < pL$ then $(0, 1)$ (full self-insurance) is a symmetric Nash equilibrium.

- *No symmetric pure protection eq.:* There is no pure-strategy symmetric Nash equilibrium for this game.

In particular, there is no *pure symmetric* protection equilibrium in this game. For one protection equilibrium to exist, we would need $b < c$ and $pL > b$. But even assuming that this is the case, as long as the game is synchronized, players endlessly oscillate between

securing as much as possible ($e_i = 1$) and free-riding ($e_i = 0$). This is due to the fact that as soon as one player secures, all others have an incentive to free-ride. Conversely, if everybody free-rides, all players have an incentive to deviate and secure as much as possible.

If $b < pL$ and $b < c$, then the asymmetric strategy in which exactly one player protects and the rest do nothing forms an asymmetric Nash equilibrium. Further, a mixed strategy that includes positive investments for protection exists. If $b < c$, then agents protect fully with probability $1 - \left(\frac{b}{pL}\right)^{\frac{1}{N-1}}$ and remain passive with probability $\left(\frac{b}{pL}\right)^{\frac{1}{N-1}}$. There are no parameter values such that there can exist more than one type of pure symmetric Nash equilibria.

Increasing number of players N: In the absence of coordination between players, the outcome of this game is globally independent of the number of players N, as there is no protection equilibrium, and the insurance equilibrium is independent of the number of players. However, the game may be stabilized by using player coordination (e.g., side payments) for low values of N, something harder to do as N grows.

2.5.4 Weakest-target security game (without mitigation)

Fix the strategy point and let $\varepsilon < \frac{pL}{2b}$. Let e_0 be the minimum effort level of any player. Then no player selects a higher effort than $e_0 + \varepsilon$ because it dominates all higher effort levels. However, any player at e_0 would prefer to switch to $e_0 + 2\varepsilon$. Then the change in her payoff is greater than $pL - 2\frac{pL}{2b}b = 0$. Because this deviation is profitable this strategy

point is not an equilibrium.[7]

Result: *In the weakest-target game with an attacker of infinite strength we find that pure Nash equilibria for non trivial values of b, p, L and c do not exist.*

Mixed strategy equilibria. While no pure Nash equilibria exist, let us explore the existence of a mixed strategy equilibrium. We use the shorthand notation $e_i = e$, $s_i = s$ here, and consider mixed strategies for choosing e. There are two cases to consider.

Case $c > pL$: If $c > pL$ then dominance arguments immediately lead to $s = 0$ meaning that nobody buys any self-insurance.

An equilibrium strategy may be parametrized by e. For a given player, the utility function U becomes a function of a single variable e. Let $f(e)$ be the probability distribution function of effort in the weakest-target game and let $F(e)$ be the cumulative distribution function of effort. Assuming only one player is at the minimum protection level, shall an attack occur, the probability of being the victim is then $(1 - F(e))^{N-1}$. (All N players choose protection levels greater than e.)

Then the utility is given by

$$U = M - pL(1 - F(e))^{N-1} - be\,. \tag{2.12}$$

In a Nash equilibrium, the first-order condition $dU/de = 0$ must hold, so that:

$$(N-1)pLf(e)[1 - F(e)]^{N-2} - b = 0$$

[7]While this proof assumes the player is initially at $(e_i, s_i) = (e_0, 0)$, it can be trivially extended to the case $(e_i, s_i) = (e_0, s)$ with $s > 0$ by picking $\varepsilon < \frac{pL}{2b}(1 - s) + \frac{cs}{2b}$ for any s in $(0, 1]$.

If we substitute $G = (1 - F(e))$ and $g = -f$ we can write $G^{N-2}dG/de = -b/p(N-1)L$, which, by integration yields

$$\int_{G(e)}^{G(0)} G^{N-2}dG = \int_{e}^{0} \frac{-b}{p(N-1)L}d\hat{e}\,,$$

that is

$$G^{N-1}\Big|_{G(e)}^{G(0)} = \frac{-b}{pL}e\,. \tag{2.13}$$

With $G(0) = 1$,

$$G(e) = \left(1 - \frac{b}{pL}e\right)^{\frac{1}{N-1}}\,.$$

Differentiating, we get

$$g(e) = -\frac{1}{N-1}\frac{b}{pL}\left(1 - \frac{b}{pL}e\right)^{-\frac{N-2}{N-1}}\,,$$

and, replacing $g = -f$ we find,

$$f(e) = \frac{1}{N-1}\frac{b}{pL}\left(1 - \frac{b}{pL}e\right)^{-\frac{N-2}{N-1}}\,, \tag{2.14}$$

as the probability distribution function of self-protection in a mixed Nash equilibrium.

Case $c \leq pL$: Now let us consider a game with insurance under the more reasonable assumption $c \leq pL$; that is, insurance is not overpriced compared to expected losses. Dominance arguments indicate that a Nash strategy must be of the form $(e, s) \in \{(e, 0), e \geq 0\} \cup \{(0, 1)\}$.

Let q be the probability that a player chooses strategy $(e, s) = (0, 1)$. That is, $F(0) = q$. Because insurance is independent of protection, we can reuse Eqn. (2.13) with the new

boundary $G(0) = 1 - q$:

$$G(e) = \left((1-q)^{N-1} - \frac{b}{pL}e\right)^{\frac{1}{N-1}} \tag{2.15}$$

However, since we are now including self-insurance, a second condition must hold. The payoff for strategy $(e, s) = (0, 1)$ must equal the payoff for all other strategies.

Specifically, we may compare payoffs for strategies $(e, s) = (\varepsilon, 0)$ and $(e, s) = (0, 1)$ which gives, by continuity as $\varepsilon \to 0$,

$$pL(1-q)^{N-1} = c\,. \tag{2.16}$$

Together Eqs. (2.15) and (2.16) yield:

$$F(e) = 1 - G(e) = 1 - \left(\frac{c - be}{pL}\right)^{\frac{1}{N-1}}\,,$$

which, differentiating, gives

$$f(e) = \frac{1}{N-1}\frac{b}{pL}\left(\frac{c-be}{pL}\right)^{\frac{1}{N-1}-1}\,. \tag{2.17}$$

This allows us to compute how often strategy $(e, s) = (0, 1)$ is played:

$$q = F(0) = 1 - \left(\frac{c}{pL}\right)^{\frac{1}{N-1}}\,. \tag{2.18}$$

Result: *In the weakest-target game with an attacker of infinite strength, a mixed Nash equilibrium strategy exists. The individual's strategy is given by Eqs. (2.17) and (2.18).*

Also note that, per Eqn. (2.17) and continuity arguments, the upper bound for protection effort is given by $e_{\max} = c/b$, which can be less than 1 when protection costs dominate insurance costs $b > c$.

Increasing number of players N: From Eqn. (2.18), we can directly infer that an increase in the number of participating players decreases the probability that a full self-insurance strategy is chosen. When N grows large, q tends to zero, which means that players increasingly prefer to gamble in order to find a protection level that leaves them unharmed.

2.5.5 Weakest-target security game (with mitigation)

Let us assume that there exists a Nash equilibrium where $0 < K < N$ players who satisfy $e_i = e_0 = \min(e_i, e_{-i})$, while $(N - K > 0)$ players satisfy $e_i > e_0$. We can show that such an equilibrium does not exist and that players rather congregate at the highest protection level if certain conditions are met. Due to space constraints, we will only sketch the analysis of this equilibrium. By computing the partial derivatives $\partial U_i/\partial s_i$ and $\partial U_i/\partial e_i$, and discriminating among values for e_i and s_i, we get the following results.

Result: In contrast to the infinite strength weakest-target game we find that a pure Nash equilibrium may exist.

- *Full protection eq.:* If $b \leq c$ we find that the full protection equilibrium $(\forall i, (e_i, s_i) = (1, 0))$ is the only possible pure Nash equilibrium.
- *For $b > c$ we can show that no pure Nash equilibrium exists.*
- *There are no pure self-insurance equilibria.*

Mixed strategy equilibrium To complement this analysis we also present the mixed strategy equilibrium. The derivation is similar to the one given by Eqs. (2.12–2.18), however,

with an additional substitution step. This gives the resulting distribution,

$$F(e) = 1 - \left(\frac{c - be}{pL(1-e)}\right)^{\frac{1}{N-1}}, \tag{2.19}$$

so that

$$f(e) = \frac{1}{N-1}\left(\frac{(b-c)pL}{pL^2(1-e)^2}\right)\left(\frac{c-be}{pL(1-e)}\right)^{-\frac{N-2}{N-1}}.$$

Interestingly, the probability of playing $(e, s) = (0, 1)$ remains

$$q = F(0) = 1 - \left(\frac{c}{pL}\right)^{\frac{1}{N-1}} \tag{2.20}$$

Note that if $c < b$ there is a zero probability that $e = 1$ will be chosen by any player. The upper bound for protection effort is given by $e_{\max} = c/b$.

Result: *In the weakest-target game with an attacker of finite strength we find that a mixed Nash equilibrium strategy exists. The relevant equations are given in Eqs. (2.19–2.20).*

2.6 Identification of social optima

Organizations and public policy actors frequently attempt to identify policies that provide the highest utility for the largest number of people. This idea has been operationalized with the social optimum analysis. It states that a system has reached the optimum when the sum of all players' utilities is maximized. That is, the social optimum is defined by the set of strategies that maximize $\sum_i U_i$. Consider N players, and denote by $\Phi(e_1, s_1, \ldots, e_N, s_N)$ the aggregate utility, $\Phi(e_1, s_1, \ldots, e_N, s_N) = \sum_i U_i(e_i, s_i)$. The social optimum maximizes $\Phi(s_i, e_i)$ over all possible $(s_i, e_i) \in [0, 1]^{2N}$. Because enforcing a social optimum

may at times be conflicting with the optimal strategy for a given (set of) individual(s), to enforce a social optimum in practice, we may need to assume the existence of a "social planner" who essentially decides, unopposed, the strategy each player has to implement.

2.6.1 Total effort security game

Summing the utility given by Eqn. (2.2) over i, we realize that $\Phi((e_i, s_i)_{i\in\{1,\dots,N\}})$ can be expressed as a function of two variables, $E = \sum_i e_i$ and $S = \sum_i s_i$. Φ is continuous and twice differentiable in E and S, and the second derivative test tells us that the only possible extrema of Φ are reached for the boundary values of E and S, that is $(E, S) \in \{0, N\}^2$. In other words, the only possible social optima are 1) passivity (for all i, $(e_i, s_i) = (0, 0)$), 2) full protection (for all i, $(e_i, s_i) = (1, 0)$), or 3) full insurance (for all i, $(e_i, s_i) = (0, 1)$). As long as one of b or c is strictly positive, a social planner will never advise agents to invest into protection and self-insurance at the same time.

By comparing the values of Φ in all three cases, we find that if $b < pL$ and $b < c$ then all agents are required to exercise maximum protection effort $(e_i, s_i) = (1, 0)$. With $c < pL$ and $c < b$ all agents will self-insure at the maximum possible $(e_i, s_i) = (0, 1)$. A social planner will not encourage players to invest in security measures if they are too expensive ($c > pL$ and $b > pL$).

Result: *In the total effort security game we observe that in the Nash equilibrium there is almost always too little protection effort exerted compared to the social optimum. In fact, for a wide range of parameter settings no protection equilibria exist while the social*

optimum prescribes protection at a very low threshold.

- *Protection:* Except for very unbalanced parameter settings (i.e., $pL > bN$ and $c > b + pL\frac{N-1}{N}$) agents refrained from full protection. Now full protection by all agents is a viable alternative.

- *Self-insurance:* Full self-insurance now has to compete with full protection effort under a wider range of parameters.

- *Passivity:* Agents remain passive if self-insurance is too expensive ($c > pL$). However, we find a substantial difference with respect to protection behavior. Agents would selfishly refrain from protection efforts if $pL < bN$ since they would only be guaranteed the N-th part of their investments as returns. Now the social planner can ensure that all agents protect equally so that it is beneficial to protect up until $b < pL$.

2.6.2 Weakest-link security game

In the weakest-link game agents are required to protect at a common effort level to be socially efficient. We compute Φ by summing Eqn. (2.3) over i, and can express Φ as a function of e_i, s_i and $e_0 = \min_i(e_i)$. In particular, for all i, we obtain $\partial\Phi/\partial s_i = pL(1 - e_0) - c$. Studying the sign of $\partial\Phi/\partial s_i$ as a function of e_0 tells us that, if $b < c$ and $b < pL$ the social planner requires all agents to protect with maximum effort $(e_i, s_i) = (1, 0)$. If $c < b$ and $c < pL$ the social planner requires all agents to self-insure $(e_i, s_i) = (0, 1)$. Finally, the Nash equilibrium and social optimum coincide when security costs are high. Agents do

not invest in protection or self-insurance if $b > pL$ or $c > pL$.

Result: *The availability of self-insurance lowers the risk of below-optimal security in the Nash equilibrium since agents have an alternative to the unstable Pareto-optimal protection equilibrium. From the analysis of the weakest-link game with many agents we know that deviation from the Pareto-optimal highest protection level is very likely. A social planner can overcome these coordination problems.*

- *Protection:* The *Pareto-optimal* Nash equilibrium coincides with socially optimal protection. However, the protection level would likely be lower in the Nash case due to coordination problems.

- *Self-insurance:* The self-insurance equilibria are equivalent for the Nash and social optimum analysis.

- *Passivity:* A social planner cannot expand the range of parameter values at which it would be socially beneficial to protect or self-insure while passivity would be prescribed in the Nash equilibrium.

2.6.3 Best shot security game

We compute the social optimum by summing U_i given in Eqn. (2.4) over i, yielding that Φ can be expressed as a function of e_i, s_i, and e^*. It is immediate that, to maximize Φ, one should pick $e_i = 0$ for all i, except for one participant j, where $e_j = e^* \geq 0$. We then get $\partial Phi/\partial s_i = pL(1 - e^*) - c$, which tells us under which conditions on e^* (and

consequently on b, c, and pL) self-insurance is desirable.

We find that if $b/c < N$ (i.e., protection is not at a prohibitive cost compared to insurance and/or there is a reasonably large number of players), the social optimum is to have one player protect as much as possible, the others not protect at all, and no one insures. In practice, this may describe a situation where all participants are safely protected behind an extremely secure firewall. If, on the other hand $b/c > N$, which means there are either few players, insurance is very cheap compared to protection, then the best strategy is to simply insure all players as much as possible.

Result: *In the best shot security Nash outcome there is almost always too little effort exerted compared to the social optimum. Exceptions are few points in which full self-insurance remains desirable for the social planner and all agents remain passive.*

- *Protection:* Surprisingly, while protection is not even a Nash strategy we find that a social planner would elect an individual to exercise full protection effort.

- *Self-insurance:* Full self-insurance by every player is only desirable if protection costs are large. Therefore, for most cases the strategy of a social planner will not coincide with the only Nash equilibrium strategy.

- *Passivity:* In the Nash equilibrium agents are also too inactive. Passivity is highly undesirable from a social planner's perspective. Only if $NpL < b$ no agent will be selected to exercise maximum protection effort (while self-insurance might remain an option).

It is important to note that the social optimum variation that requires full protection by one individual results in the whole population being unharmed, since one highly secure individual is enough to thwart all attacks. Therefore, it is easy to see that protection is extremely desirable from a planners perspective. Out of the three classical public goods games with homogeneous agents the best shot game can benefit the most from a guiding hand.

2.6.4 Weakest-target security game (without mitigation)

We compute the social optimum by using Eqn. (2.8), assuming that $1 \leq K \leq N$ players pick $e_0 = \min_i(e_i)$. By studying the variations on Φ as a function as e_i, as a function of K, and as a function of s_i (for both the K players picking e_0 and the remainder of the players), we find that in the weakest-target game without mitigation a social planner would direct a single player to exacerbate no protection effort.

Essentially, this player serves as a direct target for a potential attacker. However, as long as $c < pL$ the player would be directed to maximize self-insurance $(e_i, s_i) = (0, 1)$. If insurance is too expensive ($c > pL$) then the social planner would prefer to leave the player uninsured $(e_i, s_i) = (0, 0)$. This strategy is independent of the cost of protection. The remaining $N - 1$ players have to select their protection effort as $e_i = \varepsilon > 0$ (as small as possible). These players will not be attacked, and therefore will set their self-insurance to the possible minimum $(\varepsilon, 0)$. Passivity by all players is never an option in the social optimum.

Result: *A social planner can easily devise a strategy to overcome the coordination problems observed in the Nash analysis for the weakest-target game with mitigation. We found that no pure Nash strategy exists and, therefore, had to rely on the increased rationality requirement for entities to play a mixed strategy.*[8] *The average payoff for each player in the social optimum is considerably higher compared to the mixed Nash equilibrium.*

Understandably, without side-payments the node with the lowest protection effort is worse off compared to his peers. However, the social planner could choose to devise a so-called "honeypot" system with the sole goal of attracting the attacker while only suffering a marginal loss. A honeypot is a computer system (or another device) that is explicitly designed to attract and to be compromised by attackers. It serves usually a double purpose. First, it will detract attention from more valuable targets on the same network. Second, if carefully monitored it allows gathering of information about attacker strategies and behaviors, e.g., early warnings about new attack and exploitation trends [173].

An interesting aspect of the social optimum solution is the question how the individual is selected (if a honeypot system cannot be devised). Obviously, a social planner might be able to direct an individual to serve as a target (in particular, if $c < pL$). However, if insurance costs are large being a target requires an almost certain sacrifice (dependent on the value of p). In anthropology and economics there are several theories that relate to an individuals willingness to serve as a sacrificial lamb. Most prominently, altruism and heroism come to mind. Simon also introduced the concept of docility. This theory

[8]Economists are generally cautious regarding the assumption that individuals can detect and adequately respond to mixed strategy play by opponents [71, 142, 195].

refers to an individual's willingness to be taught or to defer to the superior knowledge of others [198].

2.6.5 Weakest-target security game (with mitigation)

We adopt the same strategy for finding Φ's maximum as in the unmitigated case – that is, summing Eqn. (2.6) over i, and then studying the variations of Φ over K, s_i and e_0.

The first observation is that the social planner might prescribe the same strategy as in the case of the weakest-target game without mitigation. However, now the planner has a second alternative. Since an attacker will not be able to compromise players if they are fully protected we find that $(e_i, s_i) = (1, 0)$ for all N players is a feasible strategy. The tipping point between the two strategies is at $Nb < c$. If this condition holds the social planner would elect to protect all machines in favor of offering one node as honeypot and investing in its self-insurance. Note that again we find that if protection and self-insurance are extremely costly the planner will elect to sacrifice one entity without insurance. Passivity is not a preferable option.

Result: *Compared to the weakest-target game without mitigation the social planner is better off if protection is cheap. Otherwise the planner has to sacrifice a node with or without self-insurance. Interestingly, while compared to the pure Nash equilibrium outcome the social planner can increase the overall utility in the network we find that security expenditures are lowered. In the Nash equilibrium agents were willing to fully protect against*

threats as long as $(b \leq c)$.

The last observation also holds for the mixed strategy case in both weakest-target games (with or without mitigation). That is, agents exert ***more*** *effort in the Nash equilibrium (except when* $Nb < c$ *for the game with mitigation).*

2.7 Practical implications

The results we obtained, and notably the disconnect between social optima and Nash equilibria we observed, lead to a number of remarks that may prove relevant to organizational strategy. However, we want to preface this discussion by pointing out that our analysis is a first comparison of different security games with two security options under common, but restrictive assumptions.[9]

Most notably, we assume agents to be risk-neutral providers of the public protection good. In our game formulation we also simplified cost of protection (and insurance) to be linear. Including different risk preferences, as well as uncertainty and limited information about important parameters of the game would be important steps towards a sensitivity analysis of our results. Shogren found, for example, that risk-averse agents will increase their contributions if information about other agents actions is suppressed [197]. Others, e.g., [188], have obtained more nuanced results. We defer a more extensive analysis of such phenomena to future work, but believe that the main trends and differentiating features

[9]Rue and Pfleeger provide an informative overview of modeling assumptions for a number of cybersecurity investment models [184].

between security games we observed remain largely unchanged.

Security scenario identification: We find that security predictions vary widely between the five different games. Similarly, policies set by a social planner do not only yield different contribution levels but may also switch the recommended security action from protection to self-insurance and vice versa. Chief Security Officers' tasks involve a careful assessment of threat models the company is faced with.

We want to emphasize that an integral part of the threat model should be an assessment of the organizational structure including system resources and employees. Similarly important is a detailed consideration whether resources are protected independently or by an overarching system policy. For example, replication, redundancy and failover systems (that automatically switch to a standby database, server or network if the primary system fails or is temporarily shut down for servicing) should most likely not be treated as independent resources.

Managers should consider how the organizational structure of resources matches potentially existing policies. For example, we can see that a policy that requires full protection by every individual is sub-optimal if the most likely threat and organizational structure fits the description of a best shot game. Contributions resources are squandered and are likely to deteriorate. Not to mention that employees may simply ignore the policy over time. See, for example, recent survey results that highlight that 35% of white-collar employees admit to violations of security policies [114].

Selection of defense posture: A security professional might be faced with an unidenti-

fiable organization and system-policy structure. However, we want to highlight that our research allows a more careful choice between security options if managers can redesign organizations and policies. For example, the choice between a system-wide firewall and intrusion detection system versus an individual alternative has important implications on how incentives drive security-relevant behavior over time. Individual systems will better preserve incentives, however, might have negative cost implications. The same choice applies between the availability of backup tools and protective measures.

Leveraging strategic uncertainty: The example of the weakest-target game shows the importance of the degree of dependency between agents. We show that in larger organizations a much lower average level of self-insurance investments will be achieved because the strategic dependence between actors is reduced. However, in turn more agents will elect to protect their resources ($e_i > 0$ for more players). In contrast, agents in small groups will respond to the increasing strategic uncertainty caused by the increased interdependency by self-insuring their resources more often.

Introducing a social planner into the weakest-target game completely removes strategic uncertainty and leads to both reduced self-insurance and protection investments. This apparent paradox emphasizes that higher security investments do not necessarily translate in higher security -- but instead that *how* the investments are made are crucial to the returns.

2.8 Summary

We consider the problem of decision-making with respect to information security investments. To that effect, we model security interactions through a careful selection of games, some established (weakest-link, best shot, and total effort) and some novel (weakest-target, with or without mitigation). All of these games offer players two independent decision parameters: a protection level, e, which determines the level of security a player chooses for his resources; and a self-insurance level, s, which mitigates losses, shall a successful attack occur. We postulate that the five games considered cover a vast majority of practical security situations, and study them both from a rational agent's perspective (Nash equilibrium analysis) and from a central planner's view (social optimum analysis).

Our main findings are that the effects of central planning compared to laissez-faire considerably differ according to the game considered. While in a number of traditional cases borrowed from the public good literature, we observe that a central planner may increase the average protection level of the network, we also note that strategic decisions are highly impacted by the level of inter-dependency between the actions of different players.

In particular, we found that the common wisdom that having a central planner who decides upon security implementation always yields higher protection contributions by individual players does not hold. Indeed, it may at times be much more advantageous from an economic standpoint to invest in self-insurance instead of protecting systems, or to select a few, unprotected, sacrificial lambs in order to divert the attention of potential attackers. This is particularly the case in situations which exhibit a "strategic uncertainty" due to a

very strong correlation between the actions of different agents, for instance, in our weakest-target game where the least secure player is always the one attacked.

With the analyses in this work we aim for a more thorough understanding of the ecology of security threats and defense functions an individual or organization faces and has to respond to. We have generalized and developed new models that represent vastly different security scenarios and will call for different actions. As Hirshleifer observed [107] a security practitioner will be presented with "all kinds of intermediate cases and combinations," e.g., social composition functions involving all of these five rules as well as other not identified yet. Some minor variations would be the "location of the top decile, or the total of the best three shots, or the average of the best and worst shots, or the variance or skewness" etc.

Chapter 3

Security diversity: Heterogeneous agents

In this chapter, we derive Nash equilibria for the five different cases of security games with the focus to understand how the inclusion of heterogeneous actors influences predictions compared to a model with representative agents.

In the modeling of economic phenomena, added complexity (e.g., adding agents with more diverse tastes) does not always change strategic predictions substantially. On the other hand, we expect that heterogeneity impacts the actions of agents in security games in different ways, for example by: 1) Negotiating the trade-off between protection and self-insurance, 2) Highlighting certain strategies and focal points due to the inherent differences in the agent population, 3) (De-)stabilizing equilibrium predictions derived in the homogeneous case. We expect several conclusions from the homogeneous case to remain relevant. But as Hartley [99] argued "representative agents models conceal heterogeneity whether it is important or not." This analysis aims at pinpointing key differences and discuss their

implications.

The first contribution of the present chapter is to discuss arguments for and against homogeneity in security models.

Second, we extend and generalize models for homogeneous agents (as given in Chapter 2) to the significantly more complex heterogeneous agents case.

The third contribution is to exploit the results from the analysis to evaluate the impact of possible (centralized or distributed) intervention policies aiming at reaching an outcome beneficial to society as a whole.

The rest of this chapter is organized as follows. We elaborate in Section 3.1 on the relationship of our work with related research, and extend our game-theoretic models to take into account agent heterogeneity in Section 3.2. We analyze Nash equilibria stemming from these games in Section 3.3, and use this analysis to look into possible intervention mechanisms in Section 3.4. We conclude in Section 3.5.

3.1 Background: Heterogeneity in system security

The salient feature of the research presented in this chapter is to consider security as a combination of private and public goods in the context of *heterogeneous* agents.

Both the homogeneous and heterogeneous cases are relevant to security analysis. Homogeneous agents are characteristic of large populations following the same practices and choices by end-users, for instance, when most security decisions (e.g., patching) are auto-

mated, and all users run similar software. The lack of diversity, in particular in the market for operating systems, lends credibility to such scenarios [79], and is cited as a strong motivator for developers of malicious code to exploit the resulting correlated risks or to cheaply repeat attacks.

However, there are strong reasons to compare our earlier findings (from Chapter 2) with a model that includes heterogeneous agents into a model of security decision making.

Security through diversity. Recent technical proposals aim to achieve higher resilience to attacks by introducing diversity in network and protocol design. For example, Zhuang *et al.* report of a set of formal analysis tools that introduce heterogeneity in multi-person communication protocols [220]. O'Donnell and Sethu develop and test distributed algorithms optimizing the distribution of distinct software modules to different nodes in a network [164]. Research in IT economics has evaluated the decision making of a firm when faced with the option of increased diversity in its software base. In Chen *et al.*, the decision for increased heterogeneity depends largely on the assumed risk attitudes of the organization [45]. Investments into heterogeneity will change the expectation of losses and attack probabilities, but they also impact the cost of protection and self-insurance.

Chameleonic threats. Increased diversity is not a sufficiently strong protection against correlated security threats anymore. Already in 1995 the first macro viruses started targeting MS office on all compatible systems.[1] Modern cross-platform malware is capable of targeting also different operating systems. For instance, Linux-Bi-A/Win-Bi-A is written in

[1]The macro virus (Winword-Concept) targeted Microsoft Word on Apple and Microsoft systems. For more details see: `http://web.textfiles.com/virus/macro003.txt`.

assembler and able to compromise Windows and Linux platforms. Malicious code is also capable of crossing the boundary between desktop and mobile devices. For example, the hybrid pathogen Nimda, a worm that can spread as a virus as well, has successfully propagated on different media such as floppies, portable hard drives, and USB pen drives [216].

Potentially even more disruptive is malware carrying multiple exploit codes at once. For example, Provos *et al.* report that Web-based malware often includes exploits that are used 'in tandem' to download, store and then execute a malware binary [174]. These trends render users vulnerable to propagated threats if owners of different IT systems perceive protection as too costly or ineffective.

Heterogeneous investments patterns. Different organizations follow distinct patterns of IT investment. Parts of organizations often depend on legacy systems including weakly protected systems, or "boat anchors" with limited value to an organization [217]. Such legacy systems can allow skilled attackers to intrude a network. More generally, organizations and end users justify security investments with different assumptions about potential losses and probabilities of being attacked. This often depends on different knowledge about threats and means of protection and insurance [3]. This diversity is reflected in users' choices and security practices [121, 162]. Similarly, security decisions can follow different security paradigms often reflected in different organizational structures, for instance remote replication vs. offsite tape storage.

Finally, heterogeneous agents have notable implications in terms of policy design. For instance, Bull *et al.* [38] observe the state of heterogeneous networks and argue that no

single security policy will be applicable to all circumstances. They argue that, for a system to be viable from a security standpoint, individuals need to be empowered to control their own resources and to make customized security trade-offs.

In this chapter, we formally explore such theses, by studying individuals' incentives in non-cooperative games. In particular, we focus on the impact of heterogeneous agents on system security in different network structures.

3.2 Modification of the basic model

Different from our previous exposition in Chapter 2, protection costs per unit are not necessarily identical for each entity, and, while in the formal analysis that follows we make the assumption that all decisions are made simultaneously, we later discuss the impact of relaxing the synchronization assumption.

Each of $N \in \mathbb{N}$ players receives an individual endowment M_i. If she is attacked and compromised successfully she faces a loss L_i. Further, $b_i \geq 0$ and $c_i \geq 0$ denote the unit cost of protection and insurance, respectively.

The generic utility function of Player i can now be represented as:

$$U_i = M_i - pL_i(1 - s_i)(1 - H(e_i, e_{-i})) - b_i e_i - c_i s_i \,. \tag{3.1}$$

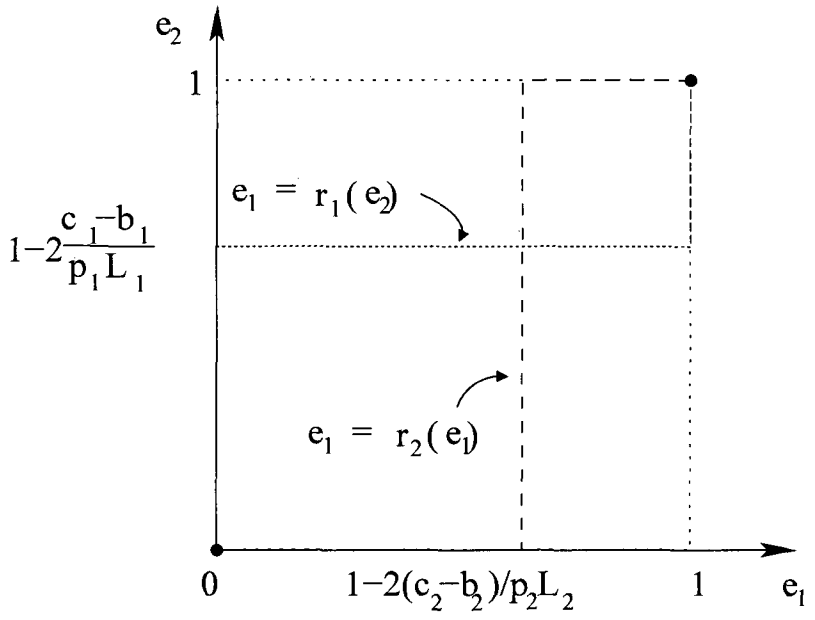

Figure 3.1: **Reaction functions for a two-player total effort game.** Bold lines and dots indicate potential Nash equilibria.

3.3 Nash equilibrium analysis

In this section, we derive Nash equilibria for the five different cases of security games with heterogeneous agents.

3.3.1 Total effort security game

The total effort game yields considerably different results depending on the number of players involved.

