Network Externalities and the Deployment of Security Features and Protocols in the Internet

Marc Lelarge
INRIA-ENS
45 rue d'Ulm
Paris, France
marc.lelarge@ens.fr

Jean Bolot
SPRINT
California, USA
bolot@sprint.com

ABSTRACT

Getting new security features and protocols to be widely adopted and deployed in the Internet has been a continuing challenge. There are several reasons for this, in particular economic reasons arising from the presence of network externalities. Indeed, like the Internet itself, the technologies to secure it exhibit network effects: their value to individual users changes as other users decide to adopt them or not. In particular, the benefits felt by early adopters of security solutions might fall significantly below the cost of adoption, making it difficult for those solutions to gain attraction and get deployed at a large scale.

Our goal in this paper is to model and quantify the impact of such externalities on the adoptability and deployment of security features and protocols in the Internet. We study a network of interconnected agents, which are subject to epidemic risks such as those caused by propagating viruses and worms, and which can decide whether or not to invest some amount to deploy security solutions. Agents experience negative externalities from other agents, as the risks faced by an agent depend not only on the choices of that agent (whether or not to invest in self-protection), but also on those of the other agents. Expectations about choices made by other agents then influence investments in self-protection, resulting in a possibly suboptimal outcome overall.

We present and solve an analytical model where the agents are connected according to a variety of network topologies. Borrowing ideas and techniques used in statistical physics, we derive analytic solutions for sparse random graphs, for which we obtain asymptotic results. We show that we can explicitly identify the impact of network externalities on the adoptability and deployment of security features. In other words, we identify both the economic and network properties that determine the adoption of security technologies. Therefore, we expect our results to provide useful guidance for the design of new economic mechanisms and for the development of network protocols likely to be deployed at a large scale.

SIGMETRICS'08, June 2–6, 2008, Annapolis, Maryland, USA.

Categories and Subject Descriptors

J.4 [**Social and Behavioral Sciences**]: Economics; G.2.2 [**Discrete Mathematics**]: Graph Theory—*Network problems*; C.4 [**Computer Systems Organization**]: Performance of Systems—*Modeling techniques*

General Terms

Theory, Economics, Design.

Keywords

Security, Game Theory, Epidemics, Economics, Price of Anarchy, Tipping.

1. INTRODUCTION

Negligent users who do not protect their computer by regularly updating their antivirus software and operating system are clearly putting their own computers at risk. But such users, by connecting to the network a computer which may become a host from which viruses can multiply and spread, also put (a potentially large number of) computers on the network at risk. In this situation, users and computers on the network face *epidemic risks*. Epidemic risks are risks which depend on the behavior of other entities in the network (such as whether or not those entities invest in security solutions to minimize their likelihood of being infected). Our goal is to analyze the strategic behavior of agents facing such epidemic risks. In particular, we characterize the incentives for agents to buy and deploy security technologies.

Understanding and optimizing the behavior of agents facing epidemic risks is an important problem, because *epidemic risks are common and their impact severe*. Epidemic risks have been faced by Internet users for some time now, and they are propagated by a variety of malware such as viruses and worms [22]. Worms in particular have been the subject of much attention since the Morris worm [17] in 1988 (with subsequent high-profile worms such as Code Red in July 2001 [15] focusing renewed attention on the problem) because they can propagate extremely fast [18], so fast that signature-based detection and prevention systems cannot build and update signatures in time to block the worms.

A key characteristic of an Internet worm is self-propagation. That is, active worms can spread rapidly by infecting computer systems and by using infected hosts to disseminate the worms in an automated fashion. There exist a wide variety of worms [21], but most of them can be conveniently divided into two types, namely scan-based and topology-based worms.

A scan-based or scanning worm finds new victims by scanning Internet addresses and then attacking whatever is at a target address. A scanning worm might probe the entire IP address space, or a subset (considered optimal by the worm in some fashion), or the routable address space. When a target address is found, the worm sends out a probe to infect it. After the target is compromised, the worm transfers a copy of itself to that host, and the newly infected host then begins to run the worm program and attempts to compromise further targets. Worms can propagate extremely rapidly: for example, the Slammer worm, one of the fastest scanning worm observed, exploited Microsoft's SQL server with a single 376-byte UDP packet and infected more than 90% of vulnerable hosts in 10 minutes on January 25, 2003 [15].

A topology-based worm spreads through topological neighbors. For example, the Morris worm [17] retrieved the neighbor list from local Unix files /ect/hosts.equiv and /.rhosts and from .forward and .rhosts files. Another topological worm is a SSH worm, which locates new targets by searching its current host for the names and addresses of other hosts which are likely to be susceptible to infection. Email worms are another example of topological worms, they use the address book to send copies of themselves to other email addresses.

The propagation of worms and viruses, but also many other phenomena in the Internet such as the propagation of alerts and patches or of routing updates, can be modeled using epidemic spreads through a network (e.g. [23, 20]). As a result, there is now a vast body of literature on epidemic spreads over a network topology from an initial set of infected nodes to susceptible nodes (see for example [7]). However, *much of that work has focused on modeling and understanding the propagation of the epidemics proper, without considering the impact of network effects and externalities.*

Recent work which did model such effects has been limited to the simple case of two agents, i.e. a two-node network. Specifically, reference [11] proposes a parametric game theoretic model for such a situation. In the model, agents decide whether or not to invest in security and agents face a risk of infection which depends on the state of other agents. The authors show the existence of two Nash equilibria (all agents invest or none invests), and suggest that taxation or insurance would be ways to provide incentives for agents to invest (and therefore reach the "good" Nash equilibrium). However, their approach does not scale to the case of n agents, and does not handle various network topologies connecting those agents. Our work specifically addresses those limitations.

We make three contributions in this paper.

First, we develop a model which captures both the propagation of epidemic risks and the existence of network externalities.

Second, we introduce sophisticated new techniques, such as recursive tree processes and recursive distributional equations [1], to tackle and solve the economic study of epidemic risks on large networks with arbitrary topologies.

Third, we show how to compute the price of anarchy to compare equilibria. Our results provide useful guidance for the design of new economic mechanisms and for the development of network protocols likely to be deployed at a large scale.

The rest of the paper is organized as follows. In Section 2, we describe our model, starting with the epidemic model, then the economic model for the agents, and finally the combined economic model for epidemic risks. In Section 3, we consider the 2-agents model and present known results for optimal self-protection. We also introduce the price of anarchy and the price of stability in our model. In Section 4, we show how the local weak convergence introduced by Aldous and Steele [2] allows us to analyze our model for sparse random graphs. We treat the case of Erdös-Rényi graphs in details. In Section 5, we apply the results obtained to compute asymptotics, as the population grows, of the price of anarchy and the price of stability of our system. We also show the existence of a tipping or cascading phenomenon, which leads to very practical advice on how to increase the overall security of networks. In Section 6, we explain how our results can be generalized to other graphs structures and conclude the paper.

2. A MODEL FOR EPIDEMIC RISKS

2.1 Epidemic model

We first describe our model for the spread of a worm, or of an attack in general. This model may seem simplistic at first but the reader must remember that we will have to construct and analyze, in a second step, an economic model on top of this epidemic model. Furthermore, our goal is not to capture the minutia of worm propagation, but to obtain a (hopefully mathematically tractable) model which captures the salient features of malware propagation in the Internet and provide actionable insights for network and security architects.

To simplify our analysis, we consider one-period probabilistic models for the risk, in which all decisions and outcomes occur in a simultaneous instant.

Let $G = (V, E)$ be a graph on a countable vertex set V. Agents are represented by vertices of the graph. For $i, j \in V$, we write $i \sim j$ if $(i, j) \in E$ and we say that agents i and j are neighbours. The state of agent i is represented by X_i; agent i is infected (respectively healthy) iff $X_i = 1$ (respectively $X_i = 0$). Then any infected agent contaminates neighbours independently of each others with probability q. There is no recovery.

It is convenient to introduce a proper probability space Ω_e for random variables describing the epidemic. We take $\Omega_e = \prod_{(i,j)\in\mathbb{N}\times\mathbb{N}}\{0, 1\}$, points of which are represented as $\omega_e = (B_{ij}, (i, j) \in \mathbb{N} \times \mathbb{N})$. For $i \sim j$, the value $B_{ij} = 1$ corresponds to possible contagion between agents i and j, and $B_{ij} = 0$ corresponds to the absence of contagion between these agents. If $(i, j) \notin E$, then the value of B_{ij} is irrelevant to the problem. We assume that the sequence $(B_{ij}, i < j)$ is a sequence of i.i.d. Bernoulli random variables with parameter q: $\mathbb{P}(B_{ij} = 1) = q = 1 - \mathbb{P}(B_{ij} = 0)$. We define also for $i > j$, $B_{ij} = B_{ji}$.