Two-player game Let us first start the discussion for the simple case $N = 2$. From the game description given by Eqn. (2.2), we get $U_1(e_1, s_1) = M_1 - pL_1(1 - s_1)(1 - (e_1 + e_2)/2) - b_1 e_1 - c_1 s_1$ for Player 1. The second partial derivative test indicates that there

is no local extremum, so that the only possible maxima of U_1 are given by $U_1(0,0) = M_1 - pL_1(1 - e_2/2)$, $U_1(1,0) = M_1 - pL_1(1/2 - e_2/2) - b_1$, $U_1(0,1) = M_1 - c_1$, or $U_1(1,1) = M_1 - b_1 - c_1$. With $b_1 > 0$, we immediately see that $U_1(0,1) > U_1(1,1)$, which tells us that fully insuring and protecting at the same time is a strictly dominated strategy for Player 1. The passivity strategy $(e_i, s_i) = (0,0)$ dominates the "protect-only" strategy $(e_i, s_i) = (1,0)$ when $b_1 > pL_1/2$.

Assuming $b_1 \leq pL_1/2$, the "protect-only" $(1,0)$ strategy dominates the "insure-only" $(0,1)$ strategy for Player 1 if and only if (all quantities being assumed to be defined):

$$e_2 > 1 - 2\frac{c_1 - b_1}{pL_1} \,. \tag{3.2}$$

A similar rationale yields the corresponding conditions for Player 2, leading to the reaction functions $e_1 = r_1(e_2)$ and $e_2 = r_2(e_1)$ plotted in Figure 3.1. By definition, Nash equilibria are characterized by fixed points $e_1 = r_1(e_2) = e_2 = r_2(e_1)$. From the above analysis summarized in Figure 3.1, this occurs for two values: when both agents fully protect and when both agents abstain from investing in protection. We note that both fixed points are stable, meaning that, if they are reached, minimal deviations in the strategy of one player are unlikely to perturb the actions of the other player.

Result: *The two-player total effort security game with heterogeneous agents presents the following equilibria:*

- *Full protection eq.:* If $b_1 \leq pL_1/2$, $b_2 \leq pL_2/2$ (protection costs are modest for both players), and the initial values $e_1(0)$ and $e_2(0)$ satisfy either $e_1(0) > 1 - 2(c_2 -$

$b_2)/(pL_2)$ or $e_2(0) > 1-2(c_1-b_1)/(pL_1)$ (at least one player is initially fairly secure, or at least one player faces very high self-insurance costs) then the (only) Nash equilibrium is defined by both players protecting but not insuring, that is, $(e_i, s_i) = (1, 0)$.

- *Multiple eq. without protection:* If the conditions above do not hold, then we have an insecure equilibria. Both players converge to $e_1 = 0$ and $e_2 = 0$. Their respective investments in self-insurance depend on whether their self-insurance premium is smaller than their potential losses: a player will fully insure if and only if $c_i < pL_i$, and will be passive otherwise.

A particularly interesting feature of the two-player version of the game is that expensive self-insurance or protection costs at *either* of the players directly condition which equilibrium can be reached. For instance, if one of the players has to pay a very high self-insurance premium in front of its protection costs, she will elect to protect, likely leading the other player to protect as well. Conversely, if either of the players faces a high protection premium ($b_i > pL_i/2$), the game will likely converge to an equilibrium without protection efforts. As we discuss later, this property can be used by some form of intervention to have the game converge to a desirable equilibrium.

More generally, in this game, each of the two players generally tracks what the other is doing. When moves are made perfectly simultaneously, this may result in oscillations between insecure and secure configurations. The only exception to this tracking behavior occurs when one player faces high security costs and a low self-insurance premium, while the other faces the opposite situation (low security costs, very high self-insurance

premium). In such a case, the game converges to the first player insuring, and the second player protecting. In short, extreme parameter values allow to remove network effects in this game.

N-player game (N large) In the more general case $N \geq 2$, we first notice that, for a security strategy to be meaningful, we need to have $b_i < pL_i/N$. This means that, as the number of players increases, individual protection costs have to become very small, or expected losses have to considerably increase. Failing that, self-insurance or passivity is always a better option.

Second, from Eqn. (2.2), we obtain that Eqn. (3.2) is generalized to

$$\frac{1}{N-1}\sum_{j\neq i} e_j > 1 - \frac{N}{N-1}\frac{c_i - b_i}{pL_i}, \tag{3.3}$$

as a condition for player i to select a protection-only strategy as opposed to an self-insurance-only strategy. Eqn. (3.3) tells us that, for large values of N, changes in a single player's protection strategy are unlikely to have much of an effect on the other players' strategies. Indeed, each player reacts to changes in the average protection level over the $(N-1)$ other players.

This observation brings the question of exactly how robust the N-player game is to a change in the strategy played by a given individual. Are "domino effects" possible, where changes in a single player's strategy, albeit with a minimal effect on all other players, lead another player to switch strategies, and eventually to large groups changing their plays?

To help us answer this question, let us consider $N > 2$, and $K \leq N$ arbitrary players that are initially (at time 0) unprotected. For instance, assume without loss of generality

that Players $1, \ldots, K$ are initially unprotected, and that

$$\frac{c_2 - b_2}{pL_2} \geq \frac{c_3 - b_3}{pL_3} \geq \ldots \geq \frac{c_K - b_K}{pL_K} .$$

Further assume that at a later time $t > 0$, Player 1 switches her strategy to full protection, that is, $e_1(t) = 1$. Assuming all players may have an incentive to protect (i.e., for all i, $b_i < pL_i/N$), Player 2 would also switch to full protection only if

$$\frac{1}{N-1} \sum_{j \neq 2} e_j(t) > 1 - \frac{1}{N-1} \frac{c_2 - b_2}{pL_2}$$

that is, only if

$$\frac{1}{N-1} \sum_{j \neq 2} e_j(0) + \frac{1}{N-1} > 1 - \frac{1}{N-1} \frac{c_2 - b_2}{pL_2} ,$$

which reduces to

$$\frac{1}{N-1} \sum_{j > K} e_j(0) + \frac{1}{N-1} > 1 - \frac{1}{N-1} \frac{c_2 - b_2}{pL_2} . \tag{3.4}$$

Player 2's switch causes Player 3 to switch too only if

$$\frac{1}{N-1} \sum_{j \neq 3} e_j(t) > 1 - \frac{1}{N-1} \frac{c_3 - b_3}{pL_3} ,$$

that is,

$$\frac{1}{N-1} \sum_{j > K} e_j(0) + \frac{2}{N-1} > 1 - \frac{1}{N-1} \frac{c_3 - b_3}{pL_3} . \tag{3.5}$$

From Eqs. (3.4) and (3.5) we get

$$\frac{c_2 - b_2}{pL_2} - \frac{c_3 - b_3}{pL_3} < 1 .$$

Iterating over the K players that are initially not protecting, we get:

$$\max_{2 \leq i \leq K} \frac{c_i - b_i}{pL_i} - \min_{2 \leq i \leq K} \frac{c_i - b_i}{pL_i} < K - 1 .$$

We can follow an identical derivation for the case where the K players switch from a protection strategy to a non-protection strategy. We then obtain the following necessary condition for "domino effects" to occur over K players, that is a switch in Player 1's strategy causing a switch in the strategy of K players:

$$\left| \max_{2\leq i\leq K} \frac{c_i - b_i}{pL_i} - \min_{2\leq i\leq K} \frac{c_i - b_i}{pL_i} \right| < K - 1 \,. \tag{3.6}$$

Result: *We have derived a stability measure of the heterogeneity of a total effort security game with N agents (Eqn. (3.6)). The more heterogeneous the players are, the more unlikely Eqn. (3.6) is to hold for large values of K. In other words, the more heterogeneous a system is, the more likely it is to be resilient to perturbations due to a single individual changing strategies.*

3.3.2 Weakest-link security game

Here again, we start by considering a two-player game. Computing partial derivatives in e_i and s_i from Eqn. (2.3), we observe that each player chooses either $(e_i, s_i) = (0, 1)$ (self-insurance strategy) or $(e_i, s_i) = (\min_{j\neq i} e_j, 0)$ (protection strategy, where in the two-player version of the game $\min_{j\neq i} e_j$ is naturally equal to the protection value chosen by the other player) in order to maximize their utility function.

Looking at the payoffs that can be obtained in both cases leads us to the reaction functions of both players, which we plot in Figure 3.2. In the figure, we see that a fixed-point is attained when $e_1 = e_2 = 0$ (self-insurance-only equilibria) and when both e_1 and e_2 are

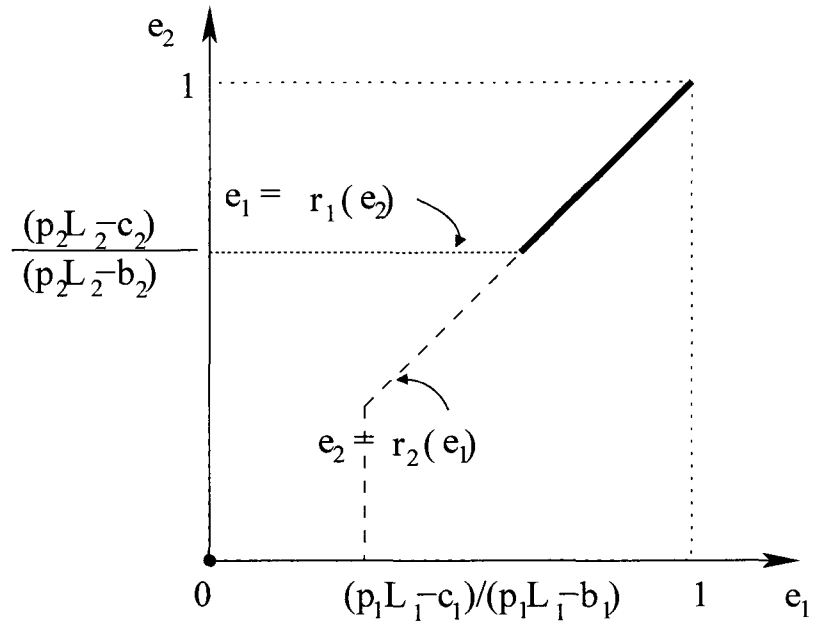

Figure 3.2: **Reaction functions for a two-player weakest-link game.** Bold lines and dots indicate potential Nash equilibria.

greater than $\max\{(pL_1 - c_1)/(pL_1 - b_1), (pL_2 - c_2)/(pL_2 - b_2)\}$.

Result: *Generalizing to N players, we obtain the following distinction for the weakest-link security game:*

- *Full protection eq.:* If, for all i, $pL_i > b_i$, and either 1) $pL_i < c_i$, or 2) $pL_i \geq c_i$ and $\hat{e}(0)$, the minimum of the security levels initially chosen by all players, satisfies

$$\hat{e}(0) > \max_{1 \leq i \leq N} \{(pL_i - c_i)/(pL_i - b_i)\} ,$$

  then we have a Nash equilibrium where everyone picks $(\hat{e}(0), 0)$.

- *Multiple eq. without protection:* All players select $e_i = 0$ if the conditions above do not hold. The value of self-insurance they select depends on their respective

valuations. Players for whom self-insurance is too expensive ($pL_i < c_i$) do not insure, with $s_i = 0$, while others choose full self-insurance, that is $s_i = 1$.

The likelihood of reaching a full protection equilibrium is conditioned by the player which has the largest difference between protection and self-insurance costs relative to its expected losses. In particular, it only takes one player with an self-insurance premium smaller than its protection cost ($b_i > c_i$) to make the full protection equilibrium unreachable. Hence, when N grows large, we expect protection equilibria to become more and more infrequently observed.

3.3.3 Best shot security game

Looking at the variations of the payload function U_i given in Eqn. (2.4) as a function of e_i and s_i tells us there are three possibilities for maximizing U_i: a passivity strategy $(0, 0)$, a secure-only strategy $(1, 0)$ and an insure-only strategy $(0, 1)$.

We get $U_i(0,0) = M_i - pL_i(1 - \max\{e_{-i}\})$, $U_i(1,0) = M_i - b_i$, and $U_i(0,1) = M_i - c_i$. We immediately notice that $b_i > c_i$ leads Player i to never invest in protection: either the player is passive, or she insures. If, on the other hand $b_i \leq c_i$, then player i chooses a protection strategy over a passivity strategy if and only if (b_i assumed greater than 0) we have $\max\{e_{-i}\} < 1 - b_i/pL_i$. We plot the reaction functions, in a two-player case, in Figure 3.3.

Result: *For the two-player best shot security game we can identify the following equilibria:*

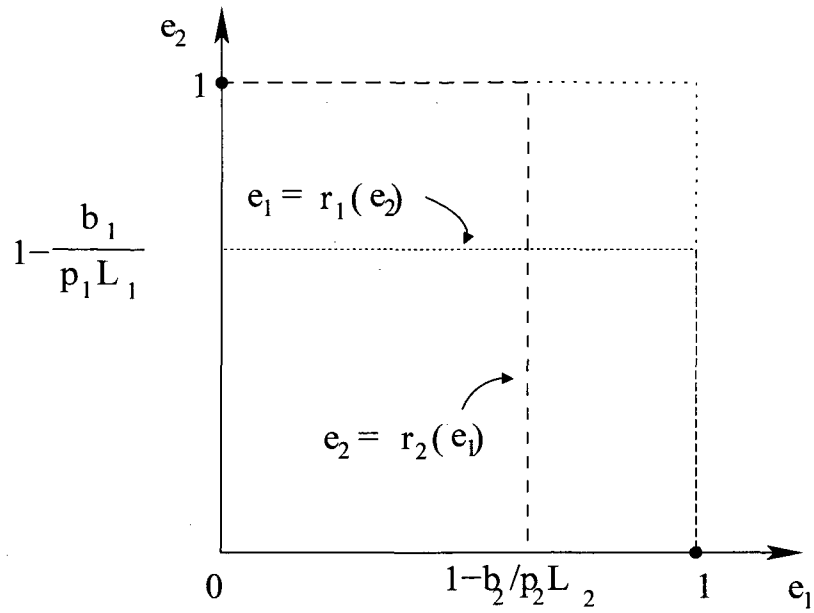

Figure 3.3: **Reaction functions for a two-player best shot game.** Bold dots indicate potential Nash equilibria. Protection costs are assumed here to be smaller than self-insurance costs for both players.

- *Protection eq.:* In contrast to the homogeneous case a protection equilibrium does exist. The Nash equilibrium is a free-riding equilibrium where one player protects, and the other does not.

- *Multiple eq. without protection:* If $b_i > c_i$ for all player i individuals will choose to self-insure or remain passive.

In the homogeneous version of the game, we had noted that these Nash equilibria were not reached in a synchronized game with N players, as players would constantly oscillate between free-riding and protecting (see Chapter 2). With heterogeneous players, however, it is possible to reach a Nash equilibrium. Indeed, if the initial protection levels chosen satisfy $\max\{e_{-i}(0)\} > 1 - b_i/pL_i$ for all players *but one*, this last player will be the

only one to secure, while everybody else will defect. Note that there should be only one player choosing to secure for a Nash equilibrium to be reached – as soon as at least two players decide to protect, each will defect in the next round hoping to free-ride on the other protecting players. In other words, if there exists a unique i for which the initial constellation of protection levels satisfies

$$\max\{e_{-i}(0)\} < 1 - b_i/pL_i \,, \tag{3.7}$$

then a Nash equilibrium where all players free-ride on player i is reached as long as $b_i < c_i$. This situation could happen when only one player faces disproportionate losses compared to other players, or her security costs are very small.

Result: *When protection levels are initially randomly set, protection equilibria in the best shot game are increasingly unlikely to happen as the number of players N grows.*

Assume that the initial protection levels, $e_i(0)$ for $1 \leq i \leq N$ are set independently and at random, that is, that they can be expressed as a random variable with cumulative distribution function F. Then for any Player k, the probability that $e_k(0) < 1 - b_i/pL_i$ is simply $F(1 - b_i/pL_i)$. It follows that Eqn. (3.7) is satisfied for Player i with probability $F(1 - b_i/pL_i)^{N-1}$.

Next, we want Eqn. (3.7) to be violated for all players other than i. Eqn. (3.7) is defeated for a given Player k with probability $1 - F(1 - b_k/pL_k)^{N-1}$. Consequently, it is defeated for all Players $j \neq i$ with probability $\prod_{j \neq i}(1 - F(1 - b_j/pL_j)^{N-1})$.

It follows that the probability ρ_i that Eqn. (3.7) is satisfied *only* for Player i is given by

$$\rho_i = F\left(1 - \frac{b_i}{pL_i}\right)^{N-1} \prod_{j \neq i} \left(1 - F\left(1 - \frac{b_j}{pL_j}\right)^{N-1}\right) .$$

Then, the probability that a protection equilibrium can be reached is given by $\sum_i \rho_i$, since the ρ_i's characterize mutually exclusive events. To simplify notations, let $x_i = F\left(1 - \frac{b_i}{pL_i}\right)$. Rearranging terms gives

$$\sum_i \rho_i = \sum_i \prod_{j \neq i} \left(1 - x_j^{N-1}\right) - N \prod_j \left(1 - x_j^{N-1}\right) .$$

Let $k = \arg\max_i \left\{\prod_{j \neq i} \left(1 - x_j^{N-1}\right)\right\}$. Then we have

$$\sum_i \rho_i \leq N \prod_{j \neq k} \left(1 - x_j^{N-1}\right) - N \prod_j \left(1 - x_j^{N-1}\right) ,$$

which gives us, after rearranging

$$\sum_i \rho_i \leq N x_k^{N-1} \prod_{j \neq k} \left(1 - x_j^{N-1}\right) ,$$

which tends to zero as N increases, as soon as $x_k = F(1 - b_k/pL_k) < 1$.

This is notably the case if we assume a function F strictly monotonous increasing on $[0, 1]$, and positive security costs ($b_i > 0$) for all players.

3.3.4 Weakest-target security games

As in the homogeneous case (Chapter 2, Nash equilibria for the weakest-target game are quite different depending on whether or not we are considering that mitigation is possible.

Without mitigation. In the weakest-target game without mitigation, we have reported in Chapter 2 that, in the homogeneous case where $b_i = b$, $c_i = c$, and $L_i = L$, there are no

pure strategy Nash equilibrium. The proof can be extended to the heterogeneous case, as we discuss next.

Let us assume that the minimum protection level over all players is set to $\hat{e} < 1$. Then, we can group players in two categories: those who play $e_i = \hat{e}$, and those who set $e_i > \hat{e}$. By straightforward dominance arguments coming from the description of the payoffs in Eqn. (2.6), players who select $e_i > \hat{e}$ select $e_i = \hat{e} + \varepsilon$, where $\varepsilon > 0$ is infinitesimally small, and $s_i = 0$. Let

$$\varepsilon < \min_i \left\{ \frac{pL_i}{2b_i}(1 - s_i) + \frac{c_i s_i}{2b_i} \right\} .$$

Players who play $e_i = \hat{e}$ would actually prefer to switch to $\hat{e} + 2\varepsilon$. Indeed, the switch in strategies allows a payoff gain of

$$U_i(\hat{e} + 2\varepsilon, 0) - U_i(\hat{e}, s_i) = -2b_i\varepsilon + pL_i(1 - s_i) + c_i s_i > 0 .$$

Hence, this strategy point is not a Nash equilibrium. It follows that the only possible equilibrium point would have to satisfy $e_i = 1$ for all e_i. However, in that case, all players are attacked, which ruins their security investments. All players therefore have an incentive to instead select $e_i = \hat{e} = 0$, which, per the above discussion, cannot characterize a Nash equilibrium.

Result: *In the weakest-target game without mitigation we find that pure Nash equilibria for non trivial values of b_i, p, L_i and c_i do not exist.*

With mitigation. In the weakest-target game with mitigation, we showed that, with homogeneous agents, a full protection Nash equilibrium exists as long as protection costs

are smaller than self-insurance costs (see Chapter 2). An exactly identical proof can be conducted in the heterogeneous case to show that a full protection equilibrium is reached if $b_i < c_i$ for all i.

On the other hand, it only takes one of the players to face high security costs to make this equilibrium collapse. Indeed, if there exists k such that $b_k > c_k$, then Player k will always prefer a full self-insurance strategy ($(e_k, s_k) = (0, 1)$) over a full-protection strategy ($(e_k, s_k) = (1, 0)$). This will immediately lead other players to try to save on security costs by picking $e_i = \varepsilon > 0$ as small as possible. We then observe an escalation as in the unmitigated version discussed above. Hence, heterogeneity actually threatens the (precarious) stability of the only possible Nash equilibrium.

Result: In contrast to the weakest-target game without mitigation we find that a pure Nash equilibrium may exist.

- *Full protection eq.:* If $b_i \leq c_i$ for all agents we find that the full protection equilibrium ($\forall i, (e_i, s_i) = (1, 0)$) is the only possible pure Nash equilibrium.
- If $b_i > c_i$ for *any* agent we can show that no pure Nash equilibrium exists.
- There are no pure self-insurance equilibria.

3.4 Intervention mechanisms

In practice system designers may not be satisfied with the outcomes predicted by non-cooperative game theory. First, equilibria may not be achievable due the complexity of the

games, which limits the understanding and accurate execution of strategies by agents. Second, planners may wish to improve upon the Nash equilibrium security practices. Below we discuss selected intervention strategies in the context of the security games to improve convergence and to achieve certain contribution targets.

Objective 1 - Help agents to identify individually rational strategy: In the five games we consider, agents will incur a loss when adequate protection or self-insurance is amiss. However, the reasons for vulnerability to a loss and eventual compromise are different. For example, in the weakest-target game without mitigation, a security breach is not solely the result of an agent's protection level, but is dependent on the ordering of contribution levels. Individual rationality presumes that agents follow a sophisticated mixed strategy (see also Chapter 2)). However, non-automated agents will only be able to follow such a strategy with difficulty [195]. Even pure strategies might require several periods of convergence [40].

One possible method of intervention to overcome complexity or coordination problems is to offer (non-binding) advice to agents in a security game. For example, Brandts and MacLeod [34] show that players might choose, in a self-enforcing manner, a strategy recommended by an external arbiter. The assignment strongly influences behavior if it does not conflict with another focal principle. In practice, individuals care about who is giving the advice. For example, the suggestion by a computer security company to protect against security breaches with a product of the same brand might be regarded as advertisement and be less influential [131]. Instruments for coordination may also take the form of financial

incentives. For example, a third party or intermediary such as an Internet Service Provider (ISP) can offer a rebate or service discount to its subscribers who demonstrably invest in an adequate level of protection.

System designers have also started to exploit individuals' preferences for status quo settings [122]. If users rarely alter default settings, the importance of choosing secure defaults on the two dimensions of self-insurance and protection is immensely high. For example, the Windows XP firewall, when first introduced to the Microsoft Windows operating system in 2001, was disabled by default. Subsequent to the Blaster worm attack, the default setting was changed to "fully enabled" with Windows Server 2003. As another example, OpenBSD's "secure by default" philosophy means that all non-essential services are disabled by default. This promotes general network security and also encourages users to learn more about potential consequences of making changes to security settings.

Objective 2 - Achieve social optimality: In the weakest-link security game, deviation of a single agent i from a full protection strategy can render all other agents' efforts meaningless or force them to self-insure or be passive. However, this decision by agent i can be individually rational if $b_i \geq c_i$ or $b_i > p_i L_i$. The traditional solution to this situation has been to involve a social planner who can mandate certain protection and self-insurance settings that optimize overall system utility. In Chapter 2, we discussed social optimum outcomes for the homogeneous agents scenario.

A different approach is to allow agents to conduct *binding* pregame communication. For example, consider a scenario in which an agent can propose to another agent that she will

only protect if the other agent agrees to reciprocate. Such two-sided communication can increase the protection contribution in the total effort game since the responding agent can internalize the potential contribution of the proposing agent (rather than merely evaluating $b_i < p_i L_i/N$). In a different scenario, an agent may commit to a high or low protection level and not require reciprocation. Given this one-way pregame communication, the protection contributions by other agents will be unaffected in a total effort game, but impacted in the best shot and weakest-target games. In the best shot game, a one-way message indicating that the sender will shirk can encourage another agent i to take action (if $b_i < p_i L_i$). In the weakest-target game, the same message would signal to other agents that the sender will bear the burden of the attack. This act of altruism is particularly likely if the agent can self-insure at low cost. Finally, binding pregame communication is largely ineffective for the weakest-link game. However, it can help to coordinate on the protection Nash equilibrium that Pareto-dominates other equilibria with a lower $\hat{e}$.

Objective 3 - Overcome free-riding and lack of protection in networks: Free-riding occurs in our games at several points. For example, in the best shot game, agents coordinate so that only one agent exercises maximum protection effort in a protection Nash equilibrium. In the social optimum for the weakest-target games with homogeneous agents, one node (that may invest in self-insurance) will bear the brunt of an attack while others shirk (see Chapter 2). Both outcomes, while maximizing utility, might result in loss of camaraderie and willingness to contribute in the future.

One strategy to probabilistically increase contributions by agents is to leverage the

strategic uncertainty when agents act independently. The coordination problems inherent in the best shot game with heterogeneous agents and in the weakest-target games may lead agents to contribute to protection levels above the social optimum. In practice such an approach might have the merit of increasing general preparedness against different types of attacks. Strategic uncertainty is often a function of network size. For example, in the weakest-target game, agents in small groups will notice the increased interdependency and risk of being the weakest-target. These agents will decide to self-insure their resources more often $((e_i, s_i) = (0, 1))$. That means with increasing network size, we would observe that more individuals contributing to protection. This result stands in contrast to the weakest-link game analysis. In the heterogeneous as well as in the homogeneous game, it becomes increasingly unlikely that contributions to protection are made. Heterogeneity can also moderate protection contributions in a different way. In the homogeneous best shot game, we do not observe individually rational protection contributions at all since agents cannot overcome the associated coordination problems. However, they can achieve higher protection levels when agents have heterogeneous tastes.

Contributions can be increased behaviorally by modifying the framing of a security situation. Framing effects occur when two logically equivalent (but not transparently equivalent) statements describing a problem drive individuals to choose dissimilar options. More specifically, such differences in the presentation may draw a subject's attention to alternative aspects of a decision situation, leading an individual to make mistakes in pursuing her underlying preferences [176]. For example, homogeneous agents can be tempted to con-

tribute in a best shot game if they receive feedback that highlights the uniqueness of their contributions [124]. Similarly, increased protection investments may arise if agents perceive a security situation as more threatening. However, underinvestment can result from resignation with respect to the complexity of the security problem.

3.5 Summary

In Chapter 2 we have studied homogeneous populations of users, where all participants have the same utility function. In practice, the homogeneity assumption is reasonable in a number of important cases, particularly when dealing with very large systems where a large majority of the population have the same aspirations. For instance, most Internet home users are expected to have vastly similar expectations and identical technological resources at their disposal; likewise, modern distributed systems, e.g., peer-to-peer or sensor networks generally treat their larger user base as equals.

However, the fact that the Internet is increasingly used as a common vector between different businesses, and even as a bridge between completely different user bases – for instance, acting as a bridge between mobile phone networks, home users, and e-commerce retailers – emphasizes the need for considering heterogeneous agents, even though the analysis may become far less tractable.

We find several key differences to the analysis with homogeneous agents. For example, we find that in the total effort game stability increases with more pronounced heterogeneity

in the agent population. The existence of a protection equilibrium in the weakest-link game is threatened if only one agent prefers to self-insure or to remain passive. In the best shot game heterogeneous agents can overcome coordination problems more easily, so that a protection equilibrium is now possible, even though reaching this equilibrium grows increasingly unlikely with a larger number of agents participating in the network.

Surprisingly, predictions for pure Nash equilibria of the weakest-target games remain unchanged. However, mixed strategies do now have to take consideration of the heterogeneity of agents.

We discuss several intervention strategies in the context of security games. We note that in each game the challenge to increase security contributions to achieve a particular objective requires a largely different approach. This versatility is confirmed by practical observations which tell us that a "one size fits all strategy" for computer security does not exist.

Chapter 4

Bounded rationality and limited information

Users frequently fail to deploy, or upgrade security technologies, or to carefully preserve and backup their valuable data [121, 162], which leads to considerable monetary losses to both individuals and corporations every year. This state of affairs can be partly attributed to economic considerations. End users may undertake a cost-benefit analysis and decide for or against certain security actions [86, 192]. However, this risk management explanation overemphasizes the rationality of the involved consumers [115]. In practice, consumers face the task to "prevent security breaches within systems that sometimes exceed their level of understanding" [21]. In other words, the amount of information users may be able to acquire and/or to process, is much more limited than is required for a fully rational choice.

We focus on decision-making in different security scenarios that pose significant challenges for average users to determine optimal security strategies, due to *interdependencies* between users (see Chapter 2). Interdependencies occur when the actions of a given user have an effect on the rest of the network, in part or as a whole (externalities), or when the status of a given user impacts that of other users. For example, consumers who open and respond to unsolicited advertisements increase the load of spam for all participants in the network. Similarly, choosing a weak password for a corporate VPN system can facilitate the compromise of many user accounts.

We anticipate the vast majority of users to be *non-expert*, and to apply approximate decision-rules that fail to accurately appreciate the impact of their decisions on others [3]. In particular, in this chapter, we assume non-expert users to conduct a simple self-centered cost-benefit analysis, and to neglect interdependencies. Such users would secure their system only if the vulnerabilities being exploited can cause significant harm or a direct annoyance to them (e.g., their machines become completely unusable), but would not act when they cannot perceive or understand the effects of their insecure behavior (e.g., when their machine is used as a relay to send moderate amounts of spam to third parties).

In contrast, an advanced, or expert user fully comprehends to which extent her and others' security choices affect the network as a whole, and responds rationally. The first contribution of this chapter is to study the strategic optimization behavior of such an expert user in an economy of inexperienced users, using three canonical security games that account for network effects (see Chapter 2).

Our approach to capture bounded-rational behaviors of end-users differs significantly from research on computability and approximation of economic equilibria [91]. We argue that models of security decision-making can benefit from a critical inquiry of the conceptual understanding users have of security problems. While experts and unsophisticated users co-exist in the same networks, they do not share the same knowledge or mental models about security problems and countermeasures [3, 14, 21, 115, 196, 215], or the same identical perfectly rational approaches to solve security issues [2,47].

The second contribution of this chapter is to address how the security choices by users are mediated by the information available on the severity of the threats the network faces. We assume that each individual faces a randomly drawn expected loss. Indeed in practice, different targets, even if they are part of a same network, are not all equally attractive to an attacker: a computer containing payroll information is, for instance, considerably more valuable than an old "boat anchor" sitting under an intern's desk. We study how the decisions of the expert and unsophisticated users differ if all draws are common knowledge, compared to a scenario where this information is only privately known. With this approach we provide two important baseline cases. We further evaluate the value of better information on the total expected payoff of the expert agent. Specifically, we study the following metric: the payoff under complete information divided by the payoff under the incomplete information condition.

By evaluating the value of information for a range of parameters in different security scenarios, we can determine which configurations can accommodate limited infor-

mation environments (i.e., when being less informed does not significantly jeopardize an expert user's payoff), as opposed to configurations where expert users and non-expert users achieve similar outcomes due to a lack of available information. This analysis has implications for network designers that want to avoid undesirable hotspots that penalize users for their lack of information about threats. Similarly, Internet Service Providers or other intermediaries may take influence on the pricing and availability of security technologies to steer users to less harmful parameter configurations.

We first discuss selected work related to our analytic model (Section 4.1). In Section 4.2, we summarize the security games framework we developed in prior work, and detail our assumptions about agent behaviors and information conditions. We present our methodology and formal analysis in Section 4.3. We discuss the results and their implications in Section 4.4. Finally, we close with concluding remarks in Section 4.5.

4.1 Background

In this chapter we conduct a decision-theoretic analysis for a sophisticated (expert) agent who interacts with a group of users that follow a simple but reasonable rule-of-thumb strategy. We structure the remainder of the review of related literature and background information into three selected areas in which we are making a research contribution.

The analysis in this chapter significantly differs from prior decision-theoretic approaches. Gordon and Loeb present a model that highlights the trade-off between perfect and cost-

effective security [87]. They consider the protection of an information set that has an associated loss if compromised, probability of attack, and probability that attack is successful. They show that an optimizing firm will not always defend highly vulnerable data, and only invest a fraction of the expected loss. Cavusoglu *et al.* [42] consider the decision-making problem of a firm when attack probabilities are externally given compared to a scenario when the attacker is explicitly modeled as a strategic player in a game-theoretic framework. Their model shows that if the firm assumes that the attacker strategically responds then in most considered cases its profit will increase. Schechter and Smith [189] consider the decision-theoretic analysis from the perspective of the potential intruder. They highlight several modeling alternatives for attacker behavior and their payoff consequences. The analytic work on security investments and level of penalties for offenses is complemented by empirical research [171, 205].

4.1.1 Bounded rationality

Acquisti and Grossklags summarize work in the area of behavioral economics and psychology that is of relevance for privacy and security decision-making [3]. Users' decisions are not only limited by cognitive and computational restrictions (i.e., bounded rationality), but are also influenced by systematic psychological deviations from rationality.

Recent research has investigated agents that overemphasize earlier costs and benefits at the expense of their future well-being [2, 4, 165]. Christin *et al.* (building on prior economic research [5, 177]) suggest that agents respond near-rationally to the complexity

of networked systems [47]. In their model individuals are satisfied with a payoff within a small margin of the optimal outcome.

Different from the above work that considers all players to act the same, this chapter studies a mixed economy, with expert and non-expert users co-existing. While expert users are as rational as possible, non-expert users deviate from rationality by adopting approximate (rules-of-thumb) decision strategies. In practice, users frequently have to rely on rules-of-thumb when a "quantitative method to measure security levels" is not available [150]. Economic analysis including rule-of-thumb choices have been discussed outside of the security context, e.g., [63] [72] [140].

4.1.2 Limited information

In the context of the value of security information, research has been mostly concerned with incentives for sharing and disclosure. Several models investigate under which conditions organizations are willing to contribute to an information pool about security breaches and investments when competitive effects may result from this cooperation [76, 88]. Empirical papers explore the impact of mandated disclosures [41] or publication of software vulnerabilities [209] on the financial market value of corporations. Other contributions to the security field include computation of Bayesian Nash outcomes for an intrusion detection game [143], and security patrol versus robber avoidance scenarios [168].

We conduct a comparative analysis of strategies and payoffs for a sophisticated agent in a security model when the expected loss from a directed attack is either common or

private knowledge. In particular, we evaluate the influence of the lack of information given different organizational dependencies [212].

4.1.3 Heterogeneous agents

In earlier chapters we analyze both the case of homogeneous (Chapter 2) and heterogeneous agents (Chapter 3). When considering heterogeneous agents, however, we have focused on differences in the costs agents may face. We assumed that users differ in the price they have to pay for protection and self-insurance, and that they have different perceived or actual losses associated with successful (uninsured) security compromises. In the present chapter we analyze the case of agents facing different attack probabilities, that may be a priori unknown to other agents.

Given certain differences in the attractiveness of a particular target the question remains how a defender is able to determine a reasonable estimate of the expected loss. Such a problem far exceeds the scope of this chapter, whose main goal is to study the impact of information (or lack thereof) on security strategies, and we refer the reader to the threat modeling literature. (See [9] for an introduction and references.)

4.2 Decision-theoretic model

In the following we highlight the model variations for the current analysis. In particular, we extend our model to the case of an economy consisting of an expert user and several

unsophisticated users.

4.2.1 Modifications to the basic model

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 [16, 132, 154, 160].

We have further conducted a sensitivity analysis with respect to the discrete choice assumption and find that, for the study in this chapter, 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 Section 4.3). We believe that focusing on these boundary cases is of limited practical applicability, and could even be misleading. For comparison, we refer to the analysis in Chapter 2 where we considered the continuous case in a full information environment.

Expected losses: If an agent 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 $[0, 1]$.[1] In prior chapters, we interpreted the parameter p_i as

[1]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

the probability of a successful attack; however in the present work we prefer to consider the expected loss, $p_i L$, as the primary heterogeneous parameter under consideration. The same familiar notation with p_i considered as a random mitigating factor as opposed to an attack probability facilitates this perspective.

This models the heterogeneous preferences that attackers have for different targets, due to their economic, political, or reputational agenda. The choice of a uniform distribution 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.

4.2.2 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 [3], and field surveys separate between high and low expertise users [201].

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 [56]. In the present context, expert users, for example, have less difficulty to conclude that the goal to avoid censorship points at a best shot scenario, whereas the protection of a corporate network frequently suggests a weakest-link situation (see our discussion in Chapter 2). Accordingly, a sophisticated user correctly

conduct our analysis, it suffices to know only these threshold probabilities, and not the full distribution.

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 [3] [201]. 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 that actually are far below their own expectations. This chapter 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 or are able to keep up with the challenges of novel threats [201]. 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 [122].

4.2.3 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 her 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 naïve users' attack threat, and her own drawn probability of being attacked p_i.

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.

4.3 Analysis methodology

In the remainder of this discussion, we will always use the index i to denote the expert player, and $j \neq i$ to denote the naïve players. For each of the three games, weakest-link, best shot, and total effort, our analysis proceeds via the following five-step procedure.

1. Determine player i's payoff within the game for selected strategies of passivity, full insurance, and full protection. As shown in Chapters 2 and 3 through a relatively

simple analysis, player i can maximize her utility only by relying on (one or more of) these three strategies.

2. Determine the conditions on the game's parameters (b, c, L, N, p_i, and if applicable, p_j for $j \neq i$) under which player i should select each strategy.

3. Determine additional conditions on the game's parameters such that the probability (relative to p_i) of each case, as well as the expected value of p_i within each case can be easily computed.

4. Determine player i's total expected payoff relative to the distribution on p_i and all other known parameters.

5. In the case of complete information, eliminate dependence on p_j for $j \neq i$ by taking, within each parameter case, an appropriate expected value.

Diligent application of this method generates a table recording the total expected payoffs for player i, given any valid assignment to the parameters b, c, L, N. In the process it also generates tables of selection conditions, probabilities, and expected payoffs for each possible strategy; and in the complete information case, gives results for total expected payoffs conditioned on the exact draws of p_j by the other players. The results are presented in Tables A.1–A.15.

In the remainder of this section we illustrate this method by considering, for each step listed above, one game and one parameter case for which we have applied the appropriate step.

Step 1 example: Passivity payoff computation. Let us consider the challenge of determining payoffs for player i's passivity in the best shot game, under the conditions of limited information and parameter constraints $b \leq c$. The general payoff function for the best shot game is obtained by substituting $H(e_i, e_{-i}) = \max(e_i, e_{-i})$ into the general utility function for all games, i.e. $U(i) = M - p_i L(1 - s_i)(1 - H(e_i, e_{-i})) - be_i - cs_i$. Doing this, we obtain $U(i) = M - p_i L(1 - s_i)(1 - \max(e_i, e_{-i}) - be_i - cs_i$. To get the payoff for player i's passivity we plug in $e_i = s_i = 0$ to obtain

$$U_i = \begin{cases} M - p_i L, \text{ if } \max_{j \neq i} e_j = 0 \\ \\ M, \text{ if } \max_{j \neq i} e_j = 1 \end{cases}.$$

Now in the incomplete information case, we do not know any of the p_j for $j \neq i$, so we do not know all the parameters to compute the required payoff. However, since we assume that the p_j are independently and uniformly distributed in $[0, 1]$, we can compute an expected value for this payoff as follows. The probability (over p_j) that none of the other players protect (i.e. that $\max_{j \neq i} p_j < b/L$) is exactly $(b/L)^{N-1}$, and in this case the payoff would be $M - p_i L$. The probability (over p_i) that at least one of the other players protect (i.e. that $b/L \leq \max_{j \neq i} p_j$) is exactly $1 - (b/L)^{N-1}$, and in this case the payoff would be M. Thus the total expected payoff for selecting the passivity strategy is $(b/L)^{N-1}(M - p_i L) + (1 - (b/L)^{N-1})M$, which simplifies to $M - p_i L(b/L)^{N-1}$. We record this as the payoff result for passivity in the incomplete game, with $b \leq c$, as can be seen in Table A.6.

Step 2 example: Strategy selection. Let us next consider the challenge of determining parameter conditions under which we should select player i's strategy in the weakest-link game. Since this is a second step, consider the game payoffs in Table A.1 as given. We are interested in determining player i's most strategic play for any given parameter case. Select for consideration the case $b \leq c$ with incomplete information. (Note: this is the most difficult case for this game).

To determine the optimal strategy for player i, we must select the maximum of the quantities Passivity: $M - p_i L$, Insurance: $M - c$, and Protection: $M - b - p_i L(1 - (1 - b/L)^{N-1})$. We should choose passivity if it is better than insurance or protection, i.e. $M - p_i L > M - c$ and $M - p_i L > M - b - p_i L(1 - (1 - b/L)^{N-1})$. We should choose insurance if it is better than passivity or protection, i.e. $M - c \geq M - p_i L$ and $M - c > M - b - p_i L(1 - (1 - b/L)^{N-1})$. We should choose protection if it is better than passivity or insurance, i.e. $M - b - p_i L(1 - (1 - b/L)^{N-1}) \geq M - p_i L$ and $M - b - p_i L(1 - (1 - b/L)^{N-1}) \geq M - c$.

Re-writing the above inequalities as linear constraints on p_i , we choose passivity if $p_i \leq c/L$ and $p_i \leq \frac{b}{L(1-(1-b/L)^{N-1})}$; we choose insurance if $p_i > c/L$ and $p_i > \frac{c-b}{L(1-(1-b/L)^{N-1})}$; and we choose protection if $\frac{c-b}{L(1-(1-b/L)^{N-1})} \leq p_i \leq \frac{b}{L(1-(1-b/L)^{N-1})}$.

For simplicity of computation, we would like to have our decision mechanism involve only a single inequality constraint on p_i. To obtain this it is necessary and sufficient to determine the ordering of the three terms: $\frac{c}{L}$, $\frac{b}{L(1-(1-b/L)^{N-1})}$, and $\frac{c-b}{L(1-(1-b/L)^{N-1})}$.

It turns out that there are only two possible orderings for these three terms. The single inequality $c < \frac{b}{(1-b/L)^{N-1}}$ determines the ordering: $\frac{c}{L} < \frac{c-b}{L(1-(1-b/L)^{N-1})} < \frac{b}{L(1-(1-b/L)^{N-1})}$; while the reverse inequality $\frac{b}{(1-b/L)^{N-1}} \leq c$ determines the reverse ordering on all three terms. This observation suggests we should add sub-cases under $b \leq c$ depending on which of these two inequalities holds. See Table A.2.

Within each new sub-case the criterion for selecting the strategy that gives the highest payoff can now be represented by a single linear inequality on p_i. If $c \leq \frac{b}{(1-b/L)^{N-1}}$, then passivity wins so long as $p_i < c/L$; (because the new case conditions also guarantee $p_i < \frac{b}{L(1-b/L)^{N-1}}$). Similarly insurance wins if $p_i \geq c/L$. Protection never wins in this case because we cannot have $\frac{c-b}{L\left(1-(1-b/L)^{N-1}\right)} \leq p_i \leq \frac{b}{L\left(1-(1-b/L)^{N-1}\right)}$ when we also have $\frac{b}{(1-b/L)^{N-1}} < \frac{c-b}{L\left(1-(1-b/L)^{N-1}\right)}$. The computations for the case $\frac{b}{(1-b/L)^{N-1}} < c$ are similar; the results are recorded in Table A.2.

Step 3 example: Case determination. Now, consider the challenge of determining additional constraints on parameters in the total effort game, so that in any given case, the total payoffs can be represented by simple closed form functions of the game's parameters. Since this is a third step, we assume the second step has been diligently carried out and consider the strategy conditions given in Table A.12 as given. For brevity, we consider only the incomplete information case under the assumption $b \leq c$.

To illustrate the problem we are about to face, consider the condition for selecting passivity in the incomplete game and case: $b + b^2(N-1)/L < c$. The condition here is that $p_i < bN/L$. This condition is possible if and only if $bN < L$. The case conditions deter-

mined thus far do not specify which of these is the case; so for subsequent computations, we will need to know which it is, and therefore must consider the two cases separately.

Going beyond this particular example, there are several other values in this table where a similar phenomenon occurs. In particular, we need new cases to determine whether each of the following relations holds: $bN/L \leq 1$, $\frac{c}{b+(L-b)/N} \leq 1$, and $\frac{c-b}{b-b/N} \leq 1$. (See Table A.12). To combine these with previous cases in a way that avoids redundancy, we rewrite the conditions involving c as linear inequalities on c; obtaining $c \leq b + (L-b)/N$ and $c \leq 2b - b/N$.

We are thus left to reconcile these additional cases with the current cases $b \leq c \leq b + \frac{b^2}{L}(N-1)$ and $b + \frac{b^2}{L}(N-1) < c$. To do this efficiently we must know the order of the terms $b + \frac{L-b}{N}$, $2b - \frac{b}{N}$, and $b + \frac{b^2}{L}(N-1)$. Fortunately, it turns out that there are only two possible orderings on these terms; and furthermore, which of the two orderings it is depends on the relation $bN < L$ which we already needed to specify as part of our case distinctions. If $bN \leq L$, then $b + \frac{b^2}{L}(N-1) \leq 2b - \frac{b}{N} \leq b + \frac{L-b}{N}$ and if $bL > N$, then the reverse relations hold.

Assuming limited information, $b \leq c$, and dividing all cases according to $bN \leq L$, it requires a total of 5 cases to determine all important relationships among important parameters for this game. We may have $bN \leq L$ and $b \leq c \leq b + \frac{b^2}{L}(N-1)$; $bN \leq L$ and $b + \frac{b^2}{L}(N-1) < c < 2b - \frac{b}{N}$; $bN \leq L$ and$2b - \frac{b}{N} \leq c$; $bN > L$ and $c \leq b + \frac{L-b}{N}$; and $bN > L$ and $b + \frac{L-b}{N} < c$. For reference, see table A.15.

Step 4 example: Total payoff computation. Let us determine the total expected payoff for the expert player with incomplete information in the best shot game with $b \leq c$. As intermediate steps we must compute the probability that each strategy is played, along with the expected payoff for each strategy. The total payoff is then given by (Probability of passivity · Expected payoff for passivity) + (Probability of insurance · Expected payoff for insurance) + (Probability of protection · Expected payoff for protection).

The expected probability of passivity in this case is 1, with a payoff of $M - p_i L(b/L)^{N-1}$. To get an expected payoff, we compute the expected value of p_i within this case. Since there is no constraint on p_i and it is drawn from a uniform distribution its expected value is $1/2$. Thus the expected payoff for this case is $M - (L/2)(b/L)^{N-1}$. The total expected payoff is thus $M - (L/2)(b/L)^{N-1}$.

Step 5 example: Eliminating dependencies on other players. Consider the challenge of examining the total expected payoff for player i, who has complete information, and rewriting this payoff in a way that is still meaningful as an expected payoff, but does not depend on any p_j for $j \neq i$. The reason we want to do this last step is so we can compare complete information payoff results with incomplete information payoff results. We can only do this if the direct dependence on privileged information is removed from the complete information case payoff. Our method of information removal involves taking an appropriate expected value.

For this example we consider the best shot game with complete information in the case $b \leq c$. Since this is a fifth step, we should assume that the fourth step – computing the

expected payoff for player i as a function of parameters that may include p_j for $j \neq i$ – has been accomplished.

Indeed, by following steps 1–4, the total expected payoffs for player i (conditioned on other players) in the case $b \leq c$ can be derived, subject to two additional sub-cases. If $\max_{j \neq i} p_j \leq b/L$, then the expected payoff is $M - c + c^2/L$; while if $b/L < \max_{j \neq i} p_j$, then the expected payoff is $M - b + b^2/L$.

To generate an appropriate "a posteriori" expected payoff over all choices of p_j, we compute the probability (over choice of p_j) that we are in case $\max_{j \neq i} p_j \leq b/L$ times the payoff for that case, plus the probability (over p_j) that we are in the case $b/L < \max_{j \neq i} p_j$ times the payoff for that case. We obtain $(b/L)^{N-1} \cdot [M - c + c^2/L] + \left[1 - (b/L)^{N-1}\right] \cdot [M - b + b^2/L]$. The end result is $M - b\,(1 - b/2L)\,(b/L)^{N-1}$. See Table A.10.

4.4 Results

4.4.1 Strategies and payoffs

Our results provide us with insights into security decision-making in networked systems. We can recognize several situations that immediately relate to practical risk choices. We start with basic observations that are relevant for all three games, before discussing the different games and information conditions in more detail.

General observations applicable to all three security games. Every scenario involves simple cost-benefit analyses for both sophisticated and naïve agents [86]. 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 $c < b$ 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. In this case self-insurance has a more subtle effect on the payoffs.

There are important differences between the two agent types. The expert agent considers the strategic interdependencies of all agents' choices. For example, consider $b < p_i L$ and $b \leq c$ (that is, protection would be the preferred choice in the absence of interdependencies) then the expert agent sometimes rather prefers to self-insure, or to remain passive while naïve agents would always protect without further consideration. The more nuanced strategies of the expert agent attest to her realization that the group protection goal is sometimes not achievable. Note that we model the agents' incentives to invest in protection in canonical scenarios when security is critically dependent on a group effort (see descriptions for tightly coupled games in Chapter 2). For example, with full cooperation of all agents the incentives to send unsolicited bulk email could be significantly reduced. However, if naïve users open, respond or otherwise interact with spam then other users have

little choice but some form of mitigation of the resulting inconveniences. Otherwise, the expert agent will commonly invest in security for a resource when its safety is not subject to peers' (in)actions (i.e., if $N = 1$).

If $b > p_j L$ for some agents j, then the naïve users do not fully internalize how the inactions of those agents can impact system-wide security. This naïveté is coming back to haunt them. In fact, surveys of average end users' security experiences show that 66 percent lost data permanently due to lacking backup provisions [121]. Similarly, 54 percent have had their computers infected by a network-propagated malicious code [162]. For example, the success of the Conflickr/Downadup worm is dependent on users not applying available patches to their operating system [199].

The naïve agents face a payoff reduction as a result of their limited understanding of correlated threats, but even the sophisticated agent can experience a similar payoff reduction due to limited information. On the one hand, she might invest in self-protection or self-insurance when it is not necessary because the naïve agents collectively or individually secured the network. On the other hand, she may fail to take a security action when a (relatively unexpected low probability) breach actually occurs. It is important to mention that she acted rationally in both situations, but these additional risks remain.

Basic payoffs for different security actions: We can immediately observe that the additional risk due to limited information results from different mechanisms for each security scenario. In the weakest-link game (Table A.1) we find that self-protection carries a risk for the expert agent with limited information that at least one naïve agent chooses not to pro-

tect. This would result in a break-down of system security and a waste of self-protection expenditure. In contrast, in the best shot game (Table A.6) the investment in preventive action always secures the network but with limited information this may be a duplicative effort. In the total effort game these risks are more balanced (Table A.11). The expert can add or withhold her N-th part of the total feasible security contribution. Depending on the cost of security she has to estimate the expected number of naïve contributors K in order to respond adequately.

Conditions for choice between different security actions: In the weakest-link game and complete information, the expert agent can utilize the lowest attack probability that any naïve agent has drawn. If this value is below the required threshold for protection, (i.e. if $\min_{j \neq i} p_j < b/L$), then the sophisticated agent will never protect. Otherwise, depending on her own draw she will make or break a successful defense. Under incomplete information she has to consider the likelihood $(1 - b/L)^{N-1}$ that all naïve agents protect. In all cases there is now a residual likelihood that she might self-insure. See Table A.2.

In the best shot game the fully informed expert can simply determine the highest likelihood of being attacked for any naïve agent to decide whether she should contribute to system protection. With full or limited information, it is obvious that she will only have to contribute very rarely, and can mostly rely on others' efforts. Nevertheless, it is surprising to find that in the incomplete information scenario the expected payoff from passivity always dominates the expected payoff for protection, even when the expected loss is near total ($p_i \sim 1$). The sophisticated user with limited information will never protect. Under

neither information condition is it optimal to self-insure if $b \leq c$. See Table A.7 for details.

Next consider the total effort game (Table A.12). Under full information with $b \leq c$, all conditions depend non-trivially on K, the number of contributors to protection. Under incomplete information the expert must compute the expected value of K, which is $(1 - b/L)(N-1)$. The case differences between complete and incomplete conditions reflect the replacement of K with $E[K]$, and subsequent simplification. In all cases, the critical factor for the decision to protect is whether the potential loss is N times greater than the cost of protection (i.e. $p_i L \geq bN$).

Case boundaries for choice between different security actions: In Figure 4.1, we plot the cases used to record total expected payoffs for the expert agent in Tables A.5, A.10, and A.15. The associated results for the probabilities of self-protection, self-insurance and passivity (within each case) are Tables A.3, A.8, and A.13.

In the weakest-link game only cases 3 and 4 allow for investments in self-protection. We find that increasing the number of agents, N, results in a shrinkage of both cases 3 and 4 to the benefit of case 2. In contrast, the determination of case boundaries in the best shot game is independent of the size of the network. Finally, in the total effort game only cases 3 and 4 allow for rational self-protection investments. Again an increase in the network size reduces the prevalence of these cases (since $bN \leq L$ is a necessary condition).

Payoffs: Tables A.5, A.10, and A.15 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

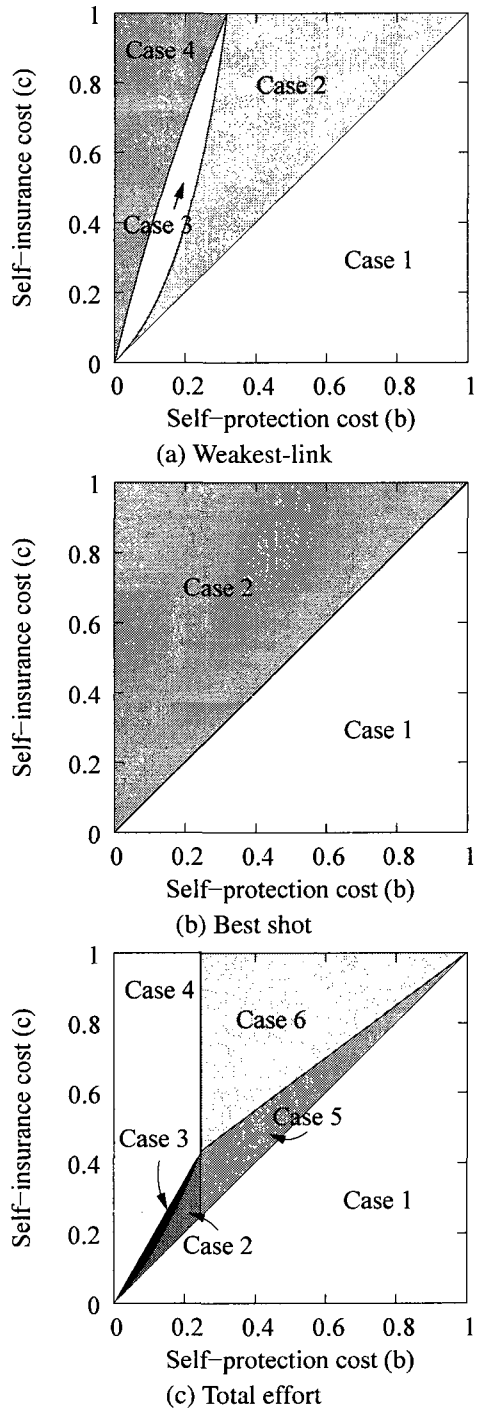

Figure 4.1: **Strategy boundaries in the incomplete information scenario for the sophisticated player.** The different cases refer to Tables A.5, A.10 and A.15. $L = M = 1$ and $N = 4$ in this set of examples.

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 $b \leq c$.

Consider the graphs in Figure 4.2. We plot the payoff functions for sophisticated agents types under the different information conditions, as well as the payoff output for the non-expert agent. 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 in protection. This reflects the fact that the naïve agent ignores the self-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.

4.4.2 Value of information

From a system design perspective it is important to select parameter settings (e.g., making available specific security technologies) that maximize user utility and are robust to changes in the environment. The security games we analyze in this chapter are a significant challenge in both aspects. In particular, from Figure 4.2 we can infer that the penalty for the lack of complete information about attack threats can be highly variable depending on the system parameters. We argue that the reduction of this disparity should be an important design goal. To further this goal we propose a mathematical formulation to measure the value of better information. We then apply this metric to the analysis of the three canonical security games.

Definition: We are interested in a mathematical measure that allows us to quantify the payoff loss due to incomplete information for sophisticated agents, that can be applied to a variety of decision-theoretic scenarios. It is nontrivial to arrive at a definitive answer for this problem statement, therefore, we consider our analysis as a first step towards this goal. We define the value of information metric as the ratio:

$$\frac{\text{Expected payoff in the complete information environment}}{\text{Expected payoff in the incomplete information environment}}$$

Observations: Consider Figure 4.3 which gives, for all three security games, a heat plot for the value of better information over all choices of b and c with L, M, N fixed at $L = M = 1$ and $N = 4$. The most remarkable feature of these graphs are the different hotspot regions. In the weakest-link game we find that higher ratios are to be found

within the boundaries of cases 3 and 4. Both cases allow for self-protection in the presence of incomplete information and therefore balance the various risks more directly than the remaining cases. (Case 1 and 2 associate zero probability with self-protection.)

In the best shot scenario the peak region is located trivially within the boundaries of case 2. We know that the expert player will never protect under incomplete information but is subject to the residual risk of a system-wide security failure. For $N = 4$ the likelihood of such a breakdown is already very small, and decreases with N. Still this outcome is feasible and most pronounced for protection costs that are about a half to two-thirds of the loss, L. For higher b the disincentive of buying self-protection and the potential loss are relatively balanced resulting in a lower penalty for limited information.

In the total effort game we observe multiple hotspot regions. Cases 4 and 6 are unaffected by limited information. They are characterized by the absence of self-insurance as a feasible strategy. This eases the decision-making problem of the expert, and reduces the likelihood of a misspent security investment.

4.5 Summary

In our work we emphasize that security decision-making is shaped by the structure of the task environment as well as the knowledge and computational capabilities of the agents. To that effect, we study security investment choices in three canonical scenarios. Decisions are made from three distinct security actions (self-protection, self-insurance or passivity)

to confront the security risks of weakest-link, best shot and total effort interdependencies. In these environments, we investigate the co-habitation of a single fully rational expert and $N-1$ naïve agents. The naïve agents fail to account for the decisions of other agents, and instead follow a simple but reasonable self-centered rule-of-thumb. We further study the impact of limited information on rational agents' choices. To guide the reader through our analysis, we provide a detailed overview and examples of our methodology to compare strategies and payoffs.

We find that in general, the naïve agents match the payoff of the expert when self-insurance is cheap, but not otherwise. Even with limited information, the sophisticated agent can generally translate her better structural understanding into decisions that minimize wasted protection investments, or an earlier retreat to the self-insurance strategy when system-wide security is (likely) failing.

A notable exception is the weakest-link game with incomplete information, where the payoff of the sophisticated agent degrades to that of the naïve agent as self-insurance becomes more expensive. A practical implication of this result is that, in corporate network access control, having a lot of information about the various potential vulnerabilities that may exist at network access points actually only marginally enhances security; the key factor is whether self-insurance (e.g., data backups) provide adequate security or not. When some items, such as trade secrets, cannot be self-insured, they simply should not be stored on a publicly accessible network. Common sense tells us that much; a contribution of this chapter is to provide a mathematical foundation to justify such policy recommendations.

Our analysis also shows that an expert user never provides a positive improvement to system-wide security (in comparison to her replacement by an unsophisticated agent). While our expert agent is rational, she is not benevolent.[2] Instead she acts selfishly, and the set of scenarios for which protection is her best option is always a subset of the set of scenarios for which the naïve agent chooses protection. In other words, assuming that competent CISOs may be interested in enhancing security at all costs may be a tall order; they may, in fact, be much more interested in finding optimal security investments, which may not result in improved security.

To complement our study we are interested in studying properties of a network with varying fractions of expert to naïve users. Further, we want to address the desire of some computer experts to sacrifice individual resources to improve system resilience to attacks, by introducing *benevolent* agents. As discussed above, our analysis thus far evidences the need for such benevolent agents. As a practical example, censorship-resilient networks are run by volunteers; without these benevolent participants, the whole network collapses. This chapter shows that there is little hope for strong security if all participants are either naïve, or selfish.

To analyze the impact of the different information conditions we have proposed a new mathematical formalization. We measure the value of complete information as the ratio of the payoff in the complete information environment to the payoff in the incomplete information environment. Our analysis of Figure 4.3 is a first step in that direction, however, a

[2]There is a related debate on vigilante defenders in the computer science literature [20, 54, 117].

more formal analysis is deferred to Chapter 5.

Finally, a system designer is not only interested in the payoffs of the network participants given different information realities (e.g., due to frequent changes in attack trends). He is also concerned with how well-fortified the organization is against attacks. To that effect we plan to include a more thorough presentation of the parameter conditions that cause attacks to fail due to system-wide protection, and when they succeed (due to coordination failures, passivity, and self-insurance).

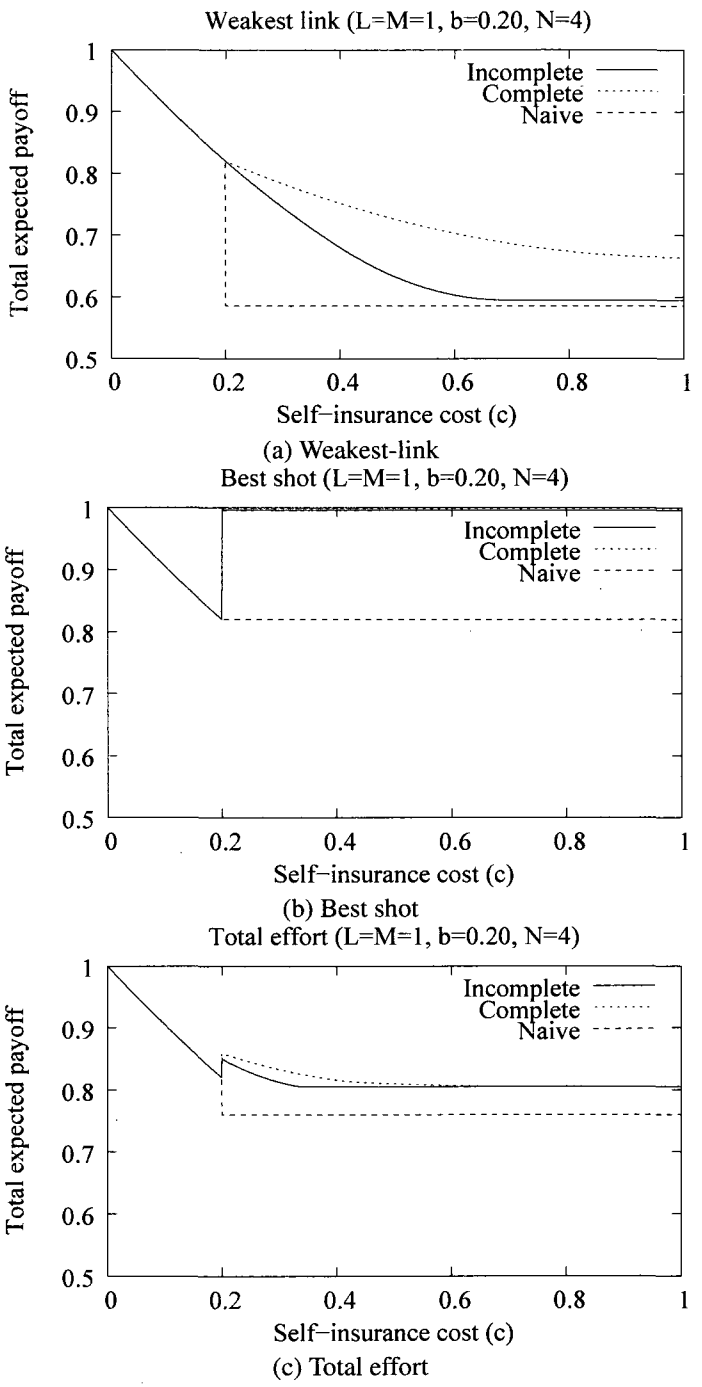

Figure 4.2: **Total expected payoffs for the strategic player under different information conditions, compared with that of the naïve agents.** $L = M = 1$, $N = 4$, and b is fixed to $b = 0.20$ in this set of examples.