In order to completely describe the epidemic, we need to specify the 'initial condition': the set of agents that are sick before the contagion process takes place. This initial state is represented by a vector χ; site i is sick (respectively healthy) before the contagion process takes place iff $\chi_i = 1$ (respectively $\chi_i = 0$). Then the fundamental recursion satisfied by the vector $X = (X_i, i \in V)$ is

$$1 - X_i = (1 - \chi_i) \prod_{j \sim i} (1 - B_{ij} X_j). \quad (1)$$

Note that the topology of the underlying graph G is arbitrary. We will discuss different specific cases for G in Sections 4 and 6.1.

For now, we need to specify the vector χ. Each χ_i is related to the economic decision of agent i concerning its investment in security and this decision depends on the behavior of other entities in the network. So the vector χ is responsible for the modeling of the network externalities. It is crucial a crucial part of our model, in fact what makes this model unique. We discuss it in detail next.

2.2 Economic model for the agents

Some economic background on the expected utility model and risk aversion is given in Appendix 8.1. The crucial notion of risk premium (denoted by π) is also defined in Appendix 8.1.

Investments in security involve either self-protection (to reduce the probability of a loss) and/or self-insurance (to reduce the size of a loss). For example, intrusion detection and prevention systems are mechanisms of self-protection. Denial-of-service mitigation systems, traffic engineering solutions, overprovisioning, and public relations companies are mechanisms of self-insurance (overprovisioning to reduce the impact of overloads or attacks, PR firms to reduce the impact of security attack on a company stock price with crafty messages to investors). It is somewhat artificial to distinguish mechanisms that reduce the probability of a loss from mechanisms that reduce the size of the loss, since many mechanisms do both. Nevertheless, we focus on self-protection mechanisms only.

Consider an economy in which each economic agent is endowed with an initial wealth w and faces a potential loss ℓ. An agent's utility $u(y)$ is a function of final wealth y. The utility function is increasing and strictly concave, i.e. agents are risk-averse (see Appendix 8.1). The agents maximize their expected utility of final wealth.

We first look at the problem of optimal self-protection for only one agent. We denote by c the cost of self-protection and by $p(c)$ the corresponding probability of loss. We expect larger investments in self-protection to translate into a lower likelihood of loss, and therefore we reasonably assume that p is a non-increasing function of c. The optimal amount of self-protection is given by the value c^* which maximizes

$$p(c)u(w - \ell - c) + (1 - p(c))u(w - c). \quad (2)$$

Consider the simple case where the loss probability is either one of two values, namely $p(c) = p^N$ if $c < c_t$ or $p(c) = p^S$ if $c > c_t$, with $p^N > p^S$. The optimization problem (2) above becomes easy to solve: indeed, the optimal expenditure is either 0 or c_t.

In the rest of the paper, we assume that the choice of an agent regarding self-protection is a binary choice: either the agent does not invest, or it invests c_t which will be denoted c for simplicity. There are two possible economic states for agent i: if it decides to invest in self-protection, we say that agent is in state S (as in Safe or Secure). If it decides not to invest in self-protection, we say that agent i is in state N (as in Not safe). In state N, the expected utility is $p^N u(w - \ell) + (1 - p^N)u(w)$; whereas in state S, the expected utility is $p^S u(w - \ell - c) + (1 - p^S)u(w - c)$. Using the definition of risk premium in Appendix 8.1, we see that these quantities are equal to $u(w - p^N \ell - \pi(p^N))$ and $u(w - c - p^S \ell - \pi(p^S))$, respectively. Therefore, the optimal strategy is for the agent to invest in self-protection only if the cost for self-protection is less than the threshold

$$c < (p^N - p^S)\ell + \pi(p^N) - \pi(p^S). \quad (3)$$

We now return to our multi-agents setting. We first introduce some notations: for agent i, w_i is the initial wealth, ℓ_i is the potential loss and c_i is the cost for self-protection. Also we denote by π_i the risk premium associated to the initial wealth w_i and the utility function of agent i. Then if p_i^N and p_i^S are the loss probabilities depending on whether agent i invests in self-protection or not, as in (3), the optimal strategy is for the agent i to invest in self-protection only if

$$c_i < (p_i^N - p_i^S)\ell_i + \pi_i(p_i^N) - \pi_i(p_i^S). \quad (4)$$

We denote $D_i = \mathbb{1}(c_i < (p_i^N - p_i^S)\ell_i + \pi_i(p_i^N) - \pi_i(p_i^S))$, so that $D_i = 1$ if the optimal strategy for agent i is to invest in self-protection and $D_i = 0$ if his optimal strategy is not to invest. We assume agents apply best-response updates, $D_i = 1$ if agent i is in state S and $D_i = 0$ if he is in state N.

2.3 Epidemic risks for interconnected agents

We now have a model for epidemic propagation between agents and an economic model for those agents. We next present a combined economic model for epidemic risks for interconnected agents.

There are several ways in which we could combine the propagation and economic models presented earlier. We present here a natural way which depends on two parameters p^- and p^+, $0 \le p^- < p^+ \le 1$. Note that there are two possible ways in which a loss can occur: it can either be caused directly by an agent itself (direct loss), or indirectly by contagion via the actions of others (indirect loss). The variable χ_i (above in Section 2.1) is equal to 1 when a direct loss occurs to agent i and 0 if no direct loss occurs; the variable X_i is equal to 1 when a direct or indirect loss occurs to agent i and 0 otherwise.

For our model, we define the probability of direct loss as follows: if an agent invests in self-protection, then the probability of direct loss is p^-. If an agent does not invest in self-protection, then the probability of direct loss is $p^+ \ge p^-$. Formally, we have

$$\chi_i = B_i^S D_i + B_i^N (1 - D_i), \quad (5)$$

where $(B_i^S, i \in \mathbb{N})$ and $(B_i^N, i \in \mathbb{N})$ are sequences of i.i.d. Bernoulli random variables independent of everything else with respective parameters p^- and p^+.

Note that the quantities p_i^N and p_i^S are given by

$$\begin{aligned} p_i^N &= \mathbb{E}[X_i | D_i = 0] \\ p_i^N &= 1 - \mathbb{E}\left[(1 - B_i^N) \prod_{j \sim i} (1 - B_{ij} X_j)\right] \end{aligned} \quad (6)$$

and

$$\begin{aligned} p_i^S &= \mathbb{E}[X_i | D_i = 1] \\ p_i^S &= 1 - \mathbb{E}\left[(1 - B_i^S)[\prod_{j \sim i} (1 - B_{ij} X_j)\right] \end{aligned} \quad (7)$$

Our model is defined by the graph G (of arbitrary topology) and the set of Equations (1,5,6,7). Let us now specify the probability space, we are working on. We already defined Ω_e which describes the propagation of the epidemic given the graph G and the vector χ. We now need to define the probability space describing the economic behavior of agents, specifically we need to specify the variables appearing in the definition of χ_i in (5). In the rest of this paper, we will make a simplifying assumption, namely we consider a heterogeneous population where agents differ only in self-protection cost and loss sizes. The cost of protection should not exceed the possible loss, hence $0 \leq c_i \leq \ell_i$. The cost c_i and the possible loss ℓ_i are known to agent i and vary among the population. Hence we model this heterogeneous population by taking the sequence $(c_i, \ell_i, i \in \mathbb{N})$ as a sequence of i.i.d. random variables in $\mathbb{R}^2$ independent of everything else. We also assume that the function $\pi_i(x) = \pi(x)$ does not depend on i. We can now introduce the probability space $\Omega_s = \prod_{i\in\mathbb{N}}\{0,1\}^2 \times \mathbb{R}^2$, points of which are represented as $\omega_s = (B_i^S, B_i^N, c_i, \ell_i, i \in \mathbb{N})$. The probability measure taken on Ω_s has been described above. Finally, we define $\Omega = \Omega_e \times \Omega_s$ equipped with the standard product σ-field and with probability the product of the measures defined on Ω_e and Ω_s. Points of Ω are called configuration and represented as $\omega = (B_{ij}, B_i^S, B_i^N, c_i, \ell_i, i, j \in \mathbb{N})$.