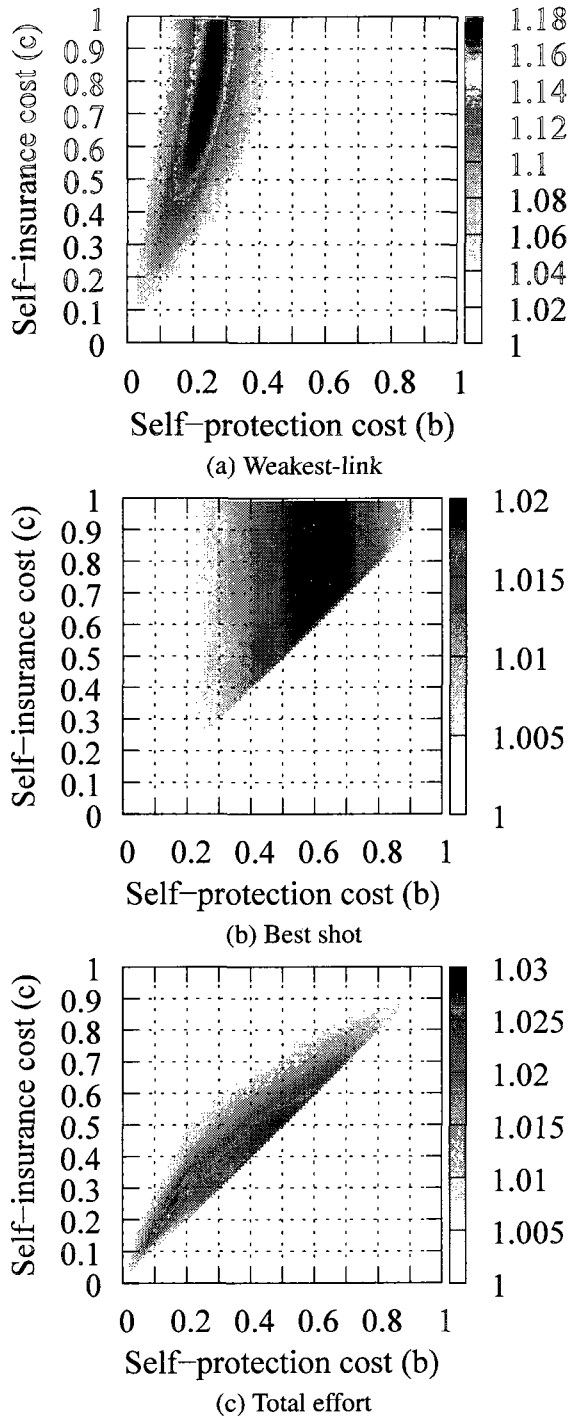

(a) Weakest-link

(b) Best shot

(c) Total effort

Figure 4.3: **The value of information for the three games.** $L = M = 1$, $N = 4$ in this set of examples.

Chapter 5

The price of uncertainty

The lack of information about security threats, response mechanisms, and associated expected losses and cost has long been recognized as important 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 [92]. Swire adds that deterring fraudsters and criminals online is hampered if we cannot correctly aggregate their offenses across different jurisdictions [207].

The question arises: how much can defenders 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 [76,88]. 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 disclosures [41] or publication of software vulnerabilities [209] on the financial market value of corporations. Some work shows that the information gathering or disclosure effect is not always unambiguously positive or negative, respectively. Choi *et al.* [46], for example, present another model on mandatory disclosure of security vulnerabilities. They present scenarios in which disclosure is and is not welfare-improving.

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 [206].

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 *infor-*

mation structure of the scenario that makes explicit the nature of the utility of information to agents [135].

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 Papadimitriou's "price of anarchy" [129].[2] We consider *difference*, *payoff-ratio*, and *cost-ratio* sub-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 a specific example. In our model for security choices, we assume 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 (see our model in Chapter 4).

We aim to understand the importance of the price of uncertainty across different canonical cases of interdependency: best shot, weakest-link and total effort (see descriptions in Chapter 2). Further, in Chapter 4, we distinguish between the roles of a fully rational expert agent and naïve end users. The latter conduct a simple self-centered cost-benefit analysis, and neglect interdependencies. In the current chapter, we analyze the price of uncertainty from the perspective of the expert agent that fully comprehends the benefits of information

[1] After our initial proposal of the price of uncertainty [97], Balcan *et al.* published a research study in which they define the price of uncertainty as the degree that small fluctuations in costs impact the result of natural best-response and improved-response dynamics [17].

[2] In the context of security, several researchers have brought forward analyses of the price of anarchy [118, 138, 146].

in the context of the interrelationship with other naïve users (see Chapter 4). This allows us to make a general observation. The value of information for the expert agent is always weakly positive [135] 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 a simple difference metric of the price of uncertainty increases linearly in losses, L, and decreases superlinearly 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. We additionally present a ratio metric that is strictly decreasing in N. Interestingly, we demonstrate that this metric is independent of the magnitude of potential losses, L. Finally, our third purely cost-based metric suggests that it might lead to misleading conclusions about the necessity of information gathering by overemphasizing 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 [143], and security patrol

versus robber avoidance scenarios [168] can benefit from a substantiation of the significance of informational assumptions by studying the price of uncertainty.

In Section 5.1, we draw the connection to our security games framework developed in earlier chapters. We present the different metrics for the price of uncertainty and describe our analysis methodology in Section 5.2. We conduct our analysis and discuss the results in Section 5.3. Finally, we close with a discussion and concluding remarks in Section 5.4.

5.1 Decision-theoretic model

Our study of the *price of uncertainty* is conducted within the context of a decision-theoretic security analysis that we have completed in Chapter 4. We studied the decision-making of a sophisticated (expert) agent who interacts with a group of users that follow a simple but reasonable rule-of-thumb strategy. We refer to Chapter 4 for a discussion of the setup of the model.

In Chapter 4 we provide the basic results for the three canonical scenarios and the decision-making of the expert and naïve agents detailed in.

Our starting point for the current analysis are the total payoff results in Tables A.5, A.10, and A.15. We will 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.

5.2 Price of uncertainty metrics

In Chapter 4 we discuss 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. the cost of protection b,
2. the cost of insurance c,
3. the magnitude of potential losses L,
4. the initial endowment M, and
5. the 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. In this work 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.

5.2.1 Three metrics for the price of uncertainty

For each of our three security games, best shot, weakest-link, and total effort, we define metrics for the price of uncertainty having the following three forms:

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

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

5.2.2 Discussion of the definitions

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 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.

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 [129]. By analogy, we defined the price of uncertainty as the worst case payoff-ratio of the expert with complete information to the expert with incomplete information, with the worst case taken over all possible prices of protection and 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.

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; and

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.

5.3 Analysis

In this section, we analyze the price of uncertainty as defined by each of our three metrics in each of our three 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 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.

5.3.1 Best shot game

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:

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

$$- \text{Best Shot Expected Payoff Incomplete}(b, c, L, M, N)]$$

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.

$$\begin{aligned}
&\text{Best Shot Expected Payoff Complete}(b, c, L, M, N) \\
&- \text{Best Shot Expected Payoff Incomplete}(b, c, L, M, N) \\
&= \left[M - b\left(1 - \frac{b}{2L}\right)\left(\frac{b}{L}\right)^{N-1}\right] - \left[M - \frac{L}{2}\left(\frac{b}{L}\right)^{N-1}\right] \\
&= \left(\frac{L}{2} - b + \frac{b^2}{2L}\right)\left(\frac{b}{L}\right)^{N-1} \\
&= \frac{L^2 - 2bL + b^2}{2L}\left(\frac{b}{L}\right)^{N-1} \\
&= \frac{(L-b)^2}{2L}\left(\frac{b}{L}\right)^{N-1}
\end{aligned}$$

This expression is maximized as a function of b when its partial derivative with respect to b is zero. So we compute:

$$
\begin{aligned}
0 &= \left(-1+\frac{b}{L}\right)\left(\frac{b}{L}\right)^{N-1} + \frac{(L-b)^2}{2L}(N-1)\left(\frac{b}{L}\right)^{N-2}\cdot\frac{1}{L} \\
0 &= -\left(\frac{b}{L}\right)^{N-1} + \left(\frac{b}{L}\right)^{N} + \frac{L^2(N-1)}{2L^2}\left(\frac{b}{L}\right)^{N-2} - \frac{2L(N-1)}{2L}\left(\frac{b}{L}\right)^{N-1} \\
&\quad + \frac{(N-1)}{2}\left(\frac{b}{L}\right)^{N} \\
0 &= \frac{N+1}{2}\cdot\left(\frac{b}{L}\right)^{N} - N\cdot\left(\frac{b}{L}\right)^{N-1} + \frac{N-1}{2}\cdot\left(\frac{b}{L}\right)^{N-2} \\
0 &= \left(\frac{N+1}{2}\right)\left(\frac{b}{L}\right)^{N-2}\left(\left(\frac{b}{L}\right)^{2} - \frac{2N}{N+1}\left(\frac{b}{L}\right) + \frac{N-1}{N+1}\right) \\
0 &= \left(\frac{N+1}{2}\right)\left(\frac{b}{L}\right)^{N-2}\left(\frac{b}{L}-1\right)\left(\frac{b}{L}-\frac{N-1}{N+1}\right)
\end{aligned}
$$

The expression is zero if and only if

$$b = 0 \text{ or } b = L \text{ or } b = L\cdot\left(\frac{N-1}{N+1}\right).$$

From the second derivative test we find that $b = 0$ and $b = L$ give local minima, hence the maximizing value of this expression for $b \in [0, L]$ occurs when $b = L\cdot\frac{N-1}{N+1}$. Figure 5.1 plots this maximizing b as a function of N. For the price of uncertainty, we have

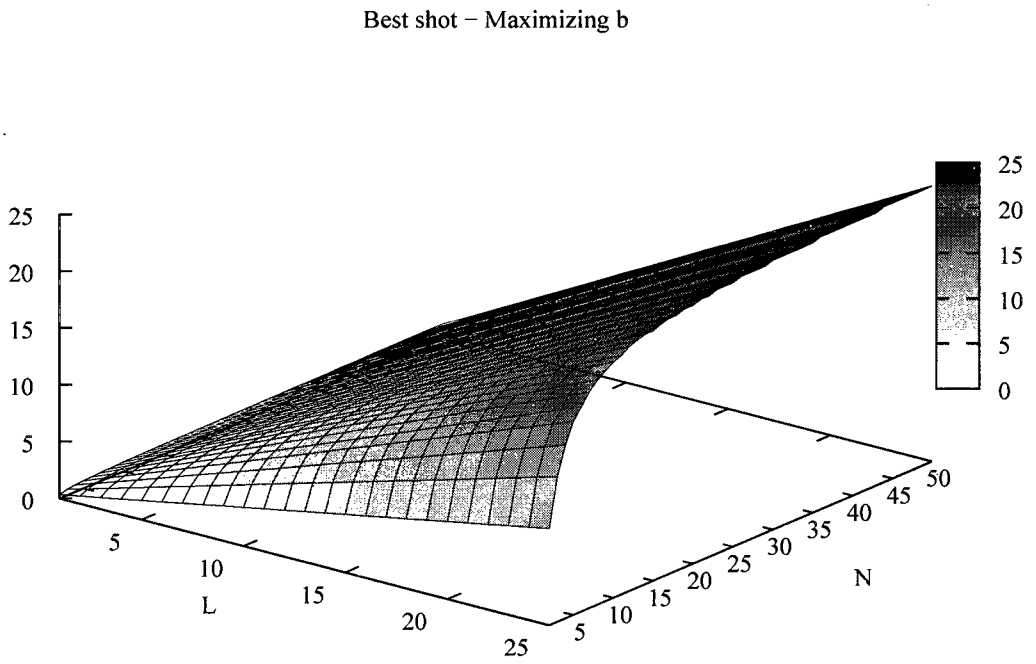

Figure 5.1: **Best shot – Difference metric: Maximizing b for $BPoU_1(L, N)$.**

$$
\begin{aligned}
&BPoU_1(L,N) \\
&= \max_{b,c\in[0,L]} [\text{Best Shot Expected Payoff Complete}(b,c,L,M,N) \\
&- \text{Best Shot Expected Payoff Incomplete}(b,c,L,M,N)] \\
&= \max_{b\in[0,L]} \left[\frac{(L-b)^2}{2L} \left(\frac{b}{L}\right)^{N-1} \right] \\
&= \frac{\left(L - L\cdot\left(\frac{N-1}{N+1}\right)\right)^2}{2L} \left(\frac{L\cdot\left(\frac{N-1}{N+1}\right)}{L}\right)^{N-1} \\
&= \frac{L^2\left(1-\left(1-\frac{2}{N+1}\right)\right)^2}{2L} \left(\frac{N-1}{N+1}\right)^{N-1} \\
&= \frac{2L}{(N+1)^2} \left(\frac{N-1}{N+1}\right)^{N-1} \\
&= 2L\cdot\frac{(N-1)^{N-1}}{(N+1)^{N+1}}
\end{aligned}
$$

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 5.2 shows a graph of the metric $BPoU_1$ as a function of N and L.

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 (unless the losses are quite high – approaching the square of the number of players).

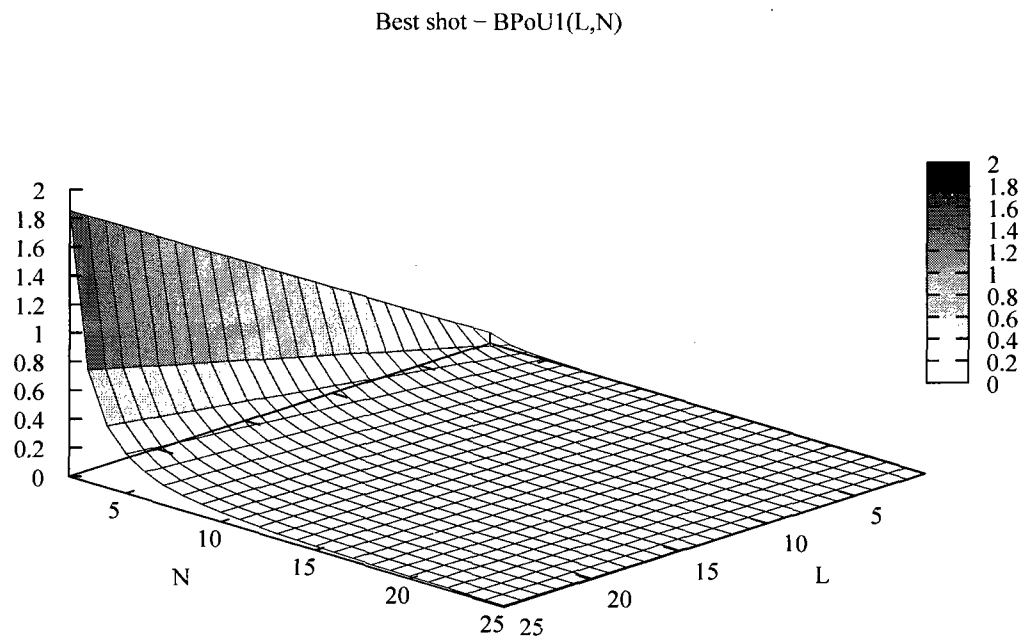

Figure 5.2: **Best shot – Difference metric:** $BPoU_1(L, N)$**.** The metric grows linearly in the potential loss L for a fixed network size N, and decreases inverse-quadratically in the network size N for a fixed loss L.

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

$$\max_{b,c\in[0,L]} \left[\frac{\text{Best Shot Expected Payoff Complete}(b, c, L, L, N)}{\text{Best Shot Expected Payoff Incomplete}(b, c, L, L, N)} \right] \tag{5.1}$$

$$\begin{aligned}
&BPoU_2(L, N) \\
&= \max_{b,c\in[0,L]} \left[\frac{\text{Best Shot Expected Payoff Complete}(b, c, L, L, N)}{\text{Best Shot Expected Payoff Incomplete}(b, c, L, L, N)} \right] \\
&= \max_{b\in[0,L]} \frac{L - b\left(1 - \frac{b}{2L}\right)\left(\frac{b}{L}\right)^{N-1}}{L - \frac{L}{2}\left(\frac{b}{L}\right)^{N-1}} \\
&= \max_{b\in[0,L]} \frac{L\left(1 - \frac{b}{L}\left(1 - \frac{b}{2L}\right)\left(\frac{b}{L}\right)^{N-1}\right)}{L\left(1 - \frac{1}{2}\left(\frac{b}{L}\right)^{N-1}\right)} \\
&= \max_{B\in[0,1]} \frac{\left(1 - B\left(1 - \frac{B}{2}\right)B^{N-1}\right)}{\left(1 - \frac{1}{2}B^{N-1}\right)} \\
&= \max_{B\in[0,1]} \frac{1 - B^N + \frac{1}{2}B^{N+1}}{1 - \frac{1}{2}B^{N-1}} \\
&= \max_{B\in[0,1]} 1 + \frac{-B^N + \frac{1}{2}B^{N+1} + \frac{1}{2}B^{N-1}}{1 - \frac{1}{2}B^{N-1}} \\
&= \max_{B\in[0,1]} 1 + \frac{\frac{1}{2}B^{N-1}(1 - B)^2}{1 - \frac{1}{2}B^{N-1}}
\end{aligned}$$

To compute the maximum, we take the derivative with respect to B and set it equal to zero. We get:

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

$$0 = \left(\frac{N-1}{2}B^{N-2}(1-B)^2 - B^{N-1}(1-B)\right)\cdot\left(1-\frac{1}{2}B^{N-1}\right)$$
$$+ \frac{N-1}{4}B^{2N-3}(1-B)^2$$

$$0 = \frac{N-1}{2}B^{N-2}(1-B)^2 - B^{N-1}(1-B) - \frac{N-1}{4}B^{2N-3}(1-B)^2$$
$$+ \frac{1}{2}B^{2N-2}(1-B) + \frac{N-1}{4}B^{2N-3}(1-B)^2$$

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

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

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

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

Both $B = 1$ and $B = 0$ are roots of this equation, but when put back into the maximizing formula, they each give the global minimum value of 1. It remains to find a solution to this derivative equation for B in $(0, 1)$. We know there is such a root because the value of $B^N - B(N+1) + N - 1$ is positive at $B = 0$ and negative at $B = 1$. Unfortunately, this root, which must maximize the $BPoU_2$ metric, is not generally expressible in closed form for $N \geq 5$. Figure 5.3 plots a graph of the maximizing $b = LB$ as a function of N and L.

It follows from our derivations that this measure of the price of uncertainty does not depend on L. Figure 5.4 plots $BPoU_2$ as a function of N. As can be seen from the graph,

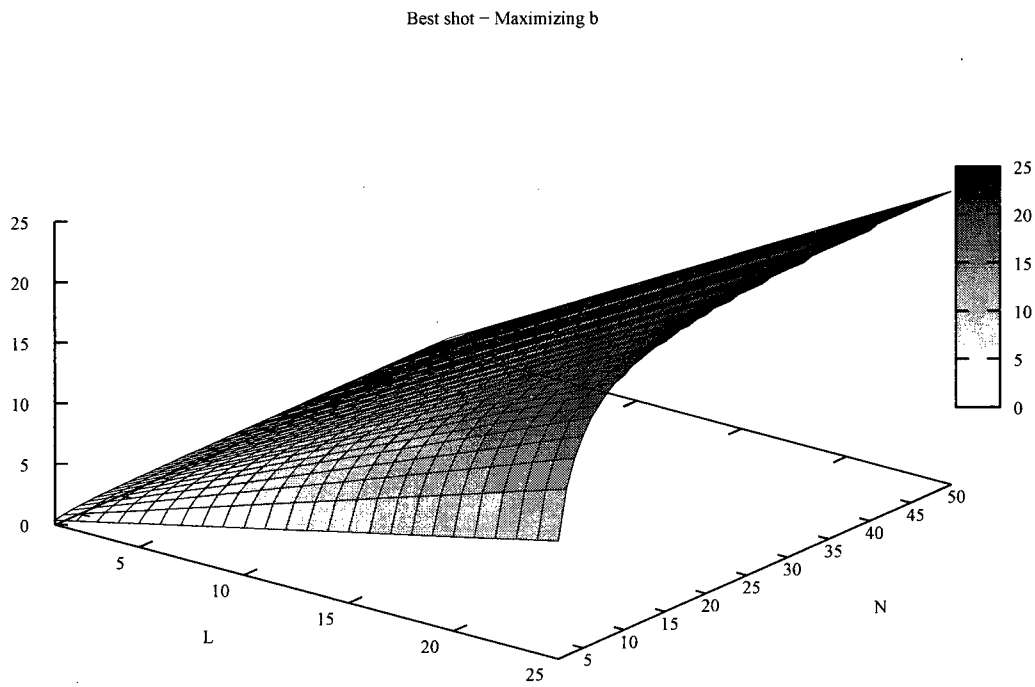

Figure 5.3: **Best shot – Payoff-ratio metric: Maximizing b for $BPoU_2(L, N)$.**

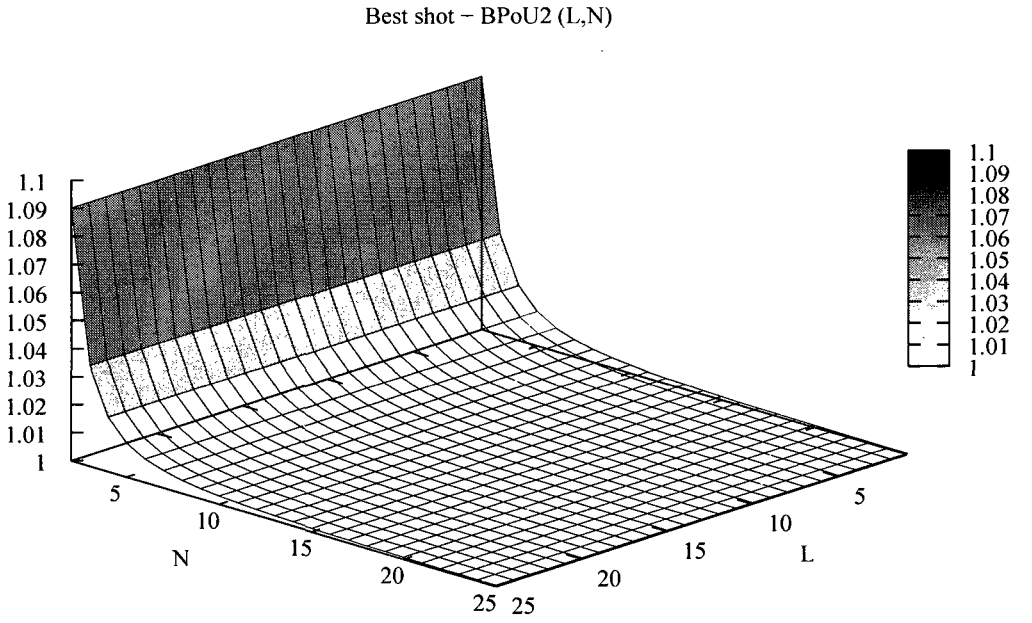

Figure 5.4: **Best shot – Payoff-ratio metric:** $BPoU_2(L, N)$**.** The metric is independent of L.

this metric 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 becomes insignificant as the number of players increases, independent of the magnitude of potential losses.

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

$$\min_{b,c\in[0,L]} \left[\frac{\text{Best Shot Expected Payoff Complete}(b, c, L, 0, N)}{\text{Best Shot Expected Payoff Incomplete}(b, c, L, 0, N)} \right] \tag{5.2}$$

This metric is expressed in terms of our payoff functions, but by starting with an initial endowment of zero, it really is 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. We have

$$
\begin{aligned}
& BPoU_3(L,N) \\
&= \min_{b,c\in[0,L]} \left[\frac{\text{Best Shot Expected Payoff Complete}(b,c,L,0,N)}{\text{Best Shot Expected Payoff Incomplete}(b,c,L,0,N)} \right] \\
&= \min_{b\in[0,L]} \frac{0 - b\left(1-\frac{b}{2L}\right)\left(\frac{b}{L}\right)^{N-1}}{0 - \frac{L}{2}\left(\frac{b}{L}\right)^{N-1}} \\
&= \min_{b\in[0,L]} \frac{2b}{L}\left(1-\frac{b}{2L}\right)
\end{aligned}
$$

Clearly the minimum value (of zero) for this expression (assuming $0 \leq b \leq L$) is achieved by taking $b = 0$. Or if the value $b = 0$ is to be avoided, the minimum is achieved by taking b arbitrarily close to zero. We observe that for the best shot game, this cost-ratio metric always measures the price of uncertainty at its greatest possible value, independent of N or L. The graphs for the maximizing b and the cost-ratio metric are both trivial but are included for consistency in Figures 5.5 and 5.6 respectively.

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.

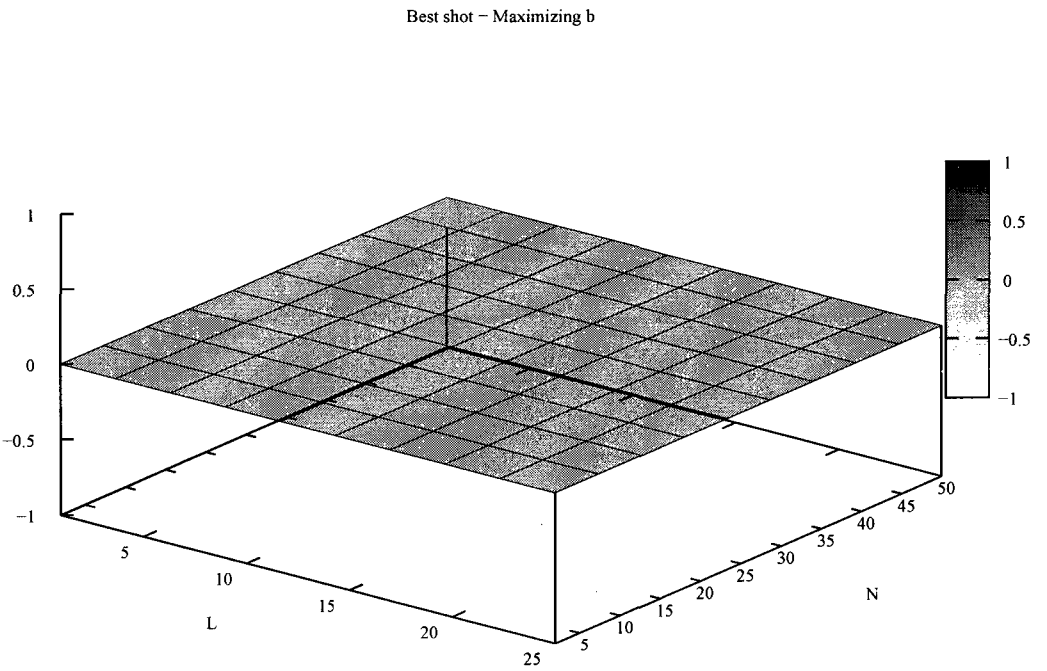

Figure 5.5: **Best shot – Cost-ratio metric: Maximizing b for $BPoU_3(L, N)$.** Here b is constantly equal to zero.

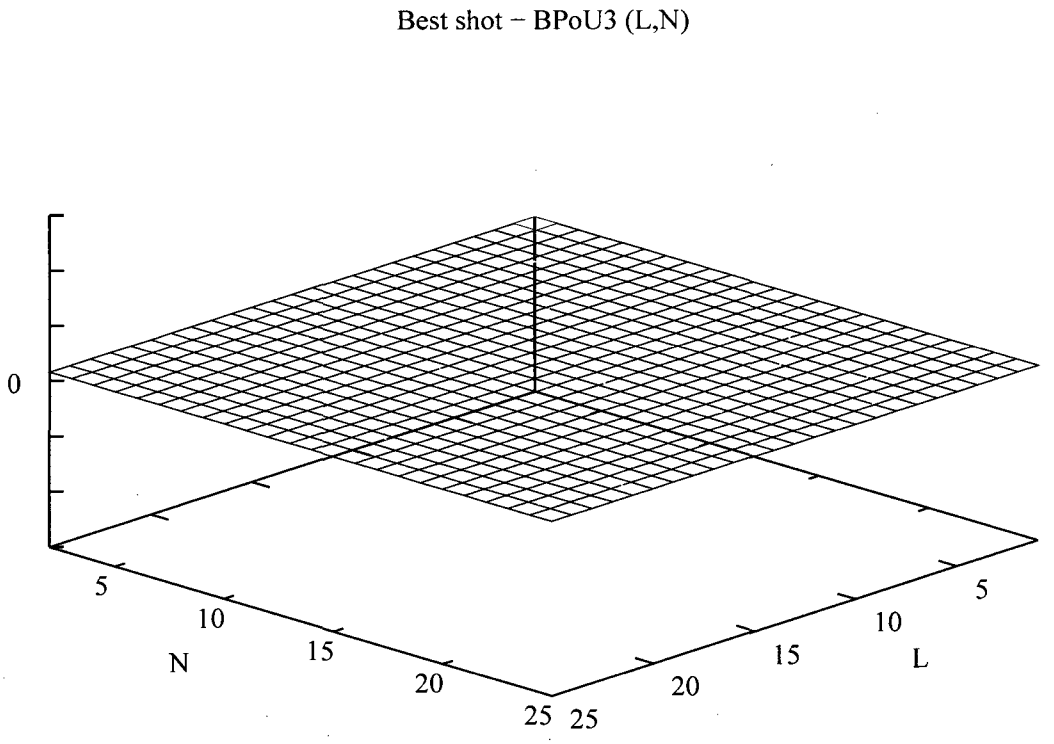

Figure 5.6: **Best shot – Cost-ratio metric:** $BPoU_3(L, N)$**.** As can be seen here, this metric is constant and equal to zero throughout the parameter space.

5.3.2 Weakest-link game

In the weakest-link game, the complete and incomplete payoffs are the same when $c < b$, but for $b \leq c$ there are a wide variety of cases to consider, and without some direction it is not clear which equations we should use. 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.