Summary of notations:

- c_i and ℓ_i are the cost of self-protection and the amount of loss for agent i;
- π is the risk premium of the agents;
- $p^- = \mathbb{P}(B_i^S = 1)$ is the probability of direct loss when investing in self-protection;
- $p^+ = \mathbb{P}(B_i^N = 1)$ is the probability of direct loss when not investing in self-protection;
- $q = \mathbb{P}(B_{ij} = 1)$ is the probability of contagion;
- Equation (1) is:

$$1 - X_i \quad = \quad (1-\chi_i)\prod_{j\sim i}(1 - B_{ij}X_j)$$

- Equation (5) is:

$$\chi_i \quad = \quad B_i^S D_i + B_i^N(1 - D_i)$$

- Equations (6) and (7) are

$$p_i^N \quad = \quad 1 - \mathbb{E}\left[(1-B_i^N)\prod_{j\sim i}(1 - B_{ij}X_j)\right]$$

$$p_i^S \quad = \quad 1 - \mathbb{E}\left[(1-B_i^S)[\prod_{j\sim i}(1 - B_{ij}X_j)\right];$$

- $D_i = \mathbb{1}(c_i < (p_i^N - p_i^S)\ell_i + \pi(p_i^N) - \pi(p_i^S))$ is one if agent i is in state S and zero otherwise.

So far, we have not yet specified the underlying graph G. We will look at cases where G is random but in all cases, the graph G is independent of the configuration ω. It is important to specify precisely Ω since it will allow us to consider natural couplings between two different graphs say G_1 and G_2 coupled by the same configuration ω.

3. INEFFICIENCY OF EQUILIBRIA

This section relates our model to existing models of the economic, computer science and statistical physics literature. We show how some notions introduced in those areas do translate into our framework. We then analyze in detail our model with 2 agents (this corresponds to the model introduced in [11] for interdependent risks) using standard tools from game theory. However, it turns out that those tools are limited to problems with only a few agents and that new methods are required to study the general case of arbitrarily large networks and large populations of agents. Those methods will be described in Section 4.

3.1 Nash equilibria and price of anarchy

There is a clear analogy between our model and disordered systems studied in statistical physics. Equation (1) describes the 'microscopic' interactions between agents, Equation (5) describes the individual action of each agent and Equations (6) and (7) introduce a coupling among these agents. From those equations we try to compute 'macroscopic' quantities such as the fraction of the population investing in self-protection or the probability of being infected averaged over the population. Goals are similar in statistical physics and our techniques are largely inspired from tools of that domain. We refer to [14] for a discussion of similarities and differences between systems of interacting players maximizing their individual payoffs and particles minimizing their interaction energy.

In the framework of statistical physics, the distribution of a vector $(X_i, i \in G)$ satisfying our model equations would be called a Gibbs measure (or equilibrium measure). It is well-known that some interactions have multiple Gibbs measures and we will see in the next section that for a certain range of parameters, our model with only two agents (i.e. with a graph $G = (\{1,2\},(1,2))$) has two possible Gibbs measures. In economic terms, these two measures correspond to two Nash equilibria. Recall that a Nash equilibrium is a strategy vector $s \in \prod_i\{S_i, N_i\}$ such that no player i can change its chosen strategy from S_i to N_i or N_i to S_i and thereby improve its payoff, assuming that all other players stick to the strategies they have chosen in s. In this paper, we consider only pure strategy equilibria where each agent deterministically plays its chosen strategy [8].

When two equilibria are possible, then it is natural to try to compare their efficiency. A metric for this purpose is particularly useful when the outcome of rational behavior by self-interested agents can be inferior to a centrally designed outcome. Indeed, a key question then is by how much the distributed outcome differs from the centrally designed outcome. The *price of anarchy*, the most popular measure of the inefficiency of equilibria, is defined as the ratio between the worst objective function value of an equilibrium of the game and that of an optimal outcome (possibly centralized in which case it will not be described by the model introduced above). The method introduced in Section 4 lets us compute the price of anarchy for large systems. In our setting, the cost incurred to agent i is $c_i + p_i^S\ell_i + \pi_i(p_i^S)$ if he invests in security and $p_i^N\ell_i + \pi_i(p_i^N)$ otherwise. So for a given equilibrium, we can compute the total cost incurred

Table 1: Probability of states

	S	N
S	$p(S_1, S_2) = 0$	$p(S_1, N_2) = pq$
N	$p(N_1, S_2) = p$	$p(N_1, N_2) = p + (1-p)pq$

to the population. The price of anarchy is the ratio of the largest (among all equilibria) such cost divided by the optimal cost. The price of anarchy is at least 1 and a value close to 1 indicates that the given outcome is approximately optimal. We refer to [16] for an introduction to the inefficiency of equilibria (in particular chapter 17).

A game with multiple equilibria has a large price of anarchy even if only one of its equilibria is highly inefficient. The *price of stability* is a measure of inefficiency designed to differentiate between games in which all equilibria are inefficient and those in which some equilibrium is inefficient. Formally, the price of stability is the ratio between the best objective function value of one of the equilibria and that of an optimal outcome. For a game with multiple equilibria (as it will be the case here), its price of stability is at least as close to 1 as its price of anarchy and it can be much closer (as we will see).

3.2 Interdependent risks

Reference [11] was the first to introduce a model for interdependent security (IDS), specifically a model for two agents faced with interdependent risks, and it proposed a parametric game-theoretic model for such a situation. We now describe it: consider a network of 2 agents sharing one link. We assume that the cost of investing in self-protection is $c_1 = c_2 = c$, and that a direct loss can be avoided with certainty when the agent invests in self-protection.

Four possible states of final wealth of an agent result: without protection, the final wealth is w in case of no loss and $w - \ell$ in case of loss. If an agent invests in protection, his final wealth is $w - c$ in case of no loss and $w - c - \ell$ in case of loss (due to possible contagion).

There are four possible economic states denoted by (d_1, d_2), where $d_i \in \{S_i, N_i\}$, d_i describes the decision of agent i where S_i means that the agent i invests in self-protection, and N_i means that the agent i does not invest in self-protection. Kunreuther and Heal [11] examine the symmetric case when the probability of a direct loss is p for both agents, where $0 < p < 1$. Knowing that one agent has a direct loss, the probability that a loss is caused indirectly by this agent to the other is q, where $0 \le q \le 1$. To completely specify their model, they assume that direct losses and contagions are independent events. Hence with the notations of our model, we have $p^- = 0$, $p^+ = p$ and q is the probability of contagion. The matrix $p(d_1, d_2)$ describing the probability of loss for agent 1, in state (d_1, d_2), is given in Table 2.

The simplest situation of interdependent risks, involving only two agents, can be analyzed using a game-theoretic framework. We now derive the payoff matrix of expected utilities for agents 1 and 2. If both agents invest in self-protection, the expected utility of each agent is $u(w - c)$. If agent 1 invests in self-protection (S_1) but not agent 2 (N_2), then agent 1 is only exposed to the indirect risk pq from agent 2. Thus the expected utility for agent 1 is $(1 - pq)u(w - c) + pqu(w - c - \ell)$ and the expected utility for agent 2 is $(1-p)u(w) + pu(w-\ell)$. If neither agent invests in self-protection, then both are exposed to the additional risk of contamination from the other. Therefore, the expected utilities for both agents are $pu(w-\ell) + (1-p)(pqu(w-\ell) + (1-pq)u(w))$. Table 2 summarizes these results and gives the expected utility of agent 1 for the different choices of the agents. Note that there is a minor difference with [11] due to our use of the concave function u which models the fact that the agents are risk adverse. To recover the model of [11], we just have to take $u(x) = x$ (and drop the assumption that agents are risk adverse), in which case we get $\pi(x) = 0$.