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:

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

$$- \text{Weakest-Link Expected Payoff Incomplete}(b, c, L, L, N)]$$

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 $\frac{b}{\left(1-\frac{b}{L}\right)^{N-1}} < c$ and $c < b + L\left(1 - \left(1 - \frac{b}{L}\right)^{N-1}\right)$. Assuming that the minimizing values of b and c do lie in this case, we can analyze the payoff equations for this case to get more specific information.

Weakest-Link Expected Payoff Complete(b, c, L, L, N)

$-$ Weakest-Link Expected Payoff Incomplete(b, c, L, L, N)

$$= \left[L - c + \frac{c^2}{2L} + (c-b)\left(1 - \frac{c+b}{2L}\right)\left(1 - \frac{b}{L}\right)^{N-1}\right]$$

$$- \left[L - c + \frac{b^2}{2L\left(1 - \frac{b}{L}\right)^{N-1}} + \frac{(c-b)^2}{2L\left(1 - \left(1 - \frac{b}{L}\right)^{N-1}\right)}\right]$$

$$= \frac{c^2}{2L} + (c-b)\left(1 - \frac{c+b}{2L}\right)\left(1 - \frac{b}{L}\right)^{N-1} - \frac{b^2}{2L\left(1 - \frac{b}{L}\right)^{N-1}}$$

$$- \frac{(c-b)^2}{2L\left(1 - \left(1 - \frac{b}{L}\right)^{N-1}\right)}$$

$$= \frac{c^2}{2L} + (c-b)\left(1 - \frac{b}{L}\right)^{N-1} - \frac{c^2 - b^2}{2L}\left(1 - \frac{b}{L}\right)^{N-1} - \frac{b^2}{2L\left(1 - \frac{b}{L}\right)^{N-1}}$$

$$- \frac{(c-b)^2}{2L\left(1 - \left(1 - \frac{b}{L}\right)^{N-1}\right)}$$

$$= \frac{c^2}{2L}\left(1 - \left(1 - \frac{b}{L}\right)^{N-1}\right) + (c-b)\left(1 - \frac{b}{L}\right)^{N-1} - \frac{b^2\left(1 - \left(1 - \frac{b}{L}\right)^{N-1}\right)}{2L\left(1 - \frac{b}{L}\right)^{N-1}}$$

$$- \frac{(c-b)^2}{2L\left(1 - \left(1 - \frac{b}{L}\right)^{N-1}\right)}$$

To find conditions on a minimum c for this expression we take the partial derivative with respect to c and set it equal to zero. We get:

$$0 = \frac{c}{L} + \left(1 - \frac{b}{L}\right)^{N-1} - \frac{c}{L}\left(1 - \frac{b}{L}\right)^{N-1} - \frac{2(c-b)}{2L\left(1 - \left(1 - \frac{b}{L}\right)^{N-1}\right)}$$

$$0 = \frac{c}{L}\left(1 - \left(1 - \frac{b}{L}\right)^{N-1} - \frac{1}{\left(1 - \left(1 - \frac{b}{L}\right)^{N-1}\right)}\right) + \left(1 - \frac{b}{L}\right)^{N-1} + \frac{b}{L \cdot \left(1 - \left(1 - \frac{b}{L}\right)^{N-1}\right)}$$

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

$$c = \frac{\frac{L\left(1 - \frac{b}{L}\right)^{N-1} \cdot \left(1 - \left(1 - \frac{b}{L}\right)^{N-1}\right) + b}{\left(1 - \left(1 - \frac{b}{L}\right)^{N-1}\right)}}{\frac{1 - \left(1 - \left(1 - \frac{b}{L}\right)^{N-1}\right)^2}{\left(1 - \left(1 - \frac{b}{L}\right)^{N-1}\right)}}$$

$$c = \frac{L\left(1 - \frac{b}{L}\right)^{N-1} \cdot \left(1 - \left(1 - \frac{b}{L}\right)^{N-1}\right) + b}{1 - \left(1 - \left(1 - \frac{b}{L}\right)^{N-1}\right)^2}$$

$$c = \frac{L\left[\left(1 - \frac{b}{L}\right)^{N-1} \cdot \left(1 - \left(1 - \frac{b}{L}\right)^{N-1}\right) + \frac{b}{L}\right]}{\left(1 - \frac{b}{L}\right)^{N-1} \cdot \left(1 - \left(1 - \frac{b}{L}\right)^{N-1}\right) + \left(1 - \frac{b}{L}\right)^{N-1}}$$

So this formula gives us the maximizing c as a function of b, L, and N. The dependence on L is quite weak in the sense that that $\frac{c}{L}$ is a function of N and $\frac{b}{L}$. By making the assumption $L = 1$ and solving for c, we immediately get cL as the maximizing solution for the same equation if L were not equal to 1.

Now to algebraically compute the maximizing b, we would just need to substitute the value of c from above into the payoff difference formula: $\frac{c^2}{2L}\left(1 - \left(1 - \frac{b}{L}\right)^{N-1}\right) + (c-b)\left(1 - \frac{b}{L}\right)^{N-1} - \frac{b^2\left(1 - \left(1 - \frac{b}{L}\right)^{N-1}\right)}{2L\left(1 - \frac{b}{L}\right)^{N-1}} - \frac{(c-b)^2}{2L\left(1 - \left(1 - \frac{b}{L}\right)^{N-1}\right)}$; then take the derivative with re-

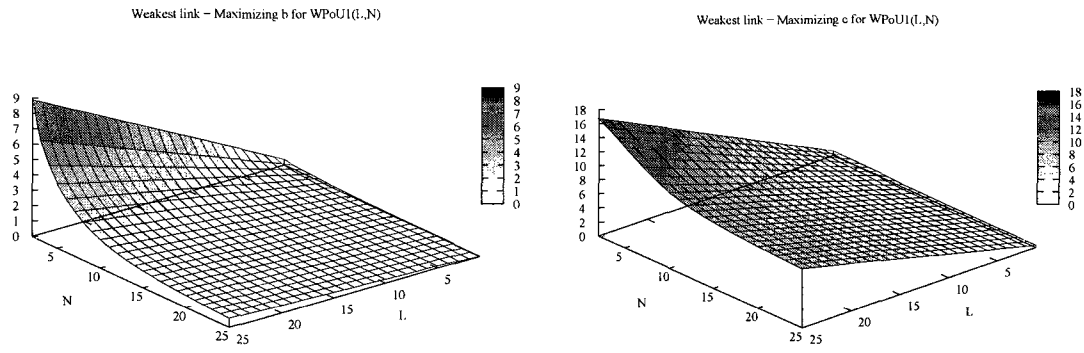

Figure 5.7: **Weakest-Link – Difference metric: Maximizing b and c for $WPoU_1(L, N)$.**

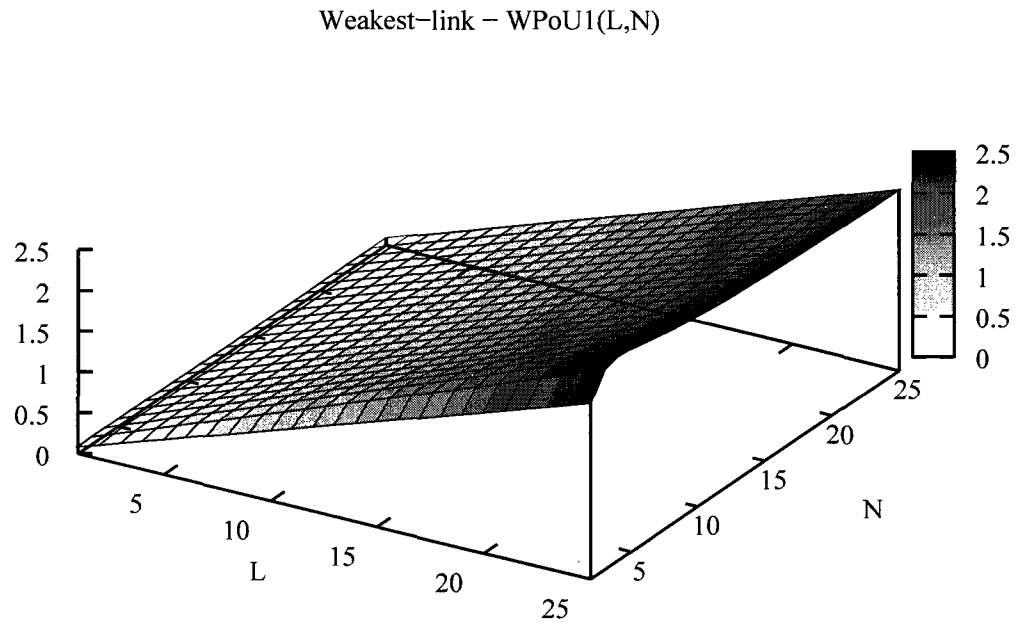

Figure 5.8: **Weakest-Link – Difference metric: $WPoU_1(L, N)$.** The metric grows linearly in the losses L and remains relatively constant for fixed L regardless of the network size N.

spect to b and find a root of this derivative in the interval $[0, L]$. We will spare the reader the computation of this derivative, as there is no closed form expression for the root of the degree 5N polynomial we would eventually need to find. Instead we refer to the graphs relevant to this metric. Figure 5.7 gives the maximizing b and c (respectively) as functions of L and N. Then Figure 5.8 gives the weakest-link difference metric $WPoU_1$ as a function of L and 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 constant 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 insurance is competitively-priced. This price of uncertainty increases directly with the potential loss, and it is not affected by the number of other players.

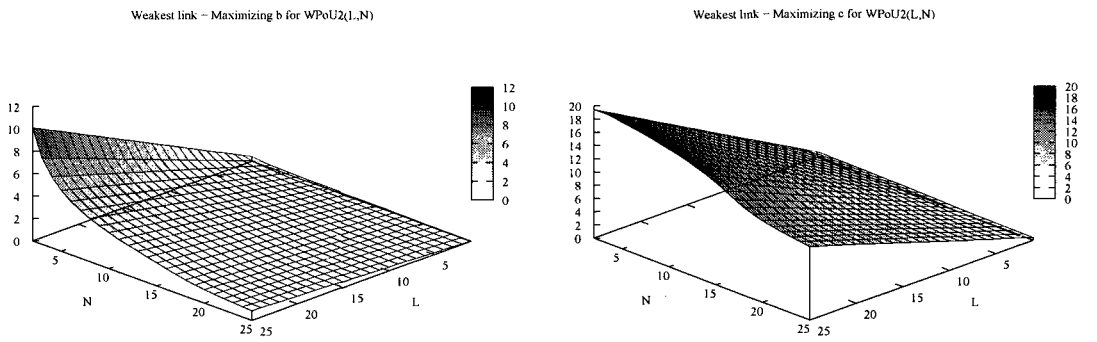

Figure 5.9: **Weakest-Link – Payoff-ratio metric: Maximizing b and c for $WPoU_2(L, N)$.** Note that the functions are actually expected to be continuous; the different "steps" that can be seen are due to sampling errors in our numerical evaluations.

The weakest-link payoff-ratio metric $PoU_2(W, L, N)$

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

$$\max_{b,c\in[0,L]} \left[\frac{\text{Weakest-Link Expected Payoff Complete}(b, c, L, L, N)}{\text{Weakest-Link Expected Payoff Incomplete}(b, c, L, L, N)} \right] \tag{5.3}$$

We begin by considering the graphs in Figure 5.9,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.

The graph for the payoff ratio metric is given in Figure 5.10. We see from the figure that it 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

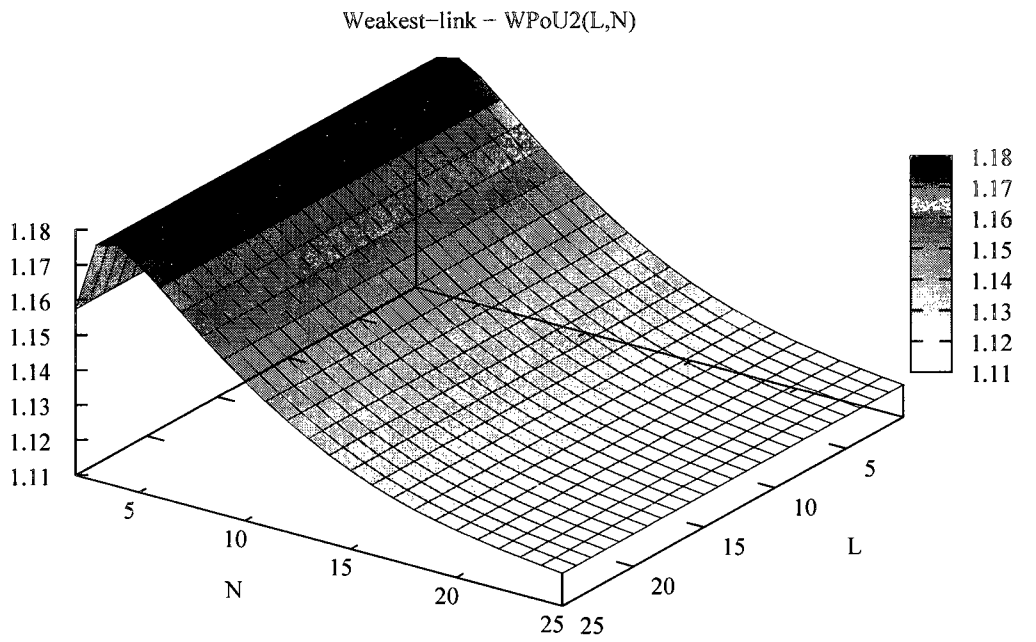

Figure 5.10: **Weakest-Link – Payoff-ratio metric:** $WPoU_2(L, N)$**.** Numeric simulations confirm the metric is independent of L.

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.

The weakest-link cost-ratio metric $PoU_3(W, L, N)$

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

$$\min_{b,c\in[0,L]} \left[\frac{\text{Weakest-Link Expected Payoff Complete}(b, c, L, 0, N)}{\text{Weakest-Link Expected Payoff Incomplete}(b, c, L, 0, N)} \right] \tag{5.4}$$

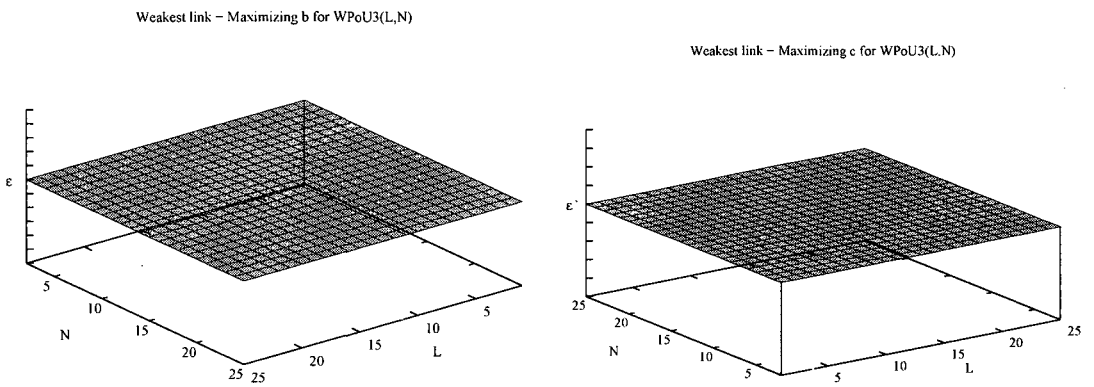

Figure 5.11: **Weakest-Link – Cost-ratio metric: Maximizing b and c for $WPoU_3(L, N)$.** ε is an extremely small positive quantity (limited by machine precision, in this case), and $\varepsilon' > \varepsilon$ is another extremely small positive quantity, barely greater than ε.

Consider the graphs in Figure 5.11, 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 maximum value for b is 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.

The graph for the payoff ratio metric is given in Figure 5.12. 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

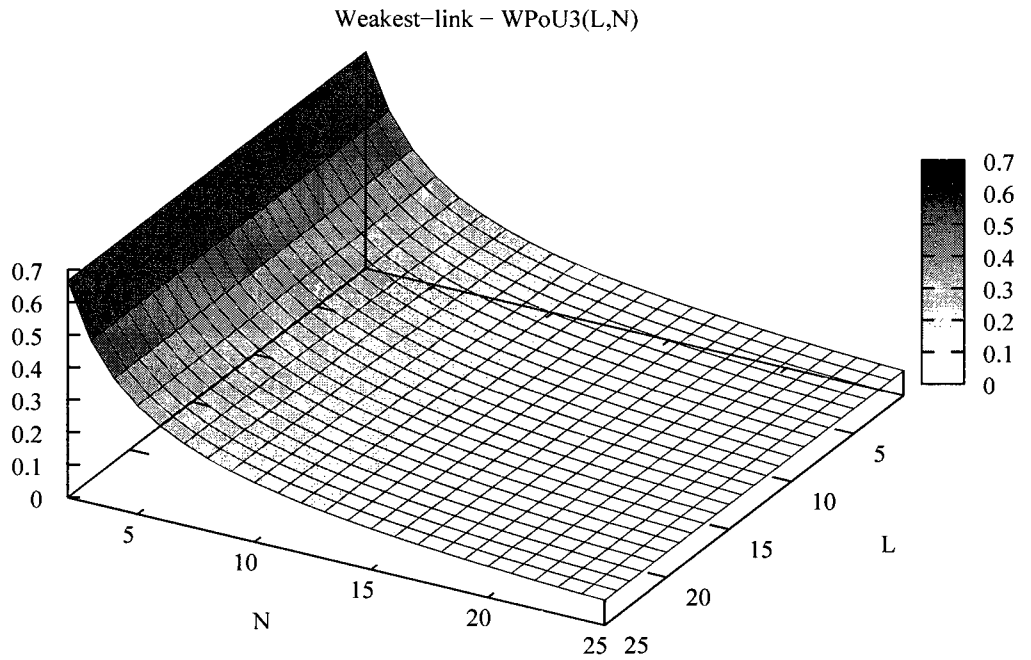

Figure 5.12: **Weakest-Link – Cost-ratio metric:** $WPoU_3(L, N)$**.** Numeric simulations confirm the metric is independent of L.

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.

5.3.3 Total effort game

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:

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

$$- \text{Total Effort Expected Payoff Incomplete}(b, c, L, M, N)]$$

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 have:

$$
\begin{aligned}
&\text{Expected Payoff Complete}(T, b, c, L, M, N) \\
&- \text{Expected Payoff Incomplete}(T, b, c, L, M, N) \\
&= \sum_{k=0}^{\lfloor N-\frac{c}{b} \rfloor} Pr[k] \cdot \left(M - c + \frac{c^2}{2L\left(1-\frac{k}{N}\right)} \right) \\
&+ \sum_{k=\lfloor N-\frac{c}{b}+1 \rfloor}^{\lfloor N-1-\frac{N}{L}(c-b) \rfloor} Pr[k] \cdot \left(M - c + \frac{b^2 N}{2L} + \frac{(c-b)^2}{2L\left(1-\frac{k+1}{N}\right)} \right) \\
&+ \sum_{k=\lfloor N-\frac{N}{L}(c-b) \rfloor}^{N-1} Pr[k] \cdot \left(M - b - \frac{L}{2}\left(1 - \frac{k+1}{N}\right) + \frac{b^2 N}{2L} \right) \\
&- \left[M - c + \frac{b^2 N}{2L} + \frac{(c-b)^2}{2\left(b-\frac{b}{N}\right)} \right] \\
&= \sum_{k=0}^{\lfloor N-\frac{c}{b} \rfloor} Pr[k] \cdot \left(\frac{c^2}{2L\left(1-\frac{k}{N}\right)} - \frac{b^2 N}{2L} \right) \\
&+ \sum_{k=\lfloor N-\frac{c}{b}+1 \rfloor}^{\lfloor N-1-\frac{N}{L}(c-b) \rfloor} Pr[k] \cdot \left(\frac{(c-b)^2}{2L\left(1-\frac{k+1}{N}\right)} \right) \\
&+ \sum_{k=\lfloor N-\frac{N}{L}(c-b) \rfloor}^{N-1} Pr[k] \cdot \left(c - b - \frac{L}{2}\left(1 - \frac{k+1}{N}\right) \right) \\
&- \frac{(c-b)^2}{2\left(b-\frac{b}{N}\right)}
\end{aligned}
$$

Now because c occurs in the terms of this expression only quadratically, we could compute an expression for the partial derivative with respect to c that is almost-everywhere valid, then set the derivative equal to zero and solve for c. In fact, we did compute this,

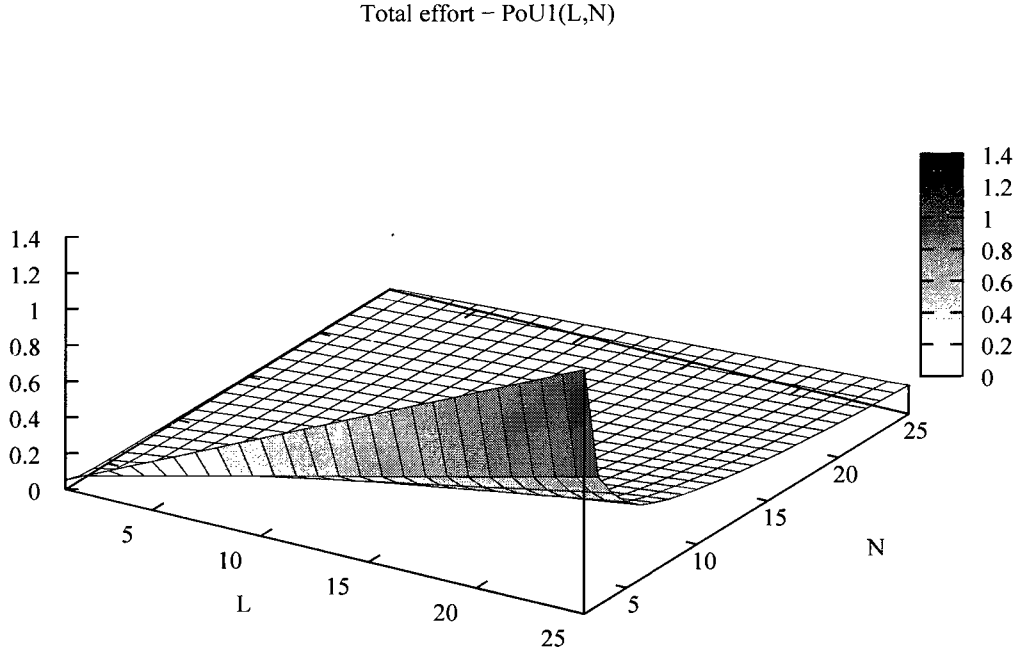

Figure 5.13: **Total effort – Difference metric:** $TPoU_1(L, N)$.

obtaining

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

The problem with this formulation in terms of an algebraic analysis is that the variable c also occurs in the terms of the summands, and it is not clear how to use algebra to get it out of there.

Proceeding with our numerical analysis, Figure 5.13 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 $\frac{1}{N}$.

Observations. The interpretation of our numerical results for this metric is that the price

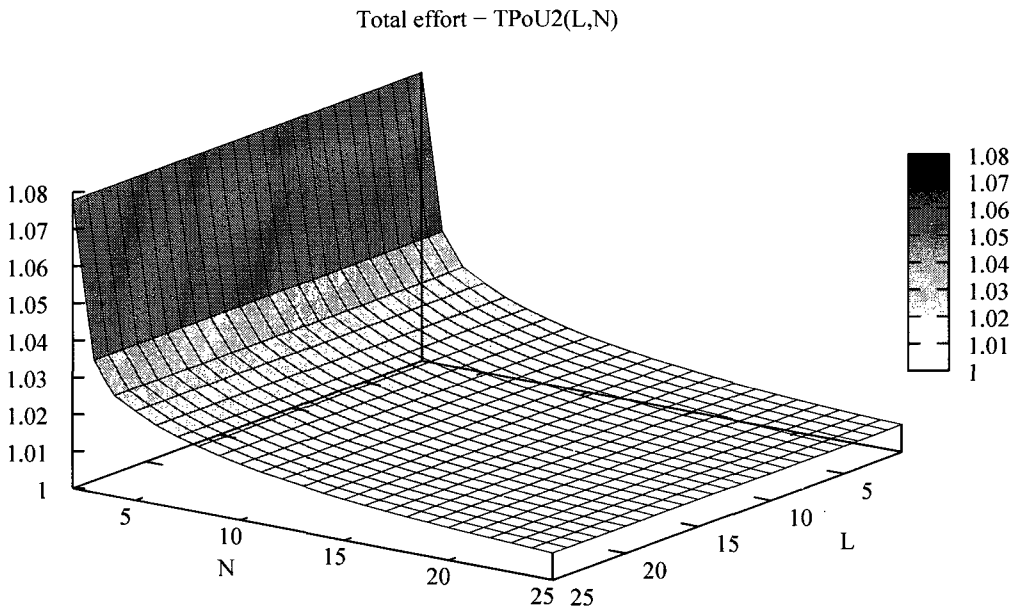

Figure 5.14: **Total effort – Payoff-ratio metric:** $TPoU_2(L, N)$**.**

of uncertainty increases with the potential losses, but as the number of players increases, the price of uncertainty diminishes quickly.

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:

$$\max_{b,c\in[0,L]} \left[\frac{\text{Total Effort Expected Payoff Complete}(b, c, L, L, N)}{\text{Total Effort Expected Payoff Incomplete}(b, c, L, L, N)} \right] \tag{5.5}$$

For the remaining total effort metrics, our analysis relies exclusively on numerical approximations. Figure 5.14 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.

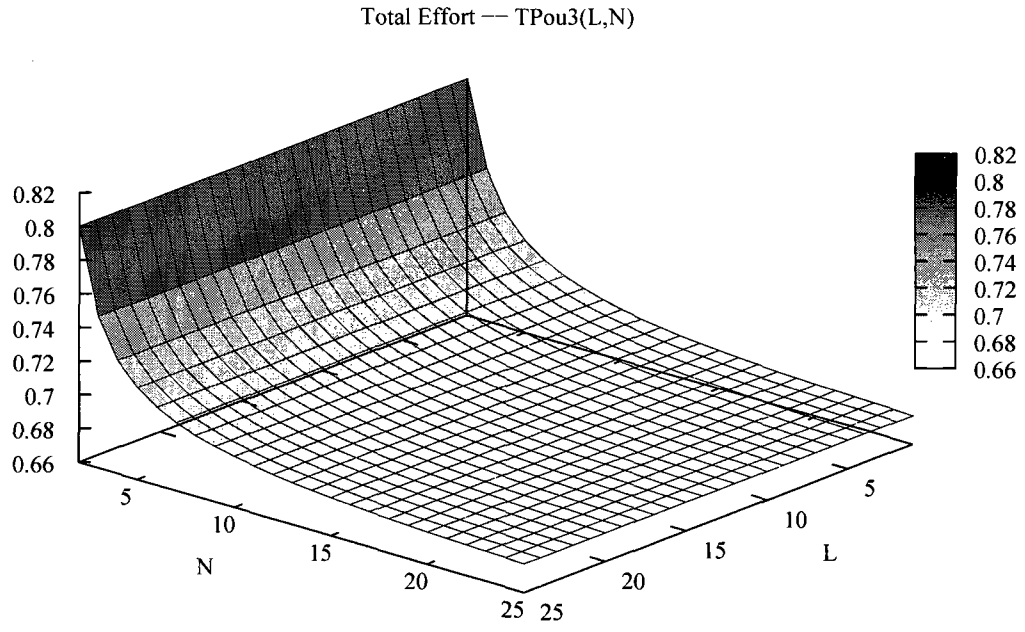

Figure 5.15: **Total effort – Cost-ratio metric:** $TPoU_3(L, N)$**.**

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.

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:

$$\max_{b,c\in[0,L]} \left[\frac{\text{Total Effort Expected Payoff Complete}(b, c, L, 0, N)}{\text{Total Effort Expected Payoff Incomplete}(b, c, L, 0, N)} \right] \tag{5.6}$$

Figure 5.15 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.

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 for the other two metrics. Again we surmise that this happens because the cost-ratio metric focuses on the cases where the cost for both complete and incomplete information scenarios are quite small, but the ratio shows a significant distinction.

5.4 Summary

In this chapter we continued 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 (see Chapter 4). 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 [116] or market-based [27]. However, the price of uncertainty is motivated by game-theory and, more specifically, by Koutsoupias and Papadimitriou's metric to evaluate worst-case equilibria [129], 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 A.5, A.10, and A.15) 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 (see Chapter 4).

- 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 [175]. 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 [134]. 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 [103, 203].

- 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 [69]. 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 this work will contribute to a serious discussion of information metrics applied to interdependent security scenarios. As we have shown here, selecting the most appropriate metric is not an straightforward choice, and several pitfalls exist.

Chapter 6

Conclusions

6.1 Contributions

In this dissertation, we have developed a framework for the analysis of decision-making in the presence of security interdependencies and multiple types of investment types. We consider three canonical tightly-coupled games that capture scenarios in which users will jointly suffer from security breaches. Further, we present two variations of the novel weakest-target game. The game captures loosely-coupled interdependencies in which agents receive differentiated results based on their security strategies.

The weakest-target game describes several security-relevant decision situations in which an attacker wants to exploit the least protected machines (e.g., to reduce the cost of amassing a large botnet for distributed denial of service attacks or to send unsolicited communications). We believe that the usefulness of the weakest-target game also extends to other

contexts. For example, shirking in workplace environments as well as showing anti-social behaviors in society can be modeled with this game.

We further distinguish between preventive and mitigative security investments. The effectiveness of users' protection investments are subject to interdependencies, while self-insurance investments (e.g., backup technologies) yield a private return. We find that in the tightly-coupled interdependencies the co-existence of different types of equilibrium strategies creates challenges for successful system security. For example, in the weakest-link game (for a wide range of parameter settings) protection and self-insurance strategies compete. As a result, users have no immediate means to infer whether other individuals will invest in protective measures rather than relying on recovery mechanisms. This problem can be a significant factor for failed security readiness, and the observed discrepancies of investment strategies in practice.

We further extend our analysis to include bounded rational user types. We assume the population to include expert as well as inexperienced users. Experts act rationally and understand the implications of system interdependencies. Non-expert users draw on rule-of-thumb strategies and neglect the role of interdependencies, i.e., the impact of their investments on others and vice versa.

Another extension is the consideration of alternative information structures. We contrast and compare security games under complete information with an incomplete information setting. Users are aware of their own security-relevant cost and threat parameters, however they do not know the extent of the potential damages that other individuals face.

We develop a decision-theoretic methodology to study games with bounded rational agents and incomplete information. We further develop the price of uncertainty metrics to efficiently compare and highlight payoff differences between the complete and incomplete information environment.

Our metrics analysis has implications for network designers that want to avoid undesirable hotspots that penalize users for their lack of information about threats. Similarly, service providers or other intermediaries may take influence on the pricing and availability of security technologies to steer users to less harmful parameter configurations.

6.2 Open questions

There are a number of topics that deserve further consideration.

Attacker ecology: We intend to study in more detail the impact of including strategic attackers in our modeling framework [75]. Many malefactors implement fully automated attack protocols for which the exogenous attacker assumption is suitable. However, other interactions are highly specialized and individualized and, therefore, attackers may react strategically to changes in important security parameters and the defenders' incentives for investments [42, 189]. Explicitly modeling the incentives of attackers also allows the meaningful inclusion of additional defense strategies, such as investments in attacker identification and enforcement. We can further capture the inherent asymmetry between attackers and defenders, i.e., targets need to successfully defend against multiple threats

while attackers can benefit from a single weakness [48]. Considering multiple attackers, we can investigate the interdependencies that exist between attackers. Offenders compete for scarce resources, for example, the least protected defenders. They may also cooperate to achieve a coordinated attack on a well-defended target (e.g., the infrastructure of a nation state or the Internet).

Behavioral study: Most of our analysis assumes that all players are selfish and perfectly rational (Chapters 2 and 3) or follow predictable patterns of nearsightedness (Chapters 4 and 5). As has been discussed elsewhere, e.g., [47], this assumption generally leads to idealized models, which deserve to be complemented by empirical studies [104, 133]. We are developing a set of laboratory experiments to conduct user studies and attempt to measure the differences between perfectly rational behavior and actual strategies devised and played.

Bibliography

[1] C. Abad. The economics of phishing: A survey of the operations of the phishing market. *First Monday*, 10(9), September 2005.

[2] A. Acquisti. Privacy in electronic commerce and the economics of immediate gratification. In *Proceedings of the Fifth ACM Conference on Electronic Commerce (EC'04)*, pages 21–29, New York, NY, May 2004.

[3] A. Acquisti and J. Grossklags. Privacy and rationality in individual decision making. *IEEE Security & Privacy*, 3(1):26–33, January–February 2005.

[4] A. Acquisti and H. Varian. Conditioning prices on purchase history. *Marketing Science*, 24(3):367–381, Summer 2005.

[5] G. Akerlof and J. Yellen. Can small deviations from rationality make significant differences to economic equilibria? *American Economic Review*, 75(4):708–720, September 1985.

[6] S. Allison, A. Schuck, and K. Lersch. Exploring the crime of identity theft: Preva-

lence, clearance rates, and victim/offender characteristics. *Journal of Criminal Justice*, 33(1):19–29, January–February 2005.

[7] R. Anderson. Liability and computer security: Nine principles. In *Proceedings of the Third European Symposium on Research in Computer Security (ESORICS 1994)*, pages 231–245, Brighton, UK, November 1994.

[8] R. Anderson. The eternity service. In *Proceedings of the 1st International Conference on the Theory and Applications of Cryptology (PRAGOCRYPT)*, pages 242–252, Prague, Czech Republic, September 1996.

[9] R. Anderson. *Security Engineering: A Guide to Building Dependable Distributed Systems*. Wiley Computer Publishing, New York, NY, second edition, 2001.

[10] R. Anderson. Why information security is hard – An economic perspective. In *Proceedings of the 17th Annual Computer Security Applications Conference (ACSAC'01)*, New Orleans, LA, December 2001.

[11] R. Anderson and T. Moore. The economics of information security. *Science*, 314(5799):610–613, October 1998.

[12] AOL/NCSA. Online safety study, October 2004. Available via archive.org: `http://web.archive.org/web/20041025192551/http://www.staysafeonline.info/news/safety_study_v04.pdf`.

[13] AOL/NCSA. Online safety study, December 2005. Available via

archive.org: http://web.archive.org/web/20051220151024/http://www.staysafeonline.org/pdf/safety_study_2005.pdf.

[14] F. Asgharpour, D. Liu, and J. Camp. Mental models of computer security risks. In *Proceedings of the Seventh Workshop on the Economics of Information Security (WEIS)*, Hanover, NH, June 2008.

[15] J. Aspnes, K. Chang, and A. Yampolskiy. Inoculation strategies for victims of viruses and the sum-of-squares partition problem. *Journal of Computer and System Sciences*, 72(6):1077–1093, September 2006.

[16] T. August and T. Tunca. Network software security and user incentives. *Management Science*, 52(11):1703–1720, November 2006.

[17] M. Balcan, A. Blum, and Y. Mansour. The price of uncertainty. In *Proceedings of the Tenth ACM Conference on Electronic Commerce (EC'09)*, pages 285–294, Stanford, CA, July 2009.

[18] V. Basili and B. Perricone. Software errors and complexity: An empirical investigation. *Communications of the ACM*, 27(1):42–52, January 1984.

[19] J. Bauer, M. van Eeten, T. Chattopadhyay, and Y. Wu. ITU study on the financial aspects of network security: Malware and spam, July 2008. A study conducted by the International Telecommunication Union, Geneva, Switzerland.

[20] L. Bellamy, D. Hutchinson, and J. Wells. User perceptions and acceptance of benevolent worms – A matter of fear? In *6th IEEE/ACIS International Conference on Computer and Information Science*, pages 29–36, Melbourne, Australia, July 2007.

[21] D. Besnard and B. Arief. Computer security impaired by legitimate users. *Computers & Security*, 23(3):253–264, May 2004.

[22] V. Bier, S. Oliveros, and L. Samuelson. Choosing what to protect: Strategic defensive allocation against an unknown attacker. *Journal of Public Economic Theory*, 9(4):563–587, August 2007.

[23] K. Birman and F. Schneider. The monoculture risk put into context. *IEEE Security & Privacy*, 7(1):14–17, January–February 2009.

[24] P. Bishop and R. Bloomfield. A conservative theory for long-term reliability-growth prediction. *IEEE Transactions on Reliability*, 45(4):550–560, December 1996.

[25] M. Blaze. Protocol failure in the escrowed encryption standard. In *Proceedings of the 2nd ACM Conference on Computer and Communications Security (CCS)*, pages 59–67, Fairfax, VA, November 1994.

[26] R. Böhme and G. Kataria. Models and measures for correlation in cyber-insurance. In *Proceedings of the Fifth Workshop on the Economics of Information Security (WEIS)*, Cambridge, UK, June 2006.

[27] R. Böhme and T. Nowey. Economic security metrics. In I. Eusgeld, F. Freiling, and

R. Reussner, editors, *Dependability Metrics (Lecture Notes in Computer Science, No. 4909)*, pages 176–187. Springer-Verlag, 2008.

[28] K. Boitmanis, U. Brandes, and C. Pich. Visualizing Internet evolution on the autonomous systems level. In *Graph Drawing, Lecture Notes in Computer Science, Volume 4875*, pages 365–376. Springer Verlag, Heidelberg, Germany, July 2008.

[29] J. Bolot and M. Lelarge. A new perspective on Internet security using insurance. In *Proceedings of the 27th Conference on Computer Communications (INFOCOM'08)*, pages 1948–1956, Phoenix, AZ, April 2008.

[30] P. Boothe, J. Hiebert, and R. Bush. Short-lived prefix hijacking on the Internet. *Presentation at the Winter 2006 NANOG Meeting (NANOG36)*, February 2006. Available from North American Network Operators Group (NANOG): `http://www.nanog.org/meetings/nanog36/presentations/boothe.pdf`.

[31] N. Borisov, I. Goldberg, and D. Wagner. Intercepting mobile communications: The insecurity of 802.11. In *Proceedings of the 7th ACM SIGMOBILE Annual International Conference on Mobile Computing and Networking*, pages 180–189, Rome, Italy, July 2001.

[32] M. Boyer and G. Dionne. Variations in the probability and magnitude of loss: Their impact on risk. *Canadian Journal of Economics*, 16(3):411–419, August 1983.

[33] R. Brady, R. Anderson, and R. Ball. Murphy's law, the fitness of evolving species,

and the limits of software reliability (Report No. 471). Technical report, University of Cambridge, Computer Laboratory, September 1999.

[34] J. Brandts and W. MacLeod. Equilibrium selection in experimental games with recommended play. *Games and Economic Behavior*, 11(1):36–63, October 1995.

[35] E. Briys, H. Schlesinger, and J. Graf v. d. Schulenburg. Reliability of risk management: Market insurance, self-insurance and self-protection reconsidered. *The Geneva Papers on Risk and Insurance Theory*, 16(1):45–58, June 1991.

[36] Bruskin Research. Nearly one in four computer users have lost content to blackouts, viruses and hackers according to new national survey, 2001. Survey conducted for Iomega Corporation. Condensed results available at: `http://www.iomega.com/about/prreleases/2001/viruses_hackers_poweroutages.html`.

[37] K. Bryant and J. Campbell. User behaviours associated with password security and management. *Australasian Journal of Information Systems*, 14(1):36–63, November 2006.

[38] J.A. Bull, L. Gong, and K. Sollins. Towards security in an open systems federation. In *Proceedings of the Second European Symposium on Research in Computer Security (ESORICS 1992), Springer LNCS No. 648*, pages 3–20, Toulouse, France, November 1992.

[39] K. Butler, T. Farley, P. McDaniel, and J. Rexford. A survey of BGP security issues and solutions. Technical report, AT&T Research Labs, August 2008. Technical Report TD-5UGJ33. Revised version available from authors' websites.

[40] C. Camerer. *Behavioral Game Theory: Experiments in Strategic Interaction.* Princeton University Press, Princeton, NJ, 2003.

[41] K. Campbell, L. Gordon, M. Loeb, and L. Zhou. The economic cost of publicly announced information security breaches: Empirical evidence from the stock market. *Journal of Computer Security*, 11(3):431–448, March 2003.

[42] H. Cavusoglu, S. Raghunathan, and W. Yue. Decision-theoretic and game-theoretic approaches to IT security investment. *Journal of Management Information Systems*, 25(2):281–304, Fall 2008.

[43] A. Chakrabarti and G. Manimaran. Internet infrastructure security: A taxonomy. *IEEE Network*, 16(6):13–21, November/December 2002.

[44] Y. Chang and I. Ehrlich. Insurance, protection from risk, and risk-bearing. *Canadian Journal of Economics*, 18(3):574–586, August 1985.

[45] P. Chen, G. Kataria, and R. Krishnan. On software diversification, correlated failures and risk management, April 2006. Available at SSRN: `http://ssrn.com/abstract=906481`.

[46] J. Choi, C. Fershtman, and N. Gandal. Network security: Vulnerabilities and disclosure policy, December 2008. Working paper.

[47] N. Christin, J. Grossklags, and J. Chuang. Near rationality and competitive equilibria in networked systems. In *Proceedings of ACM SIGCOMM'04 Workshop on Practice and Theory of Incentives in Networked Systems (PINS)*, pages 213–219, Portland, OR, August 2004.

[48] D. Clark and K. Konrad. Asymmetric conflict: Weakest link against best shot. *Journal of Conflict Resolution*, 51(3):457–469, June 2007.

[49] D. Clark, J. Wroclawski, K. Sollins, and R. Braden. Tussle in cyberspace: Defining tomorrow's Internet. In *Proceedings of ACM SIGCOMM'02*, pages 347–356, Pittsburgh, PA, August 2002.

[50] R. Clayton. Email traffic: A quantitative snapshot. In *Proceedings of the Fourth Conference on Email and Anti-Spam (CEAS)*, Mountain View, CA, August 2007.

[51] Computer Security Institute. 1997 CSI/FBI computer crime and security survey. March 1997.

[52] Consumers Union. Boom time for cybercrime: The economy and online social networks are the latest fodder for scams. *Consumer Reports Magazine*, June 2009.

[53] R. Cornes and T. Sandler. *The theory of externalities, public goods, and club goods*. Cambridge University Press, Cambridge, UK, 1986.

[54] M. Costa, J. Crowcroft, M. Castro, A. Rowstron, L. Zhou, L. Zhang, and P. Barham. Vigilante: End-to-end containment of Internet worms. In *Proceedings of the Twentieth ACM Symposium on Operating Systems Principles (SOSP)*, pages 133–147, Brighton, UK, October 2005.

[55] M. Cremonini and D. Nizovtsev. Understanding and influencing attackers decisions: Implications for security investment strategies. In *Proceedings of the Fifth Workshop on the Economics of Information Security (WEIS)*, Cambridge, UK, June 2006.

[56] D. Dörner. *The Logic Of Failure: Recognizing And Avoiding Error In Complex Situations*. Metropolitan Books, New York, NY, 1996.

[57] G. Danezis and R. Anderson. The economics of resisting censorship. *IEEE Security & Privacy*, 3(1):45–50, January–February 2005.

[58] D. Denning. Activism, hacktivism, and cyberterrorism: The Internet as a tool for influencing foreign policy. In J. Arquilla and D. Ronfeldt, editors, *Networks and Netwars: The Future of Terror, Crime, and Militancy*, pages 239–288. RAND Corporation, Santa Monica, CA, 2002.

[59] X. Dimitropoulos, D. Krioukov, G. Riley, and K. Claffy. Revealing the autonomous system taxonomy: The machine learning approach. In *Proceedings of the Passive and Active Measurement Workshop (PAM)*, Adelaide, Australia, March 2006.

[60] P. Dourish, R. Grinter, J. Delgado de la Flor, and M. Joseph. Security in the wild:

User strategies for managing security as an everyday, practical problem. *Personal and Ubiquitous Computing*, 8(6):391–401, November 2004.

[61] S. Drimer, S. Murdoch, and R. Anderson. Optimised to fail: Card readers for online banking. In *Proceedings of the 13th International Conference Financial Cryptography and Data Security (FC'09), Lecture Notes in Computer Science (LNCS), No. 5628*, pages 184–200, Christ Church, Barbados, February 2009.

[62] I. Ehrlich and G.S. Becker. Market insurance, self-insurance, and self-protection. *Journal of Political Economy*, 80(4):623–648, July 1972.

[63] A. Etzioni. On thoughtless rationality (rules-of-thumb). *Kyklos*, 40(4):496–514, November 1987.

[64] R. Faris, S. Wang, and J. Palfrey. Censorship 2.0. *Innovations*, 3(2):165–187, Spring 2008.

[65] N. Feamster, L. Gao, and J. Rexford. How to lease the Internet in your spare time. *ACM SIGCOMM Computer Communications Review*, 37(1):61–64, January 2007.

[66] Federal Trade Commission, Office of Inspector General. Review of the FTC Consumer Response Center, July 2007. `http://www.ftc.gov/oig/reports/ar0703crc.pdf`.

[67] A. Feldman, A. Halderman, and E. Felten. Security analysis of the Diebold

AccuVote-TS voting machine. In *Proceedings of the USENIX Workshop on Accurate Electronic Voting Technology (EVT'09)*, Boston, MA, August 2007.

[68] N. Fenton and M. Neil. A critique of software defect prediction models. *IEEE Transactions on Software Engineering*, 25(5):675–689, September/October 1999.

[69] D. Fetherstonhaugh, P. Slovic, S. Johnson, and J. Friedrich. Insensitivity to the value of human life: A study of psychophysical numbing. *Journal of Risk and Uncertainty*, 14(3):283–300, May 1997.

[70] R. Ford and S. Gordon. Cent, five cent, ten cent, dollar: Hitting botnets where it really hurts. In *Proceedings of the 2006 Workshop on New Security Paradigms (NSPW)*, pages 3–10, Schloss Dagstuhl, Germany, September 2006.

[71] J. Fox. The learning of strategies in a simple, two-person zero-sum game without saddlepoint. *Behavioral Science*, 17(3):300–308, May 1972.

[72] R. Frank. Shrewdly irrational. *Sociological Forum*, 2(1):21–41, December 1987.

[73] J. Franklin, V. Paxson, A. Perrig, and S. Savage. An inquiry into the nature and causes of the wealth of internet miscreants. In *Proceedings of 14th ACM Conference on Computer and Communications Security (CCS)*, pages 375–388, Alexandria, VA, October 2007.

[74] E. Friedman, M. Shor, S. Shenker, and B. Sopher. An experiment on learning with

limited information: Nonconvergence, experimentation cascades, and the advantage of being slow. *Games and Economic Behavior*, 47(2):325–352, May 2004.

[75] N. Fultz and J. Grossklags. Blue versus red: Towards a model of distributed security attacks. In *Proceedings of the 13th International Conference Financial Cryptography and Data Security (FC'09), Lecture Notes in Computer Science (LNCS), No. 5628*, pages 167–183, Christ Church, Barbados, February 2009.

[76] E. Gal-Or and A. Ghose. The economic incentives for sharing security information. *Information Systems Research*, 16(2):186–208, June 2005.

[77] M. Gallaher, B. Rowe, A. Rogozhin, and A. Link. Economic analysis of cyber security and private sector investment decisions, April 2006. Report prepared for the United States Department of Homeland Security.

[78] S. Gaw. *Ideals and Reality: Adopting Secure Technologies and Developing Secure Habits to Prevent Message Disclosure*. PhD thesis, Princeton University, June 2009.

[79] D. Geer, C. Pfleeger, B. Schneier, J. Quarterman, P. Metzger, R. Bace, and P. Gutmann. Cyberinsecurity: The cost of monopoly. How the dominance of Microsoft's products poses a risk to society, 2003. Available from Computer & Communications Industry Association at `http://www.ccianet.org/papers/cyberinsecurity.pdf`.

[80] A. Ghose, J. Grossklags, and J. Chuang. Resilient data-centric storage in wireless ad-

hoc sensor networks. In *Proceedings of the 4th International Conference on Mobile Data Management (MDM)*, pages 45–62, Melbourne, Australia, January 2003.

[81] A. Ghosh. *E-Commerce Security: Weak Links, Best Defenses (Paperback)*. John Wiley and Sons, New York, NY, 1998.

[82] J. Goeree and C. Holt. A model of noisy introspection. *Games and Economic Behavior*, 46(2):365–382, February 2004.

[83] N. Good, R. Dhamija, J. Grossklags, S. Aronovitz, D. Thaw, D. Mulligan, and J. Konstan. Stopping spyware at the gate: A user study of privacy, notice and spyware. In *Proceedings of the Symposium On Usable Privacy and Security (SOUPS 2005)*, pages 43–52, Pittsburgh, PA, July 2005.

[84] N. Good, J. Grossklags, D. Mulligan, and J. Konstan. Noticing notice: A large-scale experiment on the timing of software license agreements. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI'07)*, pages 607–616, San Jose, CA, April 2007.

[85] L. Gordon. Incentives for improving cybersecurity in the private sector: A cost-benefit perspective, October 2007. Testimony for the House Committee on Homeland Security's Subcommittee on Emerging Threats, Cybersecurity, and Science and Technology. Statement available at: `http://homeland.house.gov/SiteDocuments/20071031155020-22632.pdf`.

[86] L. Gordon and M. Loeb. *Managing Cyber-Security Resources: A Cost-Benefit Analysis*. McGraw-Hill, New York, NY, 2006.

[87] L.A. Gordon and M. Loeb. The economics of information security investment. *ACM Transactions on Information and System Security*, 5(4):438–457, November 2002.

[88] L.A. Gordon, M. Loeb, and W. Lucyshyn. Sharing information on computer systems security: An economic analysis. *Journal of Accounting and Public Policy*, 22(6):461–485, November 2003.

[89] S. Gordon. The generic virus writer. In *Proceedings of the International Virus Bulletin Conference*, pages 121 – 138, Jersey, Channel Islands, 1994.

[90] S. Gordon. Virus writers - The end of the innocence? In *10th Annual Virus Bulletin Conference (VB2000)*, Orlando, FL, September 2000. Available from IBM Research at `http://www.research.ibm.com/antivirus/SciPapers/VB2000SG.htm`.

[91] G. Gottlob, G. Greco, and F. Scarcello. Pure Nash equilibria: Hard and easy games. *Journal of Artificial Intelligence Research*, 24:357–406, July–December 2005.

[92] J. Granick. Faking it: Calculating loss in computer crime sentencing. *I/S: A Journal of Law and Policy for the Information Society*, 2(2):207–228, Spring/Summer 2006.

[93] J. Grossklags. Experimental economics and experimental computer science: A sur-

vey. In *Workshop on Experimental Computer Science (ExpCS'07), ACM Federated Computer Research Conference (FCRC)*, San Diego, CA, June 2007.

[94] J. Grossklags, N. Christin, and J. Chuang. Secure or insure? A game-theoretic analysis of information security games. In *Proceedings of the 2008 World Wide Web Conference (WWW'08)*, pages 209–218, Beijing, China, April 2008.

[95] J. Grossklags, N. Christin, and J. Chuang. Security and insurance management in networks with heterogeneous agents. In *Proceedings of the Ninth ACM Conference on Electronic Commerce (EC'08)*, pages 160–169, Chicago, IL, July 2008.

[96] J. Grossklags and B. Johnson. Uncertainty in the weakest-link security game. In *Proceedings of the International Conference on Game Theory for Networks (GameNets 2009)*, pages 673–682, Istanbul, Turkey, May 2009.

[97] J. Grossklags, B. Johnson, and N. Christin. The price of uncertainty in security games. In *Proceedings of the Eighth Workshop on the Economics of Information Security (WEIS)*, London, UK, June 2009.

[98] A. Halderman, S. Schoen, N. Heninger, W. Clarkson, W. Paul, J. Calandrino, A. Feldman, J. Appelbaum, and E. Felten. Lest we remember: Cold boot attacks on encryption keys. In *Proceedings of the 17th USENIX Security Symposium*, pages 45–60, San Jose, CA, August 2008.

[99] J. Hartley. Retrospectives: The origins of the representative agent. *The Journal of Economic Perspectives*, 10(2):169–177, Spring 1996.

[100] M. Hathaway. Remarks by Melissa E. Hathaway, Acting Senior Director for Cyberspace for the National Security and Homeland Security Councils. As prepared for delivery at the RSA Conference 2009, April 2009.

[101] G. Heal and H. Kunreuther. IDS models of airline security. *Journal of Conflict Resolution*, 49(2):201–217, April 2005.

[102] C. Herley and D. Florencio. Nobody sells gold for the price of silver: Dishonesty, uncertainty and the underground economy. In *Proceedings of the Eighth Workshop on the Economics of Information Security (WEIS)*, London, UK, June 2009.

[103] J. Hershey and J. Baron. Clinical reasoning and cognitive processes. *Medical Decision Making*, 7(4):203–211, December 1987.

[104] R. Hess, C. Holt, and A. Smith. Coordination of strategic responses to security threats: Laboratory evidence. *Experimental Economics*, 10(3):235–250, September 2007.

[105] R. Hesseling. Displacement: A review of the empirical literature. In R. Clarke, editor, *Crime Prevention Studies, Volume 3*, pages 197–230. Criminal Justice Press, Monsey, NY, 1994.

[106] L.D. Hiebert. Self insurance, self protection and the theory of the competitive firm. *Southern Economic Journal*, 50(1):160–168, July 1983.

[107] J. Hirshleifer. From weakest-link to best-shot: The voluntary provision of public goods. *Public Choice*, 41(3):371–386, January 1983.

[108] J. Hirshleifer. From weakest-link to best-shot: Correction. *Public Choice*, 46(2):221–223, January 1985.

[109] K. Hole, V. Moen, A. Klingsheim, and K. Tande. Lessons from the Norwegian ATM system. *IEEE Security & Privacy*, 5(6):25–31, November–December 2007.

[110] R. Hollinger. Hackers: Computer heroes or electronic highwaymen? *ACM SIGCAS Computers and Society*, 21(1):6–17, June 1991.

[111] T. Holz, M. Steiner, F. Dahl, E. Biersack, and F. Freiling. Measurements and mitigation of peer-to-peer-based botnets: A case study on storm worm. In *Proceedings of the 1st USENIX Workshop on Large-Scale Exploits and Emergent Threats (LEET '08)*, San Francisco, CA, April 2008.

[112] P. Honeyman, G.A. Schwartz, and A. van Assche. Interdependence of reliability and security. In *Workshop on Information Systems and Economics (WISE 2007)*, Pittsburgh, PA, June 2007.

[113] C. Hoofnagle. Measuring identity theft at top banks (Version 1.5). Working Paper Paper 45, University of California, Berkeley, eScholarship Repository, 2008.

[114] Information Systems Audit and Control Association. Telephone survey conducted by MARC Research, October 2007. Find information at `http://biz.yahoo.com/bw/071031/20071031005079.html?.v=1`.

[115] C. Jaeger, O. Renn, E. Rosa, and T. Webler. *Risk, uncertainty, and rational action.* Earthscan Publications, London, UK, 2001.

[116] A. Jaquith. *Security Metrics: Replacing Fear, Uncertainty, and Doubt.* Pearson Education, Upper Saddle River, NJ, 2007.

[117] V. Jayaswal, W. Yurcik, and D. Doss. Internet hack back: Counter attacks as self-defense or vigilantism? In *Proceedings of the International Symposium on Technology and Society (ISTAS)*, pages 380–386, Raleigh, NC, June 2002.

[118] L. Jiang, V. Anantharam, and J. Walrand. Efficiency of selfish investment in network security. In *Proceedings of the 2008 Workshop on the Economics of Networks, Systems, and Computation (NetEcon'08)*, pages 31–36, Seattle, WA, August 2008.

[119] T. Jordan and P. Taylor. A sociology of hackers. *The Sociological Review*, 46(4):757–780, November 1998.

[120] A. Juels. RFID security and privacy: A research survey. *IEEE Journal on Selected Areas in Communications*, 24(2):381–394, February 2006.

[121] Kabooza. Global backup survey: About backup habits, risk factors, worries and

data loss of home PCs, January 2009. Available at: `http://www.kabooza.com/globalsurvey.html`.

[122] D. Kahneman and A. Tversky. *Choices, values and frames*. Cambridge University Press, Cambridge, UK, 2000.

[123] C. Kanich, C. Kreibich, K. Levchenko, B. Enright, G. Voelker, V. Paxson, and S. Savage. Spamalytics: An empirical analysis of spam marketing conversion. In *Proceedings of the 15th ACM Conference on Computer and Communications Security (CCS)*, pages 3–14, Alexandria, VA, October 2008.

[124] S. Karau and K. Williams. Social loafing: A meta-analytic review and theoretical integration. *Journal of Personality and Social Psychology*, 65(4):681–706, October 1993.

[125] M. Kearns and L. Ortiz. Algorithms for interdependent security games. In S. Thrun, L. Saul, and B. Schölkopf, editors, *Advances in Neural Information Processing Systems 16*, pages 561–568. MIT Press, Cambridge, MA, 2004.

[126] O. Kerr. The fourth amendment in cyberspace: Can encryption create a reasonable expectation of privacy? *Connecticut Law Review*, 33(2):503–533, Winter 2001.

[127] J.P. Kesan, R.P. Majuca, and W.J. Yurcik. Three economic arguments for cyberinsurance. In A. Chander, L. Gelman, and M. Radin, editors, *Securing Privacy in the Internet Age*, pages 345–366. Stanford University Press, Stanford, CA, 2008.

[128] K. Konrad and S. Skaperdas. Self-insurance and self-protection: A nonexpected utility analysis. *The GENEVA Papers on Risk and Insurance – Theory*, 18(2):131–146, December 1993.

[129] E. Koutsoupias and C. Papadimitriou. Worst-case equilibria. In *Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science (STACS 99)*, pages 404–413, Trier, Germany, May 1999.

[130] C. Kreibich, C. Kanich, K. Levchenko, B. Enright, G. Voelker, V. Paxson, and S. Savage. Spamcraft: An inside look at spam campaign orchestration. In *Proceedings of the 2nd USENIX Workshop on Large-Scale Exploits and Emergent Threats (LEET '09)*, Boston, MA, April 2009.

[131] J. Kuang, R. Weber, and J. Dana. How effective is advice from interested parties? an experimental test using a pure coordination game. *Journal of Economic Behavior and Organization*, 62(4):591–604, April 2007.

[132] H. Kunreuther and G. Heal. Interdependent security. *Journal of Risk and Uncertainty*, 26(2–3):231–249, March 2003.

[133] H. Kunreuther, G. Silvasi, E. Bradlow, and D. Small. Deterministic and stochastic prisoner's dilemma games: Experiments in interdependent security. NBER Working Paper No. T0341, August 2007.

[134] J. Kwong and K. Wong. The role of ratio differences in the framing of numeri-

cal information. *International Journal of Research in Marketing*, 23(4):385–394, December 2006.

[135] J. Laffont. *The economics of uncertainty and information.* MIT Press, Cambridge, MA, 1989.

[136] C. Landwehr, A. Bull, J. McDermott, and W. Choi. A taxonomy of computer program security flaws. *ACM Computing Surveys*, 26(3):211–254, September 1994.

[137] P. Laskowski and J. Chuang. The economics of virtualization. Technical report, School of Information, University of California, Berkeley, February 2008.

[138] M. Lelarge and J. Bolot. A local mean field analysis of security investments in networks. In *Proceedings of the 2008 Workshop on the Economics of Networks, Systems, and Computation (NetEcon'08)*, pages 25–30, Seattle, WA, August 2008.

[139] M. Lelarge and J. Bolot. Network externalities and the deployment of security features and protocols in the Internet. In *Proceedings of the International Conference on Measurement and Modeling of Computer Systems (SIGMETRICS)*, pages 37–48, Annapolis, MD, June 2008.

[140] M. Lettau and H. Uhlig. Rules of thumb versus dynamic programming. *American Economic Review*, 89(1):148–174, March 1999.

[141] S. Levy. *Hackers: Heroes of the computer revolution.* Penguin Books, Harmondsworth, UK, 1984.

[142] B. Lieberman. Experimental studies of conflict in some two-person and three-person games. In J. Criswell, H. Solomon, and P. Suppes, editors, *Mathematical Models in Small Group Processes*, pages 203–220. Stanford University Press, Stanford, CA, 1962.

[143] Y. Liu, C. Comaniciu, and H. Man. A Bayesian game approach for intrusion detection in wireless ad hoc networks. In *Proceedings of the Workshop on Game Theory for Communications and Networks*, Pisa, Italy, 2006. Article No. 4.

[144] T. Loder, M. van Alstyne, and R. Wash. An economic solution to unsolicited communications. *Advances in Economic Analysis and Policy*, 6(1), 2006.

[145] S. Malphrus. The "I Love You" computer virus and the financial services industry, May 2000. Testimony before the Subcommittee on Financial Institutions of the Committee on Banking, Housing, and Urban Affairs, U.S. Senate. `http://www.federalreserve.gov/BoardDocs/testimony/2000/20000518.htm`.

[146] M. Mavronicolas, V. Papadopoulou, A. Philippou, and P. Spirakis. A graph-theoretic network security game. *International Journal of Autonomous and Adaptive Communications Systems*, 1(4):390–410, November 2008.

[147] M. Mavronicolas, V. Papadopoulou, A. Philippou, and P. Spirakis. A network game with attackers and a defender. *Algorithmica*, 51(3):315–341, July 2008.

[148] McAfee/NCSA. Online safety study, October 2007. Available at: http://download.mcafee.com/products/manuals/en-us/McAfeeNCSA_Analysis09-25-07.pdf.

[149] T. McCabe. A software complexity measure. *IEEE Transactions on Software Engineering*, SE-2(4):308–320, December 1976.

[150] J. McCalley, V. Vittal, and N. Abi-Samra. Overview of risk based security assessment. In *Proceedings of the 1999 IEEE Power Engineering Society Summer Meeting*, pages 173–178, Edmonton, Canada, July 1999.

[151] C. McDonald, P. Hawkes, and J. Pieprzyk. SHA-1 collisions now 2^{52}. In *(Rump Session) Proceedings of the 28th Annual International Conference on the Theory and Applications of Cryptographic Techniques (Eurocrypt)*, Cologne, Germany, April 2009.

[152] R. McKelvey and T. Palfrey. Quantal response equilibria for normal form games. *Games and Economic Behavior*, 10(1):6–38, July 1995.

[153] D. Meier. Changing with the times: How the government must adapt to prevent the publication of its secrets. *The Review of Litigation*, 28(1):203–240, Fall 2008.

[154] D. Meier, Y. Oswald, S. Schmid, and R. Wattenhofer. On the windfall of friendship: Inoculation strategies on social networks. In *Proceedings of the Ninth ACM Conference on Electronic Commerce (EC'08)*, pages 294–301, Chicago, IL, July 2008.

[155] S. Miltchev, S. Ioannidis, and A. Keromytis. A study of the relative costs of network security protocols. In *Proceedings of the 2002 USENIX Annual Technical Conference*, pages 41–48, San Francisco, CA, August 2002.