Assuming that both agents decide simultaneously whether or not to invest in self-protection, there is no possibility to cooperate. For investment in self-protection to be a dominant strategy, we need

$$u(w-c) \ge (1-p)u(w) + pu(w-\ell) \text{ and}$$
$$(1-pq)u(w-c) + pqu(w-c-\ell) \ge$$
$$pu(w-\ell) + (1-p)(pqu(w-\ell) + (1-pq)u(w))$$

With the notations introduced earlier, the inequalities above become:

$$\begin{aligned} c &\le p\ell + \pi(p) =: c^1, \\ c &\le p(1-pq)\ell + \pi(p + (1-p)pq) - \pi(pq) =: c^0. \end{aligned}$$

Note that in the particular case of non risk adverse agents, i.e. when $\pi \equiv 0$, we have $c^0 = p(1-pq)\ell < c^1 = p\ell$. In most practical cases, one expects that $c^0 < c^1$ still holds, and the tighter second inequality reflects the possibility of damage caused by the other agent. Therefore, the Nash equilibrium for the game is in the state (S_1, S_2) if $c \le c^0$ and (N_1, N_2) if $c > c^1$. If $c^0 < c \le c^1$, then both equilibria are possible and the solution to the game is indeterminate. More precisely, the situation corresponds to a coordination game.

3.3 Price of anarchy: the 2-agent case

To summarize, we see that the IDS model of [11] fits in our model and that even in the simple case of two agents, different cases have to be considered. We continue the analysis of this simple example and interpret previous results in our framework.

Proposition 1. *Assume that $c^0 < c^1$, then we have*

1. *if $c < c^0$, there exists an unique solution to our model. Both agents invest in self-protection and the expected utility is $u(w - c)$.*
2. *if $c > c^1$, there exists an unique solution to our model. Both agents do not invest in self-protection and the expected utility is $u(w - p^N\ell - \pi(p^N))$ with $p^N = p(1 + q - pq)$.*
3. *if $c^0 \le c \le c^1$, both solutions of points 1 and 2 are valid for our model and they are the only ones.*
4. *for any c, the price of anarchy is given by:*
$$P_a(c) = 1 \vee \mathbb{1}(c \ge c^0)\frac{p^N\ell + \pi(p^N)}{c}.$$
5. *the price of stability is given by:*
$$P_s(c) = 1 \vee \mathbb{1}(c \ge c^1)P_a(c).$$

Proof. We already proved points 1, 2 and 3. We only need to compute the expected utility. Note that by symmetry, we have $p_1^N = p_2^N$ and $p_1^S = p_2^S$ so that the vector χ is

Table 2: Expected payoff matrix for agent 1

	S_2	N_2
S_1	$u(w-c)$	$(1-pq)u(w-c)+pqu(w-c-\ell)$
N_1	$(1-p)u(w)+pu(w-\ell)$	$pu(w-\ell)+(1-p)(pqu(w-\ell)+(1-pq)u(w))$

either $(0,0)$ or $(\mathrm{Ber}_1(p), \mathrm{Ber}_2(p))$ where $\mathrm{Ber}_1(p)$ and $\mathrm{Ber}_2(p)$ are Bernoulli random variables with parameter p. The first case corresponds to (S_1,S_2) and the second one to (N_1,N_2). In case $c<c^0$, the solution to our model is $X_1=X_2=0$, there is no loss and hence the expected utility is just $u(w-c)$. In the case $c>c^1$, (S_1,S_2) is not valid and the solution to our model is (N_1,N_2) and we have

$$\begin{aligned}\mathbb{P}((X_1,X_2)=(1,1)) &= q(1-(1-p)^2)+(1-q)p^2\\ \mathbb{P}((X_1,X_2)=(0,1)) &= p(1-p)(1-q),\\ \mathbb{P}((X_1,X_2)=(0,0)) &= (1-p)^2.\end{aligned}$$

Hence we have $p^N=\mathbb{P}(X_1=1)=p(1+q-pq)$ and the expected utility is $(1-p^N)u(w)+p^Nu(w-\ell)$ which corresponds to point 2.

The last points follow from the fact that by symmetry we have only to compare states (S_1,S_2) and (N_1,N_2). □

Let simplify our model by dropping the risk adverse assumption as in [11]. Hence, we have $\pi\equiv 0$ and $c^0 = p(1-pq)\ell < c^1 = p\ell$. Figure 1 shows the corresponding price of anarchy as a function of the cost for self-protection c.

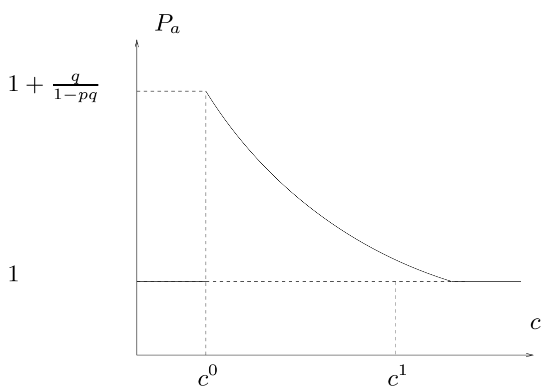

Figure 1: Price of anarchy: 2-agents case

Note that the amplitude of the price of anarchy is relatively small, of order $1+q$. It will be different in the network setting. Also note that for $c^0\le c\le c^1$, the price of stability $P_s(c)$ equals one and is strictly smaller than the price of anarchy $P_a(c)$. It corresponds to the case where both Nash equilibra are possible and only one of them is socially optimal, namely the situation where both agents choose to invest in self-proctection. When there is only one Nash equilibrium, there are some situations where the investment choices are efficient so that the decision by each of the individual agents are socially optimal. When the costs for self-protection are sufficiently low $c<c^0$, each agent wants to invest in self-protection. Also if the costs are very high $c>p\ell+pq(1-p)\ell$, then it is efficient for no one to incur them. However for $c^1<c<p\ell+pq(1-p)\ell$, the price of stability is strictly larger than one. In this case there is only one Nash equilibrium (corresponding to the case where both agents do not invest in self-protection) and this equilibirum is not socially otpimal. The costs are high enough that each agent does not want to invest in self-protection, but it would be better for society if all of them did so.

4. EPIDEMIC RISKS ON A RANDOM NETWORK

In this section we analyze our model on a large sparse random graph: $G^{(n)}=G(n,\lambda/n)$ on n nodes $\{0,1,\dots,n-1\}$, where each potential edge (i,j), $0\le i<j\le n-1$ is present in the graph with probability λ/n, independently for all $n(n-1)/2$ edges. Here $\lambda>0$ is a fixed constant independent of n. It corresponds to the case of the Erdös-Rényi graph which has received considerable attention in the past [10]. The results of this section are not restricted to this class of graphs and the analysis can be carried out for graphs with asymptotic given degree or uniform regular graphs. The main features of the solution are still valid in these different cases. We will discuss different extensions of our model in Section 6.1.

We denote by $\mathcal{G}^{(n)}$ the set of graphs on n nodes $\{0,1,\dots,n-1\}$. Our basic workspace is $\Omega\times\mathcal{G}^{(n)}$ and we denote by $X_i^{(n)}$ and $\chi_i^{(n)}$ the random variables satisfying recursions (1, 5, 6, 7) on the random graph $G^{(n)}\in\mathcal{G}^{(n)}$. First note that by exchangeability, we have for any fixed n and $i\neq j$:

$$p^N(n) := \mathbb{E}[X_i^{(n)}|D_i=1]=\mathbb{E}[X_j^{(n)}|D_j=1], \text{ and} \quad (8)$$

$$p^S(n) := \mathbb{E}[X_i^{(n)}|D_i=0]=\mathbb{E}[X_j^{(n)}|D_j=0]. \quad (9)$$

We will compute the possible limits (p^N,p^S) for the sequence $(p^N(n),p^S(n))$. In Section 4.1, we define a stochastic (tree-indexed) process $Y=(Y_i)$ which allows to construct the limiting object (as n tends to infinity) of the process $(X_i^{(n)}, i\in G^{(n)})$. We explain in Section 4.2 the required notion of convergence, namely the local weak convergence introduced by Aldous and Steele [2]. We show that it allows to compute p^N and p^S rigorously.

4.1 Exact results for trees

In this section, we suppose that $G=T$ is a tree with nodes $\emptyset,1,\dots,n$, with a fixed root $\emptyset$. For a node i, we denote by $\mathrm{gen}(i)\in\mathbb{N}$ the generation of i, i.e. the length of the minimal path from $\emptyset$ to i. Also we denote $i\to j$ if i is a children of j, i.e. $\mathrm{gen}(i)=\mathrm{gen}(j)+1$ and j is on the minimal path from $\emptyset$ to i.

Also, we fix p^N and p^S so that the random variables χ_i defined in (5) are i.i.d. Bernoulli random variables with parameter $\kappa(\gamma)=p^-\gamma+p^+(1-\gamma)$ where

$$\gamma=\mathbb{P}((p^N-p^S)\ell+\pi(p^N)-\pi(p^S)\ge c_1).$$

For an edge $(i,j)\in E$ with $i\to j$, we denote by $T_{i\to j}$ the sub-tree of T with root i when deleting edge (i,j) from T. Here the formal probability space introduced in Section 2.3 will be useful: we have a family of trees $T_{i\to j}$ and we run the epidemic model according to equation (1) with the same configuration $\omega = (B_{ij}, B_i^S, B_i^N, c_i, \ell_i, i, j \in \mathbb{N})$ on each tree. For a given configuration ω, we say that node i is infected from $T_{i\to j}$ if the node i is infected in $T_{i\to j}$ on the same configuration ω. We denote by Y_i the corresponding indicator function with value 1 if i is infected from $T_{i\to j}$ and 0 otherwise. A simple induction shows that the recursion (1) becomes:

$$1 - Y_i = (1-\chi_i) \prod_{k\to i} (1 - B_{ki} Y_k). \quad (10)$$

Note that we can compute all the Y_i recursively starting from the leaves with $Y_\ell = \chi_\ell$ for any leaf ℓ. As a consequence (and it is the main difference with (1) that makes the model on a tree tractable), the random variables Y_k with $k \to i$ in the right-hand term of (10) are independent of each others and independent of the B_{ki}. For any node $i \in T$, we just defined Y_i and the family $(Y_i, i \in T)$ is a tree-indexed process.

We now consider trees of specific types. Given a constant $\lambda > 0$, a depth-d Poisson tree $T(\lambda, d)$ with parameter λ and depth d is constructed as follows: the root node has a degree which is a random variable distributed according to a Poisson distribution with parameter λ. All the children of the root have out-degrees which are also random, distributed according to a Poisson λ distribution. We continue this process until either the process stops at some depth $d' < d$, where no nodes in level d' has any children or until we reach level d. In this case, all the children of nodes in level d are deleted and the nodes in level d become leaves. For $d = \infty$, we just denote by $T(\lambda)$ the Poisson tree.

Then Equation (10) defines a tree-indexed process $(Y_i, i \in T(\lambda, d))$, called a Recursive Tree Process (RTP) in [1]. We also introduce the associated Recursive Distributional Equation (RDE):

$$Y \stackrel{d}{=} 1 - (1-\chi) \prod_{k=1}^{N} (1 - B_k Y_k), \quad (11)$$

where N has a Poisson λ distribution, χ has a Bernoulli distribution with parameter $\kappa(\gamma)$, B_k are i.i.d. Bernoulli random variables with parameter q, Y and Y_k are i.i.d. copies and all random variables are independent of each others. RDE for RTP plays a similar role as the equation $\mu = \mu K$ for the stationary distribution of a Markov chain with kernel K, see [1].

PROPOSITION 2. *For $\gamma < 1$ or $p^- > 0$, the RDE (11) has a unique solution: $\mathbb{P}(Y = 1) = 1 - \mathbb{P}(Y = 0) = h$, with $h = h(p^+, p^-, \gamma, q, \lambda)$ the unique solution in $[0,1]$ of*

$$h = 1 - (1-\kappa(\gamma)) e^{-\lambda q h},$$

where $\kappa(\gamma) = p^-\gamma + p^+(1-\gamma)$.

PROOF. Simple calculations give

$$\begin{aligned}
1-h &= \mathbb{P}(Y=0) \\
&= \mathbb{P}((1-\chi)\prod_{k=1}^{N}(1-B_kY_k)=1) \\
&= \mathbb{P}(\chi=0)\left(\sum_n \mathbb{P}(N=n)\mathbb{P}(B_1Y_1=0)^n\right) \\
&= (1-\kappa(\gamma))\sum_{n=0}^{\infty} e^{-\lambda}\frac{\lambda^n}{n!}(1-qh)^n \\
&= (1-\kappa(\gamma))e^{-\lambda q h}.
\end{aligned}$$

It follows from Lemma 1 (in Appendix) that this fixed-point equation has a unique solution in $[0,1]$. Similar calculation for $h' = \mathbb{P}(Y=1)$ shows that $h' = (1-\kappa(\gamma))e^{-\lambda q(1-h')}$, so that $h + h' = 1$ and the proposition follows. □

As a consequence, we see that it is possible to construct an invariant version of the RTP on the tree $T(\lambda)$ (which is possibly infinite for $\lambda > 1$) where for each $k \ge 0$, the sequence $(Y_i, i \in T(\lambda), \text{gen}(i) = k)$ is a sequence of i.i.d. Bernoulli random variables with parameter h, see [1].

COROLLARY 1. *For the invariant RTP $(Y_i, i \in T(\lambda))$, we have*

$$\begin{aligned}
\mathbb{E}[Y_\varnothing | D_\varnothing = 1] &= 1 - (1-p^-)e^{-\lambda q h}, \\
\mathbb{E}[Y_\varnothing | D_\varnothing = 0] &= 1 - (1-p^+)e^{-\lambda q h},
\end{aligned}$$

where h is defined in Proposition 2.

PROOF. The random variables $(Y_i, i \in T(\lambda), \text{gen}(i) = 1)$ are independent Bernoulli random variables with parameter h. Hence the same calculation as in the proof of Proposition 2 gives the desired result. □

We introduce the fixed point equation:

$$\begin{aligned}
p^N &= 1 - (1-p^+)e^{-\lambda q h(\gamma)}, & (12)\\
p^S &= 1 - (1-p^-)e^{-\lambda q h(\gamma)}, & (13)\\
\gamma &= \mathbb{P}((p^N - p^S)\ell_1 + \pi(p^N) - \pi(p^S) \ge c_1). & (14)
\end{aligned}$$

We end this section with the following easy property:

PROPOSITION 3. *For any solution (p^N, p^S, γ) of the fixed point equation (12,13,14), the corresponding invariant RTP is a solution of our model on $T(\lambda)$.*

We will study the fixed point equation in Section 5. Before that, we show that this fixed point equation characterize the limiting behaviors of the process on the finite graph $G^{(n)}$ as $n \to \infty$.

4.2 Asymptotics for Erdos-Renyi graphs

In this section, we consider the process $(X_i^{(n)}, i \in \{0, \dots, n-1\})$ satisfying (1) on $G^{(n)}$. To simplify the analysis, we assume that for any (p^N, p^S, γ) solution of the fixed point equation (12,13,14), there exists $\epsilon > 0$, such that for any $(x, y) \in B((p^N, p^S), \epsilon)$ (where $B((p^N, p^S), \epsilon) \subset \mathbb{R}^2$ is the disk of center (p^N, p^S) and radius ϵ),

$$\begin{aligned}
&\mathbb{P}((x-y)\ell_1 + \pi(x) - \pi(y) \ge c_1) \quad (15)\\
&= \mathbb{P}((p^N - p^S)\ell_1 + \pi(p^N) - \pi(p^S) \ge c_1).
\end{aligned}$$

This assumption on the distribution of the c_i's and ℓ_i's will be satisfied in the examples treated in Section 5. The following proposition is the main technical result of this paper and explain the introduction of the RTP in previous section. Its proof relies on the Objective Method [2] and will not be needed for later analysis.

PROPOSITION 4. *For any solution (p^N, p^S, γ) of the fixed point equation (12,13,14), there exists a process $(X_i^{(n)}, i \in \{0, \dots, n-1\})$ on $G^{(n)}$ satisfying the equations of our model such that*

$$p^N = \lim_{n\to\infty} \mathbb{E}[X_0^{(n)}|D_0 = 0],$$
$$p^S = \lim_{n\to\infty} \mathbb{E}[X_0^{(n)}|D_0 = 1].$$

PROOF. The proof will be as follow: fix (p^N, p^S, γ) solution of the fixed point equation. Consider the process $(\tilde{X}_i^{(n)}, i \in \{0, \dots, n-1\})$ on $G^{(n)}$ satisfying the equations of our model but with $\chi_i^{(n)}$ replaced by χ_i which are i.i.d Bernoulli with parameter $\kappa(\gamma)$, where γ is given by (14). We will show that

$$\lim_{n\to\infty} \mathbb{E}[\tilde{X}_0^{(n)}|D_0 = 0] = p^N, \quad (16)$$
$$\lim_{n\to\infty} \mathbb{E}[\tilde{X}_0^{(n)}|D_0 = 1] = p^S. \quad (17)$$

It is then easy to see that thanks to (15), for n sufficiently large the process $(\tilde{X}_i^{(n)}, i \in \{0, \dots, n-1\})$ satisfy the equations of our model.

Let us fix a large positive integer d and let $\mathcal{N}_d(0, G^{(n)})$ denotes the neighborhood of radius d about the root 0: $\mathcal{N}_d(0, G^{(n)})$ is the collection of nodes form $G^{(n)}$ which are connected to 0 by paths with length $\le d$. For fixed d, we have as n tends to infinity $\mathcal{N}_d(0, G^{(n)}) \xrightarrow{d} T(\lambda, d)$, where $T(\lambda, d)$ is a depth-d Poisson tree [19].

By the Skorohod representation theorem, we may suppose the random variables $\mathcal{N}_d(0, G^{(n)}), T(\lambda, d)$ are all defined on a common probability space and that with probability one, there is a finite random variable N such that $\mathcal{N}_d(0, G^{(n)}) = T(\lambda, d)$ for all $n \ge N$.

We now define two depth-d RTPs as follows: let $\partial T(\lambda, d)$ denote the leaves of $T(\lambda, d)$. Define $L_i^{(d)} = \chi_i$ and $U_i^{(d)} = 1$ for $i \in \partial T(\lambda, d)$. Then we use (10) recursively to define $(L_i^{(d)}, i \in T(\lambda, d))$ and $(U_i^{(d)}, i \in T(\lambda, d))$. Since for all $i \in \partial T(\lambda, d)$, we have $L_i^{(d)} \le \tilde{X}_i^{(n)} \le U_i^{(d)}$, the monotonicity of (10) implies that for $n \ge N$,

$$L_\varnothing^{(d)} \le \tilde{X}_0^{(n)} \le U_\varnothing^{(d)}.$$

The $L_i^{(d)}$ for gen$(i) = 1$ are independent and by Lemma 2 (proved in Appendix 8.2) converge in distribution to Bernoulli random variables with parameter $h(\gamma)$. As a consequence we have for $L_\varnothing^{(d)}$:

$$\begin{aligned} &\lim_{d\to\infty} \mathbb{E}[L_\varnothing^{(d)}|D_\varnothing = 1] \\ =\ & 1 - \lim_{d\to\infty} \mathbb{P}\left((1 - B_\varnothing^S) \prod_{i\sim\varnothing} (1 - B_{i\varnothing} Y_i^{(d)}) = 1\right) \\ =\ & 1 - (1 - p^-) e^{-\lambda q h(\gamma)} = p^S. \end{aligned}$$

The same argument holds for the process U, hence there exists a function $\epsilon(d)$ with $\epsilon(d) \to 0$ as $d \to \infty$ and such that

$$p^S - \epsilon(d) \le \mathbb{E}[\tilde{X}_0^{(n)}|D_0 = 1] \le p^S + \epsilon(d).$$

We proved (16) and the proof of (17) is similar. □

5. PRICE OF ANARCHY: THE NETWORK CASE

In this section, we discuss the intuition behind, and the impact of the mathematical results obtained in the previous section. We compare the results for a network of n interconnected agents with the case of the 2-agent network described in Section 3.3. We also show the existence of a cascading phenomenon, and derive practical insight for malware protection in the Internet.

5.1 Interpretation of the mathematical results

We first discuss the intuition behind the various quantities computed in Section 4.

Consider a solution (p^N, p^S, γ) of the fixed point equation (12,13,14). By Proposition 3, this solution corresponds to an equilibrium of the game played on the random tree $T(\lambda)$. The aim of the results shown in Section 4.2 is to prove that, to any such equilibrium, one can associate an equilibrium for the game played on the (finite) Erdös-Réyni graph $G^{(n)}$ such that asymptotically in n, the characteristics of this equilibrium are given by that of the equilibrium on $T(\lambda)$ (see Proposition 4). In particular, the quantity γ corresponds exactly to the limit, as the population size n tends to infinity, of the fraction of the population investing in self-protection for the equilibrium of the game played on the Erdös-Rényi graph $G^{(n)}$. Then we have the following interpretations:

- γ is the (asymptotic) fraction of the population investing in self-protection;
- p^N is the (asymptotic) probability of loss for an agent not investing in self-protection;
- p^S is the (asymptotic) probability of loss for an agent investing in self-protection.

Hence the average (over the population) probability of loss is given by

$$\begin{aligned} \lim_{n\to\infty} \mathbb{E}[X_0^{(n)}] &= \gamma p^S + (1-\gamma) p^N \\ &= h(\gamma) \end{aligned}$$

given by Proposition 2. The fact that h corresponds exactly to the probability of loss is a special feature of Erdös-Rényi graphs; for other graphs, the probability of loss is given by $\gamma p^S + (1-\gamma)p^N$. Now the average incurred cost to the agents in the equilibrium associated to (p^N, p^S, γ) is given by $\gamma\mathbb{E}[c] + h(\gamma)\mathbb{E}[\ell]$. The first term corresponds to the cost incurred by the investment in self-protection by a fraction γ of the population and the second term corresponds to the expected cost of losses.

If there are multiple solutions to the fixed point equation (12,13,14), these solutions correspond to different equilibria of the game on $G^{(n)}$. Then the quantities computed above allow us to compare these equilibria. We consider this in the sections below.

5.2 Case of Erdos-Renyi graphs

In this section, we consider the case where $G^{(n)} = G(n, \lambda/n)$. Note that our results concerning the network case are asymptotic results as $n \to \infty$. Thus, when we write "for a certain range of parameters, all agents invest in self-protection", we mean that the fraction of agents investing in self-protection in our model on a graph with n agents tends to 1 as n tends to infinity. Similarly, the price of anarchy is the limit of the price of anarchy for the model with n agents when $n \to \infty$.

For simplicity, we drop the risk adverse condition so that we have $\pi \equiv 0$ and we assume that the cost for self-protection $c_i = c$ and the possible loss $\ell_i = \ell$ are non random parameters. Then Equation (14) becomes

$$\gamma = \mathbb{1}\left((p^N - p^S)\ell \geq c\right).$$

We define

$$c^0 := \frac{1-h(0)}{1-p^+}(p^+ - p^-)\ell,$$
$$c^1 := \frac{1-h(1)}{1-p^-}(p^+ - p^-)\ell,$$

where $h(\gamma)$ solves the fixed point equation of Proposition 2.

Proposition 5. *We have $c^0 \leq c^1$. Furthermore*

1. *If $c < c^0$ then every agent in the network invests in self-protection.*
2. *If $c > c^1$ then no agent invests in self-protection.*
3. *If $c^0 \leq c \leq c^1$ then both equilibria are possible.*
4. *The price of anarchy is given by*
$$P_a(c) = 1 \vee \mathbb{1}(c^0 < c)\frac{h(0)\ell}{c + h(1)\ell}.$$
5. *The price of stability is given by*
$$P_s(c) = 1 \vee \mathbb{1}(c^1 < c)\frac{h(0)\ell}{c + h(1)\ell}.$$

Proof. Recall that p^N and p^S are functions of γ, and we will therefore use the notation $p^{N,\gamma}$ and $p^{S,\gamma}$. We have by Proposition 4 and Equations (12) and (13),

$$c^\gamma := \left(p^{N,\gamma} - p^{S,\gamma}\right)\ell = \ell(p^+ - p^-)e^{-\lambda q h(\gamma)},$$

so that by Lemma 1, the map $\gamma \mapsto c^\gamma$ is non-decreasing and $c^0 \leq c^1$. Also by symmetry, there are only two possible equilibria: $S := (S_1, S_2, \dots)$ or $N := (N_1, N_2, \dots)$. If $c \leq c^0$, then we have $\mathbb{1}(c^\gamma \geq c) = 1$ and $\gamma = 1$ is the only solution, which corresponds to the state S. Moreover in state S (and with $\gamma = 1$), the probability of loss is given by

$$p^{S,1} = 1 - (1-p^-)e^{-\lambda q h(1)} = h(1).$$

Similarly, if $c \geq c^1$, the only solution is $\gamma = 0$ which corresponds to state N. In state N (and with $\gamma = 0$), the probability of loss is given by

$$p^{N,0} = 1 - (1-p^+)e^{-\lambda q h(0)} = h(0).$$

If $c \in (c^0, c^1)$ then both $\gamma = 1$ and $\gamma = 0$ are solutions. For a given $\gamma \in [0,1]$, the average probability of loss is given by

$$\mathbb{E}[Y_\emptyset] = \gamma p^{S,\gamma} + (1-\gamma)p^{N,\gamma} = h(\gamma).$$

Hence the average cost incurred to agents is given by $\xi^\gamma := \gamma c + h(\gamma)\ell$ which, by Lemma 1, is a concave function of γ. In particular, we have $\inf_\gamma \xi^\gamma = \min(\xi^0, \xi^1)$. Note that $\xi^1 \leq \xi^\gamma$ iff $c(1-\gamma) \leq (h(\gamma) - h(0))\ell$ so that by concavity of h, we have $\xi^1 \leq \xi^\gamma$ iff $c \leq (h(0) - h(1))\ell$. Simple calculations show that

$$c^0 < c^1 < (h(0) - h(1))\ell.$$

Points 4 and 5 follow easily. □

This proposition is similar to (and extends) Proposition 1 for the 2-agent case, and a curve similar to Figure 1 could be plotted in the network (n-agent) case. In particular, our model allows to compute the range of the parameters for which the Nash equilibrium is, or is not, socially optimal. However, there are some differences that we highlight now. The main difference is in the amplitude of the price of anarchy which is much higher here. Note that in Section 3.3, we took $p^- = 0$ and $p^+ = p$. In that case, we have $h(1) = 0$ and

$$h(0) = 1 - (1-p)e^{-\lambda q h(0)}.$$

Corollary 2. *As $p \to 0$, we have*

$$P_a(c) \sim \frac{x^*\ell}{c} \vee 1,$$

where x^* is the unique solution in $(0,1]$ of the fixed point equation

$$1 - x^* = e^{-\lambda q x^*}. \quad (18)$$

Proof. This corollary follows directly from Proposition 5 and the observation that as $p \to 0$, the fixed point equation satisfied by $h(0)$ tends to (18). □

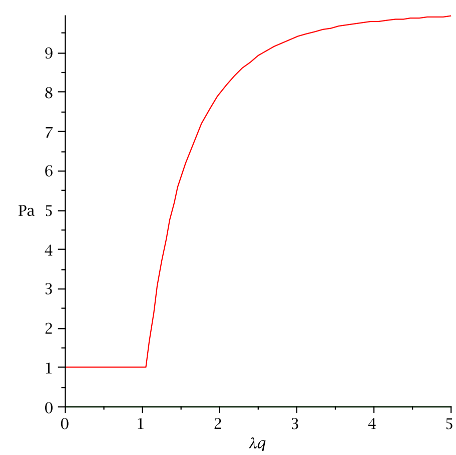

Figure 2: Price of anarchy as a function of λq with $\ell/c = 10$ and $p \to 0$.

For a fixed cost c and a fixed amount of loss ℓ, Figure 2 shows the variation of the price of anarchy with λq in the limit $p \to 0$. We see that as $\lambda q \to \infty$, the price of anarchy tends to ℓ/c which is much larger than one.

There is a qualitative interpretation of how x^* arises in our model:

- First, note that x^* has a natural interpretation in term of branching processes [4]. Consider a Poisson tree $T(\lambda q)$. Then if $\lambda q \leq 1$, this tree is finite a.s. whereas for $\lambda q > 1$, there is a positive probability that the tree is infinite. In any case, the probability that the tree is infinite is equal to x^*. In particular, $x^* = 0$ for $\lambda q \leq 1$.

- Consider now our model when every agent chooses to not invest in self-protection with a small value of $p > 0$. Then a fraction p of the population experiences direct loss. From these individuals, the epidemic propagates like a Poisson branching process with parameter λq [19].

  If $\lambda q < 1$, this branching process eventually stops. Hence, in this case, the infected population consists of disjoint clusters with one initially infected agent per cluster. As a result as $p \to 0$, we decrease the number of initially infected agents, i.e. the number of final clusters and the epidemics has asymptotically no impact.

  If $\lambda q > 1$, there is a fraction x^* of the nodes for which the branching process 'lasts forever' and those nodes belong to a single 'giant component'. For any positive value of p, an agent of this component will experience a direct loss and then contaminate the whole component. Hence we see that, as $p \to 0$, the total cost of the epidemic is of the order of $x^* n\ell$.

- Going back to the price of anarchy, consider the state where all agents are in state S (in which case there is no epidemic and the cost per agent is just c) and the state where all agents are in state N (and the average cost of the epidemic is $x^*\ell$). This gives us again Corollary 2.

In other words, we have the following situation: nobody invests in self-protection because the advantage for one agent to invest is negligible; however if all agents had chosen to invest then they would be much better off. The curve of Figure 2 quantifies exactly by how much. Note that in this case, the epidemic starts with a negligible fraction of the population and eventually reaches a significant proportion x^* of the population which fits well the observed propagation of worms [23].

5.3 Tipping phenomenon

We have shown in Section 5.2 that our model exhibits externalities and that the equilibrium is not socially optimal: there is a market failure. In order to resolve this market failure, one can actually take advantage of a tipping (or cascading) phenomenon.

We consider the framework of Section 5.2 and suppose that the price for self-protection is such that $c \in (c^0, c^1)$ so that both equilibria where everybody or nobody invests in self-protection are possible. Consider the case when the population is 'trapped' in the 'bad' state. If one 'forces' a fraction of the population to invest in self-protection, that initial core set of investors in self-protection can trigger a cascading phenomenon of adoption (of self-protection) which leads to the state where the entire population is in the 'good' state.

In other words, a possible strategy to induce agents to invest in self-protection would for example to give anti-virus software for free to a fraction of the population - this will eventually lower the probability of loss for all agents in the network. Recall (see Equation (3)) that an agent decides to invest in self protection if $c < (p^N - p^S)\ell$. In particular when γ increases, both p^N and p^S decrease but $p^N - p^S$ increases (according to (12) and (13)). Hence it is possible to increase γ in order to give an incentive to others to invest in self-protection for a fixed cost c.

Now the crucial question is to determine the minimal proportion of the population that would trigger such a cascade of adoption. The answer is given by the curve on Figure 3. For example, if $c/\ell \sim 0.14$ then if a fraction $\gamma \sim 0.2$ of the population is protected, a cascading phenomenon will automatically induce the rest of the population to invest in self-protection.

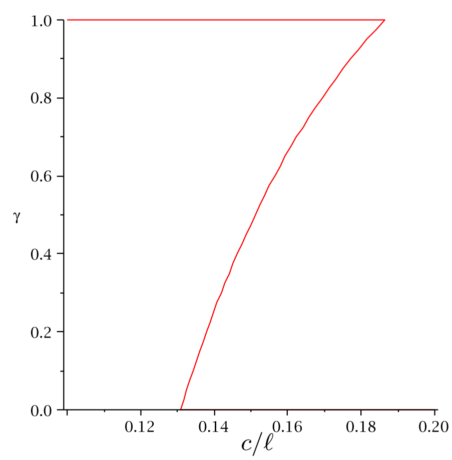

Figure 3: Tipping phenomenon: $p^- = 0$**,** $p^+ = 0.1$**,** $q = 0.5$ **and** $\lambda = 10$**.**

More formally, assume we are in the framework of Section 5.2, namely $\pi \equiv 0$ and c and ℓ are deterministic. Then for a ratio c/ℓ fixed the corresponding minimal value of γ which will trigger a cascading effect is given by

$$\arg\min_{\gamma}\{c/\ell < p^{N,\gamma} - p^{S,\gamma}\}$$

Then the results of Section 4 give the curve of Figure 3.

A similar cascading effect, although appearing in a different context, is modeled and analyzed in [3] for a different diffusion process where social interactions are local only. Also in [5] and [13], we study the question of whether insurance could be a mechanism to induce agents to invest in self-protection.

6. DISCUSSION

6.1 Extending the results to the general case of sparse graphs: the Local Mean Field model

The method developed in Section 4 relies heavily on the local weak convergence introduced by Aldous and Steele and surveyed in [2]. This method is not restricted to Erdös-Rényi graphs and is successful whenever the underlying graph is 'locally' a tree. Reference [12] presents a general framework extending the model of the present paper and introduces the

Local Mean Field (LMF) model associated to our model for a general sparse random graph.

We illustrate the approach with an example related to the propagation of worms, as considered in this paper. A specific type of worm is a scan-based worm which (in first approximation) probes the entire IP address space. We can model the propagation of a scan-based worm as follows: each newly infected agent runs the worm program and tries to compromise a random number (say I with distribution $F(k) = \mathbb{P}(I \leq k)$) of other agents taken at random among all the other agents, this with a probability of success q. In this model, the limiting object is a RTP as in Section 4.1, where the underlying tree is a Galton-Watson branching process in which the root has offspring distribution F and all other genitors have offspring distribution G where for all $k \geq 1$, $G(k-1) = \frac{kF(k)}{\sum_k kF(k)}$ (see Chapter 3 in [6]). The recursion equation (10) is still valid, and one then has to study the recursive distributional equation (11), however with N distributed according to G. Once the solution of the RDE is derived, it is easy to obtain the limits appearing in Corollary 1 (by taking care of the distinct distribution of the root). We then obtain the corresponding fixed point equation for p^N, p^S and γ which allows to compute the price of anarchy and other quantities as in previous sections. We refer to [12] for a careful treatment of this case and for extensions of the model presented here.

6.2 Summary and Conclusion

We studied a network of interconnected agents subject to epidemic risks and which can decide whether or not to invest some amount to self-protect and deploy security solutions. We introduced a general model which combines a model for the propagation of epidemics among networked agents and an economic model for agents. To the best of our knowledge, our model is the first to include an arbitrary large and structured network of agents, and to capture the situation when an agent's payoff depends on decisions made by (the entire population of) other agents. We are able to solve analytically our model which captures network externalities arising in the economic problem of security systems. In particular, we show that there is a possible market failure and we compute exactly the price of anarchy of this problem. Furthermore, we give insights on the structure of the equilibria and show that cascading phenomena are possible and could be the basis for developing strategies to increase the adoptability of security features and protocols in networked environments in general, and in the Internet in particular.

7. REFERENCES

[1] D. Aldous and A. Bandyopadhyay. A survey of max-type recursive distributional equations. *Annals of Applied Probability*, 15(2):1047-1110, 2005.

[2] D. Aldous and J.M. Steeele. The objective method: probabilistic combinatorial optimization and local weak convergence. *Probability on discrete structures*, Springer, vol. 110, pp. 1-72, 2004.

[3] H. Amini and M. Lelarge. Maximizing the Impact of Viral Marketing in a Random Network. *preprint*, 2008.

[4] K.B. Athreya and P.E. Ney. Branching processes. *Dover*, 2000.

[5] J. Bolot and M. Lelarge. A New Perspective on Internet Security using Insurance. *INFOCOM 08, Mini-Conference*, see also INRIA report 6329.

[6] R. Durrett. Random graph Dynamics. *Cambridge University Press*, 2006.

[7] A. Ganesh, L. Massoulie, D. Towsley. The effect of network topology on the spread of epidemics. *Proc. IEEE Infocom 2005*, Miami, FL, March 2005.

[8] R. Gibbons. Game theory for applied economists. *Princeton university Press*, 1992.

[9] C. Gollier. The Economics of Risk and Time. *MIT Press*, 2004.

[10] S. Janson, T. Luczak and A. Rucinski. Random graphs. *Wiley-Interscience, New York*, 2000.

[11] H. Kunreuther and G. Heal. Interdependent security: the case of identical agents. *Journal of Risk and Uncertainty*, 26(2):231–249, 2003.

[12] M. Lelarge and J. Bolot. A Local Mean Field Analysis of Security Investments in Networks. *arXiv:0803.3455 [cs.GT]*, 2008.

[13] M. Lelarge and J. Bolot. Using Insurance to Increase Security in the Internet. *preprint*, 2008.

[14] J. Miekisz. Stochastic Stability in Spatial Games. *Journal of Statistical Physics*, 117(1/2):99-110, 2004.

[15] D. Moore, V. Paxson, S. Savage, C. Shannon and N. Weaver. Inside the Slammer worm. *IEEE Security and Privacy*, 1(4):33–39, 2003.

[16] N. Nisan, T. Roughgarden, E. Tardos and V.V. Vazirani (eds). Algorithmic game theory. *Cambridge University Press*, 2007.

[17] H. Orman. The Morris worm: a fifteen-year perspective. *IEEE Security and Privacy Magazine*, Sept-Oct 2003.

[18] S. Saniford et al. The top speed of flash worms. *Proc. ACM Workshop on Rapid Malcode WORM04*, Fairfax, VA, Oct. 2004.

[19] J. Spencer. Ten lectures on the probabilistic method *SIAM*, vol. 64, 1994

[20] M. Vojnovic and A. Ganesh. On the race of worms, alerts and patches. *Proc. ACM Workshop on Rapid Malcode WORM05*, Fairfax, VA, Nov. 2005.

[21] N. Weaver, V. Paxson, S. Staniford, R. Cunningham. A taxonomy of computer worms. *Proc. First ACM Workshop on Rapid Malcode (WORM 2003)*, Washington DC, Oct. 2003.

[22] V. Yegneswaran, P. Barford, J. Ullrich. Internet intrusions: global characteristics and prevalence. *Proc. ACM Sigmetrics*, June 2003.

[23] C. Zou, W. Gong, D. Towsley. Code Red worm propagation modeling and analysis. *Proc. 9th ACM Conf. Computer Comm. Security CCS'02.*, Washington, DC, Nov 2002.

8. APPENDIX

8.1 Expected utility model and risk aversion

The classical expected utility model is named thus because it considers agents that attempt to maximize some kind of expected utility function u. In this paper, we assume that agents are rational and that they are risk averse, i.e. their utility function is concave (see Proposition 2.1 in [9]). Risk averse agents dislike mean-preserving spreads in the distribution of their final wealth.

We denote by w the initial wealth of the agent. The *risk* *premium* π is the maximum amount of money that one is ready to pay to escape a pure risk X, where a pure risk X is a random variable such that $\mathbb{E}[X] = 0$. The risk premium corresponds to an amount of money paid (thereby decreasing the wealth of the agent from w_0 to $w_0 - \pi$) which covers the risk; hence, π is given by the following equation:

$$u(w - \pi) = \mathbb{E}[u(w_0 + X)]$$

Each agent faces a potential loss ℓ, which we take in this paper to be a fixed (non-random) value. We denote by p the probability of loss or damage. There are two possible final states for the agent: a good state, in which the final wealth of the agent is equal to its initial wealth w_0, and a bad state in which the final wealth is $w - \ell$. If the probability of loss is $p > 0$, the risk is clearly not a pure risk. The amount of money m the agent is ready to invest to escape the risk is given by the equation:

$$pu(w - \ell) + (1 - p)u(w) = u(w - m) \qquad (19)$$

We clearly have $m > p\ell$ thanks to the concavity of u. We can actually relate m to the risk premium defined above:

$$m = p\ell + \pi[p].$$

8.2 Technical lemmas

For $0 \le p^- < p^+ < 1$ and $\lambda q > 0$, we define the function

$$f(x, \gamma) = 1 - (1 - p^-\gamma - p^+(1 - \gamma))e^{-\lambda q x}.$$

LEMMA 1. *For $\gamma < 1$ or $p^- > 0$, the fixed point equation*

$$h = f(h, \gamma)$$

has a unique solution in $[0, 1]$ denoted $h(\gamma)$. Moreover the function $\gamma \mapsto h(\gamma)$ is non-increasing and concave.

PROOF. The function $x \mapsto f(x, \gamma)$ is continuous, non-increasing and concave. Note that $f(0, \gamma) = p^-\gamma + p^+(1 - \gamma) > 0$ and $f(1, \gamma) = 1-(1-p^-\gamma-p^+(1-\gamma))e^{-\lambda q} \le 1$, hence first point of the lemma follows. The monotonicity of h in γ follows from the fact that the function f is non-increasing in γ so that for $\gamma_1 < \gamma_2$, we have

$$f(h(\gamma_1), \gamma_2) < f(h(\gamma_1), \gamma_1) = h(\gamma_1),$$

and iterating f, we get in the limit $h(\gamma_2) \