[156] K. Mitnick and W. Simon. *The Art of Deception: Controlling the Human Element of Security.* Wiley Publishing, Indianapolis, IN, 2002.

[157] D. Moore, V. Paxson, S. Savage, C. Shannon, S. Staniford, and N. Weaver. Inside the Slammer worm. *IEEE Security & Privacy*, 1(4):33–39, July 2003.

[158] D. Moore, C. Shannon, and J. Brown. Code-Red: A case study on the spread and victims of an internet worm. In *Proceedings of 2nd ACM/USENIX Internet Measurement Workshop*, pages 273–284, Marseille, France, November 2002.

[159] T. Moore and R. Clayton. Examining the impact of website take-down on phishing. In *Proceedings of the Anti-Phishing Working Group 2nd Annual eCrime Researchers Summit*, pages 1–13, Pittsburgh, PA, October 2007.

[160] T. Moscibroda, S. Schmid, and R. Wattenhofer. When selfish meets evil: Byzantine players in a virus inoculation game. In *Proceedings of the 25th Annual ACM Symposium on Principles of Distributed Computing (PODC'06)*, pages 35–44, Denver, CO, July 2006.

[161] J. Nash. Non-cooperative games. *Annals of Mathematics*, 54(2):286–295, September 1951.

[162] NCSA/Symantec. Home user study, October 2008. Available at: `http://staysafeonline.org/`.

[163] P. Neumann. Risks to the public. *ACM SIGSOFT Software Engineering Notes*, 34(2):15–24, March 2009.

[164] A. O'Donnell and H. Sethu. On achieving software diversity for improved network security using distributed coloring algorithms. In *Proceedings of the 11th ACM Conference on Computer and Communications Security (CCS)*, pages 121–131, Washington, DC, October 2004.

[165] T. O'Donoghue and M. Rabin. Doing it now or later. *American Economic Review*, 89(1):103–124, March 1999.

[166] Organisation for Economic Co-operation and Development. Measuring security and trust in the online environment: A view using official data, January 2008. Document was prepared by Martin Schaaper, published under the responsibility of the Secretary-General of the OECD, and discussed by the Working Party on Indicators for the Information Society.

[167] A. Ozment and S. Schechter. Bootstrapping the adoption of internet security protocols. In *Proceedings of the Fifth Workshop on the Economics of Information Security (WEIS)*, Cambridge, UK, June 2006.

[168] P. Paruchuri, J. Pearce, J. Marecki, M. Tambe, F. Ordonez, and S. Kraus. Play-

ing games for security: An efficient exact algorithm for solving Bayesian Stackelberg games. In *Proceedings of the 7th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2008)*, pages 895–902, Estoril, Portugal, May 2008.

[169] C. Pautasso and E. Wilde. Why is the web loosely coupled? A multi-faceted metric for service design. In *Proceedings of the 18th International Conference on World Wide Web (WWW2009)*, pages 911–920, Madrid, Spain, April 2009.

[170] F. Piessens. A taxonomy of causes of software vulnerabilities in Internet software. In *Supplementary Proceedings of the 13th International Symposium on Software Reliability Engineering (ISSRE)*, pages 47–52, Annapolis, MD, November 2002.

[171] I. Png, C. Wang, and Q. Wang. The deterrent and displacement effects of information security enforcement: International evidence. *Journal of Management Information Systems*, 25(2):125–144, Fall 2008.

[172] B. Preneel. The state of cryptographic hash functions. In I. Damgård, editor, *Lectures on Data Security: Modern Cryptology in Theory and Practice (Lecture Notes in Computer Science; Vol. 1561)*, pages 158–182. Springer Verlag, Heidelberg, Germany, 1999.

[173] N. Provos. A virtual honeypot framework. In *Proceedings of the 13th USENIX Security Symposium*, pages 1–14, San Diego, CA, August 2004.

[174] N. Provos, D. McNamee, P. Mavrommatis, K. Wang, and N. Modadugu. The ghost in the browser: Analysis of web-based malware. In *Proceedings of the First USENIX Workshop on Hot Topics in Understanding Botnets (HotBots'07)*, Cambridge, MA, April 2007.

[175] G. Quattrone and A. Tversky. Contrasting rational and psychological analyses of political choice. *The American Political Science Review*, 82(3):719–736, September 1988.

[176] M. Rabin. Psychology and economics. *Journal of Economic Literature*, 36(1):11–46, March 1998.

[177] R. Radner. Collusive behavior in noncooperative epsilon-equilibria of oligopolies with long but finite lives. *Journal of Economic Theory*, 22:136–154, April 1980.

[178] M. Rajab, J. Zarfoss, F. Monrose, and A. Terzis. My botnet is bigger than yours (maybe, better than yours): Why size estimates remain challenging. In *Proceedings of the First Workshop on Hot Topics in Understanding Botnets (HotBots)*, Cambridge, MA, April 2007.

[179] M. Randazzo, M. Keeney, E. Kowalski, D. Cappelli, and A. Moore. Insider threat study: Illicit cyber activity in the banking and finance sector. Technical report, CERT Coordination Center, Software Engineering Institute, June 2005. Technical report No. CMU/SEI-2004-TR-021. Available at `http://www.sei.cmu.edu/pub/documents/04.reports/pdf/04tr021.pdf`.

[180] R. Richardson. 2008 CSI Computer Crime and Security Survey, 2008. Conducted and published by the Computer Security Institute, San Francisco, CA.

[181] N. Rosasco and D. Larochelle. How and why more secure technologies succeed in legacy markets: Lessons from the success of SSH. In *Proceedings of the Second Annual Workshop on Economics and Information Security (WEIS)*, College Park, MD, May 2003.

[182] B. Rowe, D. Reeves, and M. Gallaher. The role of Internet service providers in security, June 2009. Prepared in cooperation with the Institute of Homeland Security Solutions, Research Triangle Park, NC.

[183] A. Rowstron and P. Druschel. Storage management and caching in PAST, a large-scale, persistent peer-to-peer storage utility. *ACM SIGOPS Operating Systems Review*, 35(5):188–201, December 2001.

[184] R. Rue and S.L. Pfleeger. Making the best use of cybersecurity economic models. *IEEE Security & Privacy*, 7(4):52–60, July–August 2009.

[185] J. Saltzer, D. Reed, and D. Clark. End-to-end arguments in system design. *ACM Transactions on Computer Systems*, 2(4):277–288, November 1984.

[186] R. Sandhu. Good-enough security: Toward a pragmatic business-driven discipline. *IEEE Internet Computing*, 7(1):66–68, January–February 2003.

[187] T. Sandler and K. Hartley. Economics of alliances: The lessons for collective action. *Journal of Economic Literature*, 39(3):869–896, September 2001.

[188] T. Sandler, F. Sterbenz, and J. Posnett. Free riding and uncertainty. *Economic Review*, 31(8):1605–1617, December 1987.

[189] S. Schechter and M. Smith. How much security is enough to stop a thief? In *Proceedings of the Seventh International Financial Cryptography Conference (FC'03)*, pages 122–137, Gosier, Guadeloupe, January 2003.

[190] T. Schelling. *The Strategy of Conflict.* Oxford University Press, Oxford, UK, 1965.

[191] T. Schelling. *Arms and Influence (The Henry L. Stimson Lectures Series).* Yale University Press, New Haven, CT, 1967.

[192] B. Schneier. *Beyond Fear.* Springer Verlag, New York, NY, 2006.

[193] B. Schneier and A. Shostack. Breaking up is hard to do: Modeling security threats for smart cards. In *Proceedings of the USENIX Workshop on Smartcard Technology*, Chicago, IL, May 1999.

[194] R. Selten. What is bounded rationality? In G. Gigerenzer and R. Selten, editors, *Bounded rationality: The adaptive toolbox*, pages 13–36. The MIT Press, Cambridge, MA, 2002.

[195] J. Shachat and J.T. Swarthout. Do we detect and exploit mixed strategy play by

opponents? *Mathematical Methods of Operations Research*, 59(3):359–373, July 2004.

[196] L. Sheeran, A. Sasse, J. Rimmer, and I. Wakeman. How web browsers shape users' understanding of networks. *The Electronic Library*, 20(1):35–42, 2002.

[197] J.F. Shogren. On increased risk and the voluntary provision of public goods. *Social Choice and Welfare*, 7(3):221–229, September 1990.

[198] H. Simon. Altruism and economics. *American Economic Review*, 83(2):156–161, May 1993.

[199] A. Sood. From vulnerability to patch: The window of exposure. *Network Security*, 2009(2), February 2009.

[200] S. Spiekermann, J. Grossklags, and B. Berendt. E-privacy in 2nd generation e-commerce: Privacy preferences versus actual behavior. In *Proceedings of the Third ACM Conference on Electronic Commerce (EC'01)*, pages 38–47, Tampa, FL, October 2001.

[201] J. Stanton, K. Stam, P. Mastrangelo, and J. Jolton. Analysis of end user security behaviors. *Computers & Security*, 2(24):124–133, March 2005.

[202] B. Sterling. *The Hacker Crackdown: Law and Disorder on the Electronic Frontier*. Bantam Books, New York, NY, 1992.

[203] E. Stone, F. Yates, and A. Parker. Risk Communication: Absolute versus Relative Expressions of Low-Probability Risks. *Organizational Behavior and Human Decision Processes*, 3(60):387–408, December 1994.

[204] B. Stone-Gross, M. Cova, L. Cavallaro, B. Gilbert, M. Szydlowski, R. Kemmerer, C. Kruegel, and G. Vigna. Your botnet is my botnet: Analysis of a botnet takeover. Technical report, University of California, Santa Barbara, April 2009. Available at `http://www.cs.ucsb.edu/~seclab/projects/torpig/torpig.pdf`.

[205] D. Straub. Effective IS Security: An Empirical Study. *Information Systems Research*, 3(1):255–276, September 1990.

[206] P. Swire. 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, Fall 2004.

[207] P. Swire. No Cop on the Beat: Underenforcement in E-Commerce and Cybercrime. *Journal on Telecommunications and High Technology Law*, 7(1):107–126, Winter 2009.

[208] Cisco Systems. 2008 Annual security report, December 2008. Available available for download at `http://www.cisco.com/go/securityreport`.

[209] R. Telang and S. Wattal. An empirical analysis of the impact of software vulnerabil-

ity announcements on firm stock price. *IEEE Transactions on Software Engineering*, 33(8):544–557, August 2007.

[210] The Honeynet Project. Know your enemy: The tools and methodologies of the script-kiddie, July 2000. Available online at `http://project.honeynet.org/papers/enemy/`.

[211] J.B. Van Huyck, R.C. Battallio, and R.O. Beil. Tacit coordination games, strategic uncertainty, and coordination failure. *American Economic Review*, 80(1):234–248, March 1990.

[212] H.R. Varian. System reliability and free riding. In L.J. Camp and S. Lewis, editors, *Economics of Information Security (Advances in Information Security, Volume 12)*, pages 1–15. Kluwer Academic Publishers, Dordrecht, The Netherlands, 2004.

[213] R. Wagner and J. Pescatore. Management update: Increase security in desktop computing through diversity, 2003. Available from Gartner Research, Inc.

[214] X. Wang, Y. Yin, and H. Yu. Finding collisions in the full SHA-1. In *Proceedings of the 25th Annual International Cryptology Conference (CRYPTO 2005)*, pages 17–36, Santa Barbara, CA, August 2005.

[215] R. Wash. Mental models of home computer security. In *Proceedings of the Symposium On Usable Privacy and Security (SOUPS 2008), Poster Session*, Pittsburgh, PA, July 2008.

[216] N. Weaver, D. Ellis, S. Staniford, and V. Paxson. Worms vs. perimeters: The case for hard-LANs. In *Proceedings of the 12th Annual IEEE Symposium on High Performance Interconnects*, pages 70–76, Stanford, CA, August 2004.

[217] N. Weaver and V. Paxson. A worst-case worm. In *Proceedings of the Third Annual Workshop on Economics and Information Security (WEIS)*, Minneapolis, MN, May 2004. Available at `http://www.dtc.umn.edu/weis2004/weaver.pdf`.

[218] E. Weyuker. Evaluating software complexity measures. *IEEE Transactions on Software Engineering*, 14(9):1357–1365, September 1988.

[219] A. Yannacopoulos, C. Lambrinoudakis, S. Gritzalis, S. Xanthopoulos, and S. Katsikas. Modeling privacy insurance contracts and their utilization in risk management for ICT firms. In *Proceedings of the 13th European Symposium on Research in Computer Security (ESORICS 2008)*, pages 207–222, Málaga, Spain, October 2008.

[220] L. Zhuang, J. D. Tygar, and R. Dhamija. Injecting heterogeneity through protocol randomization. *International Journal of Network Security*, 4(1):45–58, January 2007.

[221] M. Zviran and W. Haga. Password security: An empirical study. *Journal of Management Information Systems*, 15(4):161–186, Spring 1999.

Appendix A

Derivations and tables for complete/incomplete information security game

A.1 Derivations for weakest-link game

Weakest-link security game. Derivations for total expected game payoffs, conditioned on other players: The following derivations refer to Table A.4.

Cases (WC1) and (WC2a) and (WI1):

$$\text{Payoff}[passivity] \cdot Pr[passivity] + \text{Payoff}[insurance] \cdot Pr[insurance]$$

$$+ \text{Payoff}[protection] \cdot Pr[protection]$$

$$= [M - E[p_i] \cdot L] \cdot \left[\frac{c}{L}\right] + [M - c] \cdot \left[1 - \frac{c}{L}\right] + [M - b - E[p_i] \cdot L] \cdot [0]$$

$$= \left[M - \left(\frac{c}{2L}\right) \cdot L\right] \cdot \left[\frac{c}{L}\right] + [M - c] \cdot \left[1 - \frac{c}{L}\right]$$

$$= M - \frac{c^2}{2L} - c + \frac{c^2}{L}$$

$$= M - c + \frac{c^2}{2L}$$

Case (WC2b):

$$\text{Payoff}[passivity] \cdot Pr[passivity] + \text{Payoff}[insurance] \cdot Pr[insurance]$$

$$+ \text{Payoff}[protection] \cdot Pr[protection]$$

$$= [M - E[p_i] \cdot L] \cdot \left[\frac{b}{L}\right] + [M - c] \cdot [0] + [M - b] \cdot \left[1 - \frac{b}{L}\right]$$

$$= \left[M - \left(\frac{b}{2L}\right) \cdot L\right] \cdot \left[\frac{b}{L}\right] + [M - b] \cdot \left[1 - \frac{b}{L}\right]$$

$$= M - \frac{b^2}{2L} - b + \frac{b^2}{L}$$

$$= M - b + \frac{b^2}{2L}$$

Case (WI2):

$$\text{Payoff}[passivity] \cdot Pr[passivity] + \text{Payoff}[insurance] \cdot Pr[insurance]$$

$$+ \text{Payoff}[protection] \cdot Pr[protection]$$

$$= [M - E[p_i] \cdot L] \cdot \left[\frac{c}{L}\right] + [M - c] \cdot \left[1 - \frac{c}{L}\right]$$

$$+ \left[M - b - E[p_i] \cdot L \left(1 - \left(1 - \frac{b}{L}\right)^{N-1}\right)\right] \cdot [0]$$

$$= \left[M - \left(\frac{c}{2L}\right) \cdot L\right] \cdot \left[\frac{c}{L}\right] + [M - c] \cdot \left[1 - \frac{c}{L}\right]$$

$$= M - \frac{c^2}{2L} - c + \frac{c^2}{L}$$

$$= M - c + \frac{c^2}{2L}$$

Case (WI3):

$$\text{Payoff}[passivity] \cdot Pr[passivity] + \text{Payoff}[insurance] \cdot Pr[insurance]$$

$$+ \text{Payoff}[protection] \cdot Pr[protection]$$

$$= [M - E[p_i] \cdot L] \cdot \left[\frac{b}{L\left(1 - \frac{b}{L}\right)^{N-1}}\right] + [M - c] \cdot \left[1 - \frac{c - b}{L\left(1 - (1 - \frac{b}{L})^{N-1}\right)}\right]$$

$$+ \left[M - b - E[p_i] \cdot L \left(1 - \left(1 - \frac{b}{L}\right)^{N-1}\right)\right] \cdot \left[\frac{c - b}{L\left(1 - (1 - \frac{b}{L})^{N-1}\right)} - \frac{b}{L\left(1 - \frac{b}{L}\right)^{N-1}}\right]$$

$$= \left[M - \left(\frac{b}{2L\left(1-\frac{b}{L}\right)^{N-1}}\right)\cdot L\right]\cdot\left[\frac{b}{L\left(1-\frac{b}{L}\right)^{N-1}}\right]$$

$$+ [M-c]\cdot\left[1-\frac{c-b}{L\left(1-(1-\frac{b}{L})^{N-1}\right)}\right]$$

$$+\left[M-b-\frac{1}{2}\left(\frac{c-b}{L\left(1-(1-\frac{b}{L})^{N-1}\right)}+\frac{b}{L\left(1-\frac{b}{L}\right)^{N-1}}\right)\cdot L\left(1-\left(1-\frac{b}{L}\right)^{N-1}\right)\right]$$

$$\cdot\left[\frac{c-b}{L\left(1-(1-\frac{b}{L})^{N-1}\right)}-\frac{b}{L\left(1-\frac{b}{L}\right)^{N-1}}\right]$$

$$= M-\frac{b^2}{2L\left(1-\frac{b}{L}\right)^{2N-2}}-c+\frac{c^2-bc}{L\left(1-(1-\frac{b}{L})^{N-1}\right)}$$

$$-b\cdot\left[\frac{c-b}{L\left(1-(1-\frac{b}{L})^{N-1}\right)}-\frac{b}{L\left(1-\frac{b}{L}\right)^{N-1}}\right]$$

$$-\frac{L}{2}\left(1-(1-\frac{b}{L})^{N-1}\right)\cdot\left[\left(\frac{c-b}{L\left(1-\left(1-\frac{b}{L}\right)^{N-1}\right)}\right)^2-\left(\frac{b}{L\left(1-\frac{b}{L}\right)^{N-1}}\right)^2\right]$$

$$= M-\frac{b^2}{2L\left(1-\frac{b}{L}\right)^{2N-2}}-c+\frac{c^2-bc}{L\left(1-(1-\frac{b}{L})^{N-1}\right)}-\frac{bc-b^2}{L\left(1-\left(1-\frac{b}{L}\right)^{N-1}\right)}$$

$$+\frac{b^2}{L\left(1-\frac{b}{L}\right)^{N-1}}-\frac{(c-b)^2}{2L\left(1-\left(1-\frac{b}{L}\right)^{N-1}\right)}+\frac{b^2}{2L\left(1-\frac{b}{L}\right)^{2N-2}}$$

$$= M-c+\frac{c^2-2bc+b^2}{L\left(1-(1-\frac{b}{L})^{N-1}\right)}+\frac{b^2}{L\left(1-\frac{b}{L}\right)^{N-1}}-\frac{(c-b)^2}{2L\left(1-\left(1-\frac{b}{L}\right)^{N-1}\right)}$$

$$= M-c+\frac{b^2}{2L\left(1-\frac{b}{L}\right)^{N-1}}+\frac{(c-b)^2}{2L\left(1-\left(1-\frac{b}{L}\right)^{N-1}\right)}$$

Case (WI4):

$$\text{Payoff}[passivity] \cdot Pr[passivity] + \text{Payoff}[insurance] \cdot Pr[insurance]$$

$$+ \text{Payoff}[protection] \cdot Pr[protection]$$

$$= [M - E[p_i] \cdot L] \cdot \left[\frac{b}{L\left(1-\frac{b}{L}\right)^{N-1}}\right] + [M-c] \cdot [0]$$

$$+ \left[M - b - E[p_i] \cdot L\left(1-\left(1-\frac{b}{L}\right)^{N-1}\right)\right] \cdot \left[1 - \frac{b}{L\left(1-\frac{b}{L}\right)^{N-1}}\right]$$

$$= \left[M - \left(\frac{b}{2L\left(1-\frac{b}{L}\right)^{N-1}}\right) \cdot L\right] \cdot \left[\frac{b}{L\left(1-\frac{b}{L}\right)^{N-1}}\right]$$

$$+ \left[M - b - \frac{1}{2}\left(1 + \frac{b}{L\left(1-\frac{b}{L}\right)^{N-1}}\right) \cdot L\left(1-\left(1-\frac{b}{L}\right)^{N-1}\right)\right] \cdot \left[1 - \frac{b}{L\left(1-\frac{b}{L}\right)^{N-1}}\right]$$

$$= M - \frac{b^2}{2L\left(1-\frac{b}{L}\right)^{2N-2}} - b \cdot \left[1 - \frac{b}{L\left(1-\frac{b}{L}\right)^{N-1}}\right]$$

$$- \frac{L}{2}\left(1-(1-\frac{b}{L})^{N-1}\right) \cdot \left[1 - \left(\frac{b}{L\left(1-\frac{b}{L}\right)^{N-1}}\right)^2\right]$$

$$= M - \frac{b^2}{2L\left(1-\frac{b}{L}\right)^{2N-2}} - b + \frac{b^2}{L\left(1-\frac{b}{L}\right)^{N-1}} - \frac{L}{2}\left(1-\left(1-\frac{b}{L}\right)^{N-1}\right)$$

$$+ \frac{L}{2}\left(1-(1-\frac{b}{L})^{N-1}\right)\left(\frac{b^2}{L^2\left(1-\frac{b}{L}\right)^{2N-2}}\right)$$

$$= M - \frac{b^2}{2L\left(1-\frac{b}{L}\right)^{2N-2}} - b + \frac{b^2}{L\left(1-\frac{b}{L}\right)^{N-1}} - \frac{L}{2}\left(1-\left(1-\frac{b}{L}\right)^{N-1}\right)$$

$$+ \frac{b^2}{2L\left(1-\frac{b}{L}\right)^{2N-2}} - \frac{b^2}{2L\left(1-\frac{b}{L}\right)^{N-1}}$$

$$= M - b - \frac{L}{2}\left(1-\left(1-\frac{b}{L}\right)^{N-1}\right) + \frac{b^2}{2L\left(1-\frac{b}{L}\right)^{N-1}}$$

Weakest-link security game. Derivations for total expected game payoffs, not conditioned on other players: The following derivation refers to Table A.5. To remove dependence on p_j for $j \neq i$ in case WC2, we simply take a weighted sum of the total payoffs for cases WC2a and WC2b, where the weight is determined by the probability of $\min_{j \neq i} p_j < \frac{b}{L}$ assuming that each p_j is drawn from the uniform distribution over $[0, 1]$ (and assuming $b \leq c$).

We have:

Case (WC2):

Probability[Case (WC2a)] · ExPayoff[Case (WC2a)]

+ Probability[Case (WC2b)] · ExPayoff[Case (WC2b)]

$$= \left(1 - \left(1 - \frac{b}{L}\right)^{N-1}\right) \cdot \left[M - c + \frac{c^2}{2L}\right] + \left[\left(1 - \frac{b}{L}\right)^{N-1}\right] \cdot \left[M - b + \frac{b^2}{2L}\right]$$

$$= M - c + \frac{c^2}{2L} + \left(c - \frac{c^2}{2L}\right)\left(1 - \frac{b}{L}\right)^{N-1} - \left(b - \frac{b^2}{2L}\right)\left(1 - \frac{b}{L}\right)^{N-1}$$

$$= M - c + \frac{c^2}{2L} + \left(c - b - \frac{c^2 - b^2}{2L}\right)\left(1 - \frac{b}{L}\right)^{N-1}$$

$$= M - c + \frac{c^2}{2L} + (c - b)\left(1 - \frac{c+b}{2L}\right)\left(1 - \frac{b}{L}\right)^{N-1}$$

A.2 Derivations for best shot game

Best shot security game. Derivations for total expected game payoffs, conditioned on other players: The following derivations refer to Table A.9.

Cases (BC1) and (BI1):

$$\text{Payoff}[passivity] \cdot Pr[passivity] + \text{Payoff}[insurance] \cdot Pr[insurance]$$

$$+ \text{Payoff}[protection] \cdot Pr[protection]$$

$$= [M - E[p_i] \cdot L] \cdot \left[\frac{c}{L}\right] + [M - c] \cdot \left[1 - \frac{c}{L}\right] + [M - b] \cdot [0]$$

$$= \left[M - \left(\frac{c}{2L}\right) \cdot L\right] \cdot \left[\frac{c}{L}\right] + [M - c] \cdot \left[1 - \frac{c}{L}\right]$$

$$= M - \frac{c^2}{2L} - c + \frac{c^2}{L}$$

$$= M - c + \frac{c^2}{2L}$$

Case (BC2a):

$$\text{Payoff}[passivity] \cdot Pr[passivity] + \text{Payoff}[insurance] \cdot Pr[insurance]$$

$$+ \text{Payoff}[protection] \cdot Pr[protection]$$

$$= [M - E[p_i] \cdot L] \cdot \left[\frac{b}{L}\right] + [M - c] \cdot [0] + [M - b] \cdot \left[1 - \frac{b}{L}\right]$$

$$= \left[M - \left(\frac{b}{2L}\right) \cdot L\right] \cdot \left[\frac{b}{L}\right] + [M - b] \cdot \left[1 - \frac{b}{L}\right]$$

$$= M - \frac{b^2}{2L} - b + \frac{b^2}{L}$$

$$= M - b + \frac{b^2}{2L}$$

Case (BC2b):

$$\text{Payoff}[passivity] \cdot Pr[passivity] + \text{Payoff}[insurance] \cdot Pr[insurance]$$

$$+ \text{Payoff}[protection] \cdot Pr[protection]$$

$$= [M] \cdot [1] + [M - c] \cdot [0] + [M - b] \cdot [0]$$

$$= M$$

Case (BI2):

$$\text{Payoff}[passivity] \cdot Pr[passivity] + \text{Payoff}[insurance] \cdot Pr[insurance]$$

$$+ \text{Payoff}[protection] \cdot Pr[protection]$$

$$= \left[M - E[p_i] \cdot L \left(\frac{b}{L} \right)^{N-1} \right] \cdot [1] + [M - c] \cdot [0] + [M - b] \cdot [0]$$

$$= M - \left(\frac{1}{2} \right) \cdot L \left(\frac{b}{L} \right)^{N-1}$$

$$= M - \frac{L}{2} \left(\frac{b}{L} \right)^{N-1}$$

Best shot security game. Derivations for total expected game payoffs, not conditioned on other players: The following derivation refers to Table A.10. To remove dependence on p_j for $j \neq i$ in case BC2, we simply take a weighted sum of the total payoffs for cases BC2a and BC2b, where the weight is determined by the probability of $\min_{j \neq i} p_j < \frac{b}{L}$ assuming that each p_j is drawn from the uniform distribution over $[0, 1]$.

We have:

Case (BC2):

Probability[Case (BC2a)] · ExPayoff[Case (BC2a)]

+ Probability[Case (BC2b)] · ExPayoff[Case (BC2b)]

$$= \left(\frac{b}{L}\right)^{N-1} \cdot \left[M - b + \frac{b^2}{2L}\right] + \left[1 - \left(\frac{b}{L}\right)^{N-1}\right] \cdot [M]$$

$$= M - b\left(\frac{b}{L}\right)^{N-1} + \frac{b^2}{2L}\left(\frac{b}{L}\right)^{N-1}$$

$$= M - b\left(1 - \frac{b}{2L}\right)\left(\frac{b}{L}\right)^{N-1}$$

A.3 Derivations for total effort game

Total Effort security game. Derivations for total expected game payoffs, conditioned on other players: The following derivations refer to Table A.14.

Case (TC1): $c < b$

$$\text{Payoff}[passivity] \cdot Pr[passivity] + \text{Payoff}[insurance] \cdot Pr[insurance]$$

$$+ \text{Payoff}[protection] \cdot Pr[protection]$$

$$= [M - E[p_i] \cdot L] \cdot \left[\frac{c}{L}\right] + [M - c] \cdot \left[1 - \frac{c}{L}\right] + \left[M - b - E[p_i] \cdot L\left(1 - \frac{1}{N}\right)\right] \cdot [0]$$

$$= \left[M - \left(\frac{c}{2L}\right) \cdot L\right] \cdot \left[\frac{c}{L}\right] + [M - c] \cdot \left[1 - \frac{c}{L}\right]$$

$$= M - \frac{c^2}{2L} - c + \frac{c^2}{L}$$

$$= M - c + \frac{c^2}{L}$$

Cases (TC2) and (TC5):

$$\text{Payoff}[passivity] \cdot Pr[passivity] + \text{Payoff}[insurance] \cdot Pr[insurance]$$

$$+ \text{Payoff}[protection] \cdot Pr[protection]$$

$$= \left[M - E[p_i] \cdot L\left(1 - \frac{K}{N}\right)\right] \cdot \left[\frac{c}{L\left(1 - \frac{K}{N}\right)}\right] + [M - c] \cdot \left[1 - \frac{c}{L\left(1 - \frac{K}{N}\right)}\right]$$

$$+ \left[M - b - E[p_i] \cdot L\left(1 - \frac{K+1}{N}\right)\right] \cdot [0]$$

$$= \left[M - \left(\frac{c}{2L\left(1 - \frac{K}{N}\right)}\right) \cdot L\left(1 - \frac{K}{N}\right)\right] \cdot \left[\frac{c}{L\left(1 - \frac{K}{N}\right)}\right] + [M - c] \cdot \left[1 - \frac{c}{L\left(1 - \frac{K}{N}\right)}\right]$$

$$= M - \frac{c^2}{2L\left(1 - \frac{K}{N}\right)} - c + \frac{c^2}{L\left(1 - \frac{K}{N}\right)}$$

$$= M - c + \frac{c^2}{2L\left(1 - \frac{K}{N}\right)}$$

Case (TC3):

$$\text{Payoff}[passivity] \cdot Pr[passivity] + \text{Payoff}[insurance] \cdot Pr[insurance]$$

$$+ \text{Payoff}[protection] \cdot Pr[protection]$$

$$= \left[M - E[p_i] \cdot L\left(1 - \frac{K}{N}\right)\right] \cdot \left[\frac{bN}{L}\right] + [M - c] \cdot \left[1 - \frac{c-b}{L\left(1 - \frac{K+1}{N}\right)}\right]$$

$$+ \left[M - b - E[p_i] \cdot L\left(1 - \frac{K+1}{N}\right)\right] \cdot \left[\frac{c-b}{L\left(1 - \frac{K+1}{N}\right)} - \frac{bN}{L}\right]$$

$$= \left[M - \left[\frac{bN}{2L}\right] \cdot L\left(1 - \frac{K}{N}\right)\right] \cdot \left[\frac{bN}{L}\right] + [M - c] \cdot \left[1 - \frac{c-b}{L\left(1 - \frac{K+1}{N}\right)}\right]$$

$$+ \left[M - b - \frac{1}{2}\left[\frac{c-b}{L\left(1 - \frac{K+1}{N}\right)} + \frac{bN}{L}\right] \cdot L\left(1 - \frac{K+1}{N}\right)\right] \cdot \left[\frac{c-b}{L\left(1 - \frac{K+1}{N}\right)} - \frac{bN}{L}\right]$$

$$= M - \frac{b^2N^2}{2L}\left(1 - \frac{K}{N}\right) - c + \frac{c^2 - bc}{L\left(1 - \frac{K+1}{N}\right)}$$

$$- b\left(\frac{c-b}{L\left(1 - \frac{K+1}{N}\right)} - \frac{bN}{L}\right) - \frac{L}{2}\left(1 - \frac{K+1}{N}\right)\left(\left(\frac{c-b}{L\left(1 - \frac{K+1}{N}\right)}\right)^2 - \left(\frac{bN}{L}\right)^2\right)$$

$$= M - \frac{b^2N^2}{2L}\left(1 - \frac{K}{N}\right) - c + \frac{c^2 - bc}{L\left(1 - \frac{K+1}{N}\right)} - b\left(\frac{c-b}{L\left(1 - \frac{K+1}{N}\right)}\right) + \frac{b^2N}{L}$$

$$- \frac{(c-b)^2}{2L\left(1 - \frac{K+1}{N}\right)} + \frac{b^2N^2}{2L}\left(1 - \frac{K+1}{N}\right)$$

$$= M - c + \frac{(c-b)^2}{L\left(1 - \frac{K+1}{N}\right)} + \frac{b^2N}{L} - \frac{(c-b)^2}{2L\left(1 - \frac{K+1}{N}\right)} - \frac{b^2N}{2L}$$

$$= M - c + \frac{b^2N}{2L} + \frac{(c-b)^2}{2L\left(1 - \frac{K+1}{N}\right)}$$

Case (TC4):

$$\text{Payoff}[passivity] \cdot Pr[passivity] + \text{Payoff}[insurance] \cdot Pr[insurance]$$

$$+ \text{Payoff}[protection] \cdot Pr[protection]$$

$$= \left[M - E[p_i] \cdot L\left(1 - \frac{K}{N}\right)\right] \cdot \left[\frac{bN}{L}\right] + [M - c] \cdot [0]$$

$$+ \left[M - b - E[p_i] \cdot L\left(1 - \frac{K+1}{N}\right)\right] \cdot \left[1 - \frac{bN}{L}\right]$$

$$= \left[M - \left[\frac{bN}{2L}\right] \cdot L\left(1 - \frac{K}{N}\right)\right] \cdot \left[\frac{bN}{L}\right] + [M - c] \cdot [0]$$

$$+ \left[M - b - \left[\frac{1}{2}\left(1 + \frac{bN}{L}\right)\right] \cdot L\left(1 - \frac{K+1}{N}\right)\right] \cdot \left[1 - \frac{bN}{L}\right]$$

$$= M - \frac{b^2N^2}{2L}\left(1 - \frac{K}{N}\right) - b\left(1 - \frac{bN}{L}\right) - \frac{L}{2}\left(1 - \frac{K+1}{N}\right)\left(1 - \frac{b^2N^2}{L^2}\right)$$

$$= M - \frac{b^2N^2}{2L}\left(1 - \frac{K}{N}\right) - b + \frac{b^2N}{L} - \frac{L}{2}\left(1 - \frac{K+1}{N}\right) + \frac{b^2N^2}{2L}\left(1 - \frac{K+1}{N}\right)$$

$$= M - b + \frac{b^2N}{L} - \frac{L}{2}\left(1 - \frac{K+1}{N}\right) - \frac{b^2N}{2L}$$

$$= M - b - \frac{L}{2}\left(1 - \frac{K+1}{N}\right) + \frac{b^2N}{2L}$$

Case (TC6):

$$\text{Payoff}[passivity] \cdot Pr[passivity] + \text{Payoff}[insurance] \cdot Pr[insurance]$$

$$+ \text{Payoff}[protection] \cdot Pr[protection]$$

$$= \left[M - E[p_i] \cdot L\left(1 - \frac{K}{N}\right)\right] \cdot [1] + [M - c] \cdot [0]$$

$$+ \left[M - b - E[p_i] \cdot L\left(1 - \frac{K+1}{N}\right)\right] \cdot [0]$$

$$= M - \frac{L}{2}\left(1 - \frac{K}{N}\right)$$

Case (TI1): $c < b$

$$\text{Payoff}[passivity] \cdot Pr[passivity] + \text{Payoff}[insurance] \cdot Pr[insurance]$$

$$+ \text{Payoff}[protection] \cdot Pr[protection]$$

$$= [M - E[p_i] \cdot L] \cdot \left[\frac{c}{L}\right] + [M - c] \cdot \left[1 - \frac{c}{L}\right] + \left[M - b - E[p_i] \cdot L\left(1 - \frac{1}{N}\right)\right] \cdot [0]$$

$$= \left[M - \left(\frac{c}{2L}\right) \cdot L\right] \cdot \left[\frac{c}{L}\right] + [M - c] \cdot \left[1 - \frac{c}{L}\right]$$

$$= M - \frac{c^2}{2L} - c + \frac{c^2}{L}$$

$$= M - c + \frac{c^2}{L}$$

Cases (TI2) and (TI5):

$$\text{Payoff}[passivity] \cdot Pr[passivity] + \text{Payoff}[insurance] \cdot Pr[insurance]$$

$$+ \text{Payoff}[protection] \cdot Pr[protection]$$

$$= \left[M - E[p_i]\left(b + \frac{L-b}{N}\right)\right] \cdot \left[\frac{c}{b + \frac{L-b}{N}}\right] + [M-c] \cdot \left[1 - \frac{c}{b + \frac{L-b}{N}}\right]$$

$$+ \left[M - b - E[p_i]\left(b - \frac{b}{N}\right)\right] \cdot [0]$$

$$= \left[M - \left(\frac{c}{2\left(b + \frac{L-b}{N}\right)}\right) \cdot \left(b + \frac{L-b}{N}\right)\right] \cdot \left[\frac{c}{b + \frac{L-b}{N}}\right] + [M-c] \cdot \left[1 - \frac{c}{b + \frac{L-b}{N}}\right]$$

$$= M - \frac{c^2}{2\left(b + \frac{L-b}{N}\right)} - c + \frac{c^2}{b + \frac{L-b}{N}}$$

$$= M - c + \frac{c^2}{2\left(b + \frac{L-b}{N}\right)}$$

Case (TI3)

$$\text{Payoff}[passivity] \cdot Pr[passivity] + \text{Payoff}[insurance] \cdot Pr[insurance]$$

$$+ \text{Payoff}[protection] \cdot Pr[protection]$$

$$= \left[M - E[p_i]\left(b + \frac{L-b}{N}\right)\right] \cdot \left[\frac{bN}{L}\right] + [M-c] \cdot \left[1 - \frac{c-b}{b-\frac{b}{N}}\right]$$

$$+ \left[M - b - E[p_i]\left(b - \frac{b}{N}\right)\right] \cdot \left[\frac{c-b}{b-\frac{b}{N}} - \frac{bN}{L}\right]$$

$$= \left[M - \left(\frac{bN}{2L}\right) \cdot \left(b + \frac{L-b}{N}\right)\right] \cdot \left[\frac{bN}{L}\right] + [M-c] \cdot \left[1 - \frac{c-b}{b-\frac{b}{N}}\right]$$

$$+ \left[M - b - \frac{1}{2}\left(\frac{bN}{L} + \frac{c-b}{b-\frac{b}{N}}\right) \cdot \left(b - \frac{b}{N}\right)\right] \cdot \left[\frac{c-b}{b-\frac{b}{N}} - \frac{bN}{L}\right]$$

$$= M - \frac{b^2N^2}{2L^2}\left(b + \frac{L-b}{N}\right) - c + c\left(\frac{c-b}{b-\frac{b}{N}}\right)$$

$$- b\left(\frac{c-b}{b-\frac{b}{N}} - \frac{bN}{L}\right) - \frac{1}{2}\left(b - \frac{b}{N}\right)\left(\frac{(c-b)^2}{\left(b-\frac{b}{N}\right)^2} - \frac{b^2N^2}{L^2}\right)$$

$$= M - \frac{b^2N^2}{2L^2}\left(b - \frac{b}{N}\right) - \frac{b^2N}{2L} - c + c\left(\frac{c-b}{b-\frac{b}{N}}\right)$$

$$- b\left(\frac{c-b}{b-\frac{b}{N}}\right) + \frac{b^2N}{L} - \frac{(c-b)^2}{2\left(b-\frac{b}{N}\right)} + \frac{b^2N^2}{2L^2}\left(b - \frac{b}{N}\right)$$

$$= M - c + c\left(\frac{c-b}{b-\frac{b}{N}}\right) - b\left(\frac{c-b}{b-\frac{b}{N}}\right) + \frac{b^2N}{2L} - \frac{(c-b)^2}{2\left(b-\frac{b}{N}\right)}$$

$$= M - c + \frac{(c-b)^2}{b-\frac{b}{N}} + \frac{b^2N}{2L} - \frac{(c-b)^2}{2\left(b-\frac{b}{N}\right)}$$

$$= M - c + \frac{b^2N}{2L} + \frac{(c-b)^2}{2\left(b-\frac{b}{N}\right)}$$

Case (TI4)

$$\text{Payoff}[passivity] \cdot Pr[passivity] + \text{Payoff}[insurance] \cdot Pr[insurance]$$

$$+ \text{Payoff}[protection] \cdot Pr[protection]$$

$$= \left[M - E[p_i]\left(b + \frac{L-b}{N}\right)\right] \cdot \left[\frac{bN}{L}\right] + [M-c] \cdot [0]$$

$$+ \left[M - b - E[p_i]\left(b - \frac{b}{N}\right)\right] \cdot \left[1 - \frac{bN}{L}\right]$$

$$= \left[M - \left(\frac{bN}{2L}\right)\left(b + \frac{L-b}{N}\right)\right] \cdot \left[\frac{bN}{L}\right]$$

$$+ \left[M - b - \left(\frac{bN}{2L} + \frac{1}{2}\right)\left(b - \frac{b}{N}\right)\right] \cdot \left[1 - \frac{bN}{L}\right]$$

$$= M - \frac{b^2N^2}{2L^2}\left(b - \frac{b}{N}\right) - \frac{b^2N}{2L} - b\left(1 - \frac{bN}{L}\right) - \frac{bN}{2L}\left(b - \frac{b}{N}\right)\left(1 - \frac{bN}{L}\right)$$

$$- \frac{1}{2}\left(b - \frac{b}{N}\right)\left(1 - \frac{bN}{L}\right)$$

$$= M - \frac{b^2N^2}{2L^2}\left(b - \frac{b}{N}\right) - \frac{b^2N}{2L} - b + \frac{b^2N}{L} - \frac{bN}{2L}\left(b - \frac{b}{N}\right) + \frac{b^2N^2}{2L^2}\left(b - \frac{b}{N}\right)$$

$$- \frac{1}{2}\left(b - \frac{b}{N}\right) + \frac{bN}{2L}\left(b - \frac{b}{N}\right)$$

$$= M + \frac{b^2N}{2L} - b - \frac{1}{2}\left(b - \frac{b}{N}\right)$$

$$= M - b - \frac{1}{2}\left(b - \frac{b}{N}\right) + \frac{b^2N}{2L}$$

Case (TI6):

$$\text{Payoff}[passivity] \cdot Pr[passivity] + \text{Payoff}[insurance] \cdot Pr[insurance]$$

$$+ \text{Payoff}[protection] \cdot Pr[protection]$$

$$= \left[M - E[p_i]\left(b + \frac{L-b}{N}\right)\right] \cdot [1] + [M - c] \cdot [0] + \left[M - b - E[p_i]\left(b - \frac{b}{N}\right)\right] \cdot [0]$$

$$= M - \frac{1}{2}\left(b + \frac{L-b}{N}\right)$$

Total effort security game. Derivations for total expected game payoffs, not conditioned on other players: The following derivation refers to Table A.15. For the total effort game, the dependence on other players is noted in terms of the integer K, the number of players other than player i who choose protection. To remove dependence on this K we must compute an appropriate expected value. To begin we rewrite each of the case expressions as a linear constraint on K. After doing this it becomes clear that cases TC2 through TC4 are mutually exclusive and exhaustive in terms of K, and similarly for cases TC5 and TC6. We define case TC2-4 to be the union of cases TC2, TC3, and TC4. similarly, we define case TC5-6 to be the union of cases TC5 and TC6. Now to compute an expected payoff for case TC2-4, we take the sum, over all possible values k for K, of the probability that exactly k players protect, times the payoff for this k (considering the case TC2, TC3, or TC4, that such a choice of $K = k$ determines). We proceed similarly to compute the expected payoff for case TC5-6.

To obtain the expected payoff for TC2-4 we compute:

Case (TC2-4):

$$\sum_{k=0}^{N-1} Pr[k] \cdot \text{ Payoff assuming TC2-4 and that } K = k$$

$$= \sum_{k=0}^{\lfloor N-\frac{c}{b} \rfloor} Pr[k] \cdot \left(M - c + \frac{c^2}{2L\left(1-\frac{k}{N}\right)} \right)$$

$$+ \sum_{k=\lfloor N-\frac{c}{b}+1 \rfloor}^{\lfloor N-1-\frac{N}{L}(c-b) \rfloor} Pr[k] \cdot \left(M - c + \frac{b^2 N}{2L} + \frac{(c-b)^2}{2L\left(1-\frac{k+1}{N}\right)} \right)$$

$$+ \sum_{k=\lfloor N-\frac{N}{L}(c-b) \rfloor}^{N-1} Pr[k] \cdot \left(M - b - \frac{L}{2}\left(1 - \frac{k+1}{N}\right) + \frac{b^2 N}{2L} \right)$$

and to obtain the expected payoff for TC5-6 we compute:

Case (TC5-6):

$$\sum_{k=0}^{N-1} Pr[k] \cdot \text{ Payoff assuming TC5-6 and that } K = k$$

$$= \sum_{k=0}^{\lfloor N-\frac{cN}{L} \rfloor} Pr[k] \cdot \left(M - c + \frac{c^2}{2L\left(1-\frac{k}{N}\right)} \right)$$

$$+ \sum_{k=\lfloor N-\frac{cN}{L}+1 \rfloor}^{N-1} Pr[k] \cdot \left(M - \frac{L}{2N}(N-k) \right)$$

where as before,

$$Pr[k] = \binom{N-1}{k} \left(1-\frac{b}{L}\right)^k \left(\frac{b}{L}\right)^{N-1-k}$$

is the probability that exactly k players other than player i choose protection.

A.4 Tabulated results

In the following, we provide the tabulated results for the complete and incomplete information analysis conducted in Chapter 4 and utilized in Chapter 5.

Table A.1: Weakest-link security game: Payoffs for different strategies under different information conditions

Case	Information Type	Payoff Passivity	Payoff Self-Insurance	Payoff Protection
$c < b$	Complete	$M - p_i L$	$M - c$	$M - b - p_i L$
$b \leq c$ and $\min_{j \neq i} p_j < \frac{b}{L}$	Complete	$M - p_i L$	$M - c$	$M - b - p_i L$
$b \leq c$ and $\frac{b}{L} \leq \min_{j \neq i} p_j$	Complete	$M - p_i L$	$M - c$	$M - b$
$c < b$	Incomplete	$M - p_i L$	$M - c$	$M - b - p_i L$
$b \leq c$	Incomplete	$M - p_i L$	$M - c$	$M - b$ $-p_i L \left(1 - \left(1 - \frac{b}{L}\right)^{N-1}\right)$

Table A.2: Weakest-link security game: Conditions to select protection, self-insurance or passivity strategies

Case	Information Type	Conditions Passivity	Conditions Self-Insurance	Conditions Protection
$c < b$	Complete	$p_i < \frac{c}{L}$	$p_i \geq \frac{c}{L}$	NEVER!
$b \leq c$ and $\min_{j \neq i} p_j < \frac{b}{L}$	Complete	$p_i < \frac{c}{L}$	$p_i \geq \frac{c}{L}$	NEVER!
$b \leq c$ and $\frac{b}{L} \leq \min_{j \neq i} p_j$	Complete	$p_i < \frac{b}{L}$	NEVER!	$p_i \geq \frac{b}{L}$
$c < b$	Incomplete	$p_i < \frac{c}{L}$	$p_i > \frac{c}{L}$	NEVER!
$b \leq c \leq \frac{b}{\left(1-\frac{b}{L}\right)^{N-1}}$	Incomplete	$p_i < \frac{c}{L}$	$p_i \geq \frac{c}{L}$	NEVER!
$\frac{b}{\left(1-\frac{b}{L}\right)^{N-1}} < c$	Incomplete	$p_i < \frac{b}{L\left(1-\frac{b}{L}\right)^{N-1}}$	$p_i > \frac{c-b}{L\left(1-\left(1-\frac{b}{L}\right)^{N-1}\right)}$	$\frac{b}{L\left(1-\frac{b}{L}\right)^{N-1}} \leq p_i$, $p_i \leq \frac{c-b}{L\left(1-(1-\frac{b}{L})^{N-1}\right)}$

Table A.3: Weakest-link security game: Probabilities to select protection, self-insurance or passivity strategies

	Case	Information Type	Probability Passivity	Probability Self-Insurance	Probability Protection
WC1	$c < b$	Complete	$\frac{c}{L}$	$1 - \frac{c}{L}$	0
WC2a	$b \leq c$ and $\min_{j \neq i} p_j < \frac{b}{L}$	Complete	$\frac{c}{L}$	$1 - \frac{c}{L}$	0
WC2b	$b \leq c$ and $\frac{b}{L} \leq \min_{j \neq i} p_j$	Complete	$\frac{b}{L}$	0	$1 - \frac{b}{L}$
WI1	$c < b$	Incomplete	$\frac{c}{L}$	$1 - \frac{c}{L}$	0
WI2	$b \leq c \leq \frac{b}{\left(1-\frac{b}{L}\right)^{N-1}}$	Incomplete	$\frac{c}{L}$	$1 - \frac{c}{L}$	0
WI3	$\frac{b}{\left(1-\frac{b}{L}\right)^{N-1}} < c$ and $c < b + L\left(1 - \left(1 - \frac{b}{L}\right)^{N-1}\right)$	Incomplete	$\frac{b}{L\left(1-\frac{b}{L}\right)^{N-1}}$	$1 - \frac{c-b}{L\left(1-\left(1-\frac{b}{L}\right)^{N-1}\right)}$	$\frac{c-b}{L\left(1-(1-\frac{b}{L})^{N-1}\right)} - \frac{b}{L\left(1-\frac{b}{L}\right)^{N-1}}$
WI4	$\frac{b}{\left(1-\frac{b}{L}\right)^{N-1}} < c$ and $b + L\left(1 - \left(1 - \frac{b}{L}\right)^{N-1}\right) \leq c$	Incomplete	$\frac{b}{L\left(1-\frac{b}{L}\right)^{N-1}}$	0	$1 - \frac{b}{L\left(1-\frac{b}{L}\right)^{N-1}}$

Table A.4: Weakest-link security game: Total expected game payoffs, conditioned on other players

	Case	Information Type	Total Expected Payoff for player i (conditioned on other players)
WC1	$c < b$	Complete	$M - c + \frac{c^2}{2L}$
WC2a	$b \leq c$ and $\min_{j \neq i} p_j < \frac{b}{L}$	Complete	$M - c + \frac{c^2}{2L}$
WC2b	$b \leq c$ and $\frac{b}{L} \leq \min_{j \neq i} p_j$	Complete	$M - b + \frac{b^2}{2L}$
WI1	$c < b$	Incomplete	$M - c + \frac{c^2}{2L}$
WI2	$b \leq c \leq \frac{b}{\left(1-\frac{b}{L}\right)^{N-1}}$	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	$\frac{b}{\left(1-\frac{b}{L}\right)^{N-1}} < b + L\left(1 - \left(1 - \frac{b}{L}\right)^{N-1}\right) \leq c$	Incomplete	$M - b - \frac{L}{2}\left(1 - \left(1 - \frac{b}{L}\right)^{N-1}\right) + \frac{b^2}{2L\left(1-\frac{b}{L}\right)^{N-1}}$

Table A.5: Weakest-link security game: Total expected game payoffs, not conditioned on other players

	Case	Information Type	Total Expected Payoff for player i (not conditioned on other players)
WC1	$c < b$	Complete	$M - c + \frac{c^2}{2L}$
WC2	$b \leq c$	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	$b \leq c \leq \frac{b}{\left(1-\frac{b}{L}\right)^{N-1}}$	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	$\frac{b}{\left(1-\frac{b}{L}\right)^{N-1}} < b + L\left(1 - \left(1 - \frac{b}{L}\right)^{N-1}\right) \leq c$	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	$M - c + \frac{c^2}{2L}$
WN2	$b \leq c$	Naive	$M - b + \frac{b^2}{2L} - \frac{L}{2}\left(1 - \frac{b^2}{L^2}\right)\left(1 - \left(1 - \frac{b}{L}\right)^{N-1}\right)$

Table A.6: Best shot security game: Payoffs for different strategies under different information conditions

	Case	Information Type	Payoff Passivity	Payoff Self-Insurance	Payoff Protection
	$c < b$	Complete	$M - p_i L$	$M - c$	$M - b$
	$b \leq c$ and $\max_{j \neq i} p_j < \frac{b}{L}$	Complete	$M - p_i L$	$M - c$	$M - b$
	$b \leq c$ and $\frac{b}{L} \leq \max_{j \neq i} p_j$	Complete	M	$M - c$	$M - b$
	$c < b$	Incomplete	$M - p_i L$	$M - c$	$M - b$
	$b \leq c$	Incomplete	$M - p_i L \left(\frac{b}{L}\right)^{N-1}$	$M - c$	$M - b$

Table A.7: Best shot security game: Conditions to select protection, self-insurance or passivity strategies

	Case	Information Type	Conditions Passivity	Conditions Self-Insurance	Conditions Protection
	$c < b$	Complete	$p_i < c/L$	$p_i \geq c/L$	NEVER!
	$b \leq c$ and $\max_{j \neq i} p_j < b/L$	Complete	$p_i < b/L$	NEVER!	$p_i \geq b/L$
	$b \leq c$ and $b/L \leq \max_{j \neq i} p_j$	Complete	ALWAYS!	NEVER!	NEVER!
	$c < b$	Incomplete	$p_i < c/L$	$p_i \geq c/L$	NEVER!
	$b \leq c$	Incomplete	ALWAYS!	NEVER!	NEVER!

Table A.8: Best shot security game: Probabilities to select protection, self-insurance or passivity strategies

	Case	Information Type	Probability Passivity	Probability Self-Insurance	Probability Protection
BC1	$c < b$	Complete	$\frac{c}{L}$	$1 - \frac{c}{L}$	0
BC2a	$b \leq c$ and $\max_{j \neq i} p_j < \frac{b}{L}$	Complete	$\frac{b}{L}$	0	$1 - \frac{b}{L}$
BC2b	$b \leq c$ and $\frac{b}{L} \leq \max_{j \neq i} p_j$	Complete	1	0	0
BI1	$c < b$	Incomplete	$\frac{c}{L}$	$1 - \frac{c}{L}$	0
BI2	$b \leq c$	Incomplete	1	0	0

Table A.9: Best shot security game: Total expected game payoffs, conditioned on other players

	Case	Information Type	Total Expected Payoff
BC1	$c < b$	Complete	$M - c + \frac{c^2}{2L}$
BC2a	$b \leq c$ and $\max_{j \neq i} p_j < \frac{b}{L}$	Complete	$M - b + \frac{b^2}{2L}$
BC2b	$b \leq c$ and $\frac{b}{L} \leq \max_{j \neq i} p_j$	Complete	M
BI1	$c < b$	Incomplete	$M - c + \frac{c^2}{2L}$
BI2	$b \leq c$	Incomplete	$M - \frac{L}{2}\left(\frac{b}{L}\right)^{N-1}$

Table A.10: Best shot security game: Total expected game payoffs, not conditioned on other players

	Case	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	$M - c + \frac{c^2}{2L}$
BN2	$b \leq c$	Naive	$M - b + \frac{b^2}{2L}$

Table A.11: Total effort security game: Payoffs for different strategies under different information conditions

	Case	Information Type	Payoff Passivity	Payoff Self-Insurance	Payoff Protection
	$c < b$	Complete	$M - p_i L$	$M - c$	$M - b - p_i L\,(1 - 1/N)$
	$b \leq c$	Complete	$M - p_i L\,(1 - K/N)$	$M - c$	$M - b - p_i L\,(1 - (K+1)/N)$
	$c < b$	Incomplete	$M - p_i L$	$M - c$	$M - b - p_i L\,(1 - 1/N)$
	$b \leq c$	Incomplete	$M - p_i\,(b + (L - b)/N)$	$M - c$	$M - b - p_i\,(b - b/N)$

Table A.12: Total effort security game: Conditions to select protection, self-insurance or passivity strategies

	Case	Information Type	Conditions Passivity	Conditions Self-Insurance	Conditions Protection
	$c < b$	Complete	$p_i < \frac{c}{L}$	$p_i \geq \frac{c}{L}$	NEVER!
	$b \leq c \leq b(N-K)$	Complete	$p_i < \frac{c}{L\left(1-\frac{K}{N}\right)}$	$p_i \geq \frac{c}{L\left(1-\frac{K}{N}\right)}$	NEVER!
	$b(N-K) < c$	Complete	$p_i < \frac{bN}{L}$	$p_i > \frac{c-b}{L\left(1-\frac{K+1}{N}\right)}$	$\frac{bN}{L} \leq p_i \leq \frac{c-b}{L\left(1-\frac{K+1}{N}\right)}$
	$c < b$	Incomplete	$p_i < \frac{c}{L}$	$p_i \geq \frac{c}{L}$	NEVER!
	$b \leq c \leq b + \frac{b^2}{L}(N-1)$	Incomplete	$p_i < \frac{c}{b+\frac{L-b}{N}}$	$p_i \geq \frac{c}{b+\frac{L-b}{N}}$	NEVER!
	$b + \frac{b^2}{L}(N-1) < c$	Incomplete	$p_i < \frac{bN}{L}$	$p_i > \frac{c-b}{b-\frac{b}{N}}$	$\frac{bN}{L} \leq p_i \leq \frac{c-b}{b-\frac{b}{N}}$

Table A.13: Total effort security game: Probabilities to select protection, self-insurance or passivity strategies

	Case	Information Type	Probability Passivity	Probability Self-Insurance	Probability Protection
TC1	$c < b$	Complete	$\frac{c}{L}$	$1 - \frac{c}{L}$	0
TC2	$bN \leq L$ and $b \leq c \leq b(N-K)$	Complete	$\frac{c}{L\left(1-\frac{K}{N}\right)}$	$1 - \frac{c}{L\left(1-\frac{K}{N}\right)}$	0
TC3	$bN \leq L$ and $b(N-K) < c < b + L\left(1 - \frac{K+1}{N}\right)$	Complete	$\frac{bN}{L}$	$1 - \frac{c-b}{L\left(1-\frac{K+1}{N}\right)}$	$\frac{c-b}{L\left(1-\frac{K+1}{N}\right)} - \frac{bN}{L}$
TC4	$bN \leq L$ and $b + L\left(1 - \frac{K+1}{N}\right) \leq c$	Complete	$\frac{bN}{L}$	0	$1 - \frac{bN}{L}$
TC5	$L < bN$ and $b \leq c < L\left(1 - \frac{K}{N}\right)$	Complete	$\frac{c}{L\left(1-\frac{K}{N}\right)}$	$1 - \frac{c}{L\left(1-\frac{K}{N}\right)}$	0
TC6	$L < bN$ and $L\left(1 - \frac{K+1}{N}\right) < c$	Complete	1	0	0
TI1	$c < b$	Incomplete	$\frac{c}{L}$	$1 - \frac{c}{L}$	0
TI2	$bN \leq L$ and $b \leq c \leq b + \frac{b^2}{L}(N-1)$	Incomplete	$\frac{c}{b+\frac{L-b}{N}}$	$1 - \frac{c}{b+\frac{L-b}{N}}$	0
TI3	$bN \leq L$ and $b + \frac{b^2}{L}(N-1) < c < 2b - \frac{b}{N}$	Incomplete	$\frac{bN}{L}$	$1 - \frac{c-b}{b-\frac{b}{N}}$	$\frac{c-b}{b-\frac{b}{N}} - \frac{bN}{L}$
TI4	$bN \leq L$ and $2b - \frac{b}{N} \leq c$	Incomplete	$\frac{bN}{L}$	0	$1 - \frac{bN}{L}$
TI5	$L < bN$ and $b \leq c < b + \frac{L-b}{N}$	Incomplete	$\frac{c}{b+\frac{L-b}{N}}$	$1 - \frac{c}{b+\frac{L-b}{N}}$	0
TI6	$L < bN$ and $b + \frac{L-b}{N} \leq c$	Incomplete	1	0	0

Table A.14: Total Effort security game: Total expected game payoffs, conditioned on other players

	Case	Information Type	Total Expected Payoff
TC1	$c < b$	Complete	$M - c + \frac{c^2}{2L}$
TC2	$bN \leq L$ and $b \leq c \leq b(N-K)$	Complete	$M - c + \frac{c^2}{2L\left(1-\frac{K}{N}\right)}$
TC3	$bN \leq L$ and $b(N-K) < c < b + L\left(1 - \frac{K+1}{N}\right)$	Complete	$M - c + \frac{b^2N}{2L} + \frac{(c-b)^2}{2L\left(1-\frac{k+1}{N}\right)}$
TC4	$bN \leq L$ and $b + L\left(1 - \frac{K+1}{N}\right) \leq c$	Complete	$M - b - \frac{L}{2}\left(1 - \frac{K+1}{N}\right) + \frac{b^2N}{2L}$
TC5	$L < bN$ and $b \leq c \leq L\left(1 - \frac{K}{N}\right)$	Complete	$M - c + \frac{c^2}{2L\left(1-\frac{K}{N}\right)}$
TC6	$L < bN$ and $L\left(1 - \frac{K}{N}\right) < c$	Complete	$M - \frac{L}{2}\left(1 - \frac{K}{N}\right)$
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	$M - \frac{1}{2}\left(b + \frac{L-b}{N}\right)$

Table A.15: Total effort security game: Total expected game payoffs, not conditioned on other players

	Case	Information Type	Total Expected Payoff
TC1	$c < b$	Complete	$M - c + \frac{c^2}{2L}$
TC2-4	$bN \leq L$ and $b \leq c$	Complete	$*\sum_{k=0}^{\lfloor N-\frac{c}{b} \rfloor} Pr[k] \cdot \left(M - c + \frac{c^2}{2L\left(1-\frac{k}{N}\right)} \right)$ $+ \sum_{k=\lfloor N-\frac{c}{b}+1 \rfloor}^{\lfloor N-1-\frac{N}{L}(c-b) \rfloor} Pr[k] \cdot \left(M - c + \frac{b^2 N}{2L} + \frac{(c-b)^2}{2L\left(1-\frac{k+1}{N}\right)} \right)$ $+ \sum_{k=\lfloor N-\frac{N}{L}(c-b) \rfloor}^{N-1} Pr[k] \cdot \left(M - b - \frac{L}{2}\left(1 - \frac{k+1}{N}\right) + \frac{b^2 N}{2L} \right)$
TC5-6	$L < bN$ 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	$M - c + \frac{c^2}{2L}$
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	$M - c + \frac{b^2 N}{2L} + \frac{(c-b)^2}{2\left(b-\frac{b}{N}\right)}$
TI4	$bN \leq L$ and $2b - \frac{b}{N} \leq c$	Incomplete	$M - b - \frac{1}{2}\left(b - \frac{b}{N}\right) + \frac{b^2 N}{2L}$
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	$M - c + \frac{c^2}{2L}$
TN2	$b \leq c$	Naive	$M - b - \frac{1}{2}\left(b - \frac{b}{N}\right) + \frac{b^2}{L}\left(1 - \frac{1}{2N}\right)$

$*Pr[k] = \binom{N-1}{k}\left(1-\frac{b}{L}\right)^k \left(\frac{b}{L}\right)^{N-1-k}$ is the probability that exactly k players other than i choose protection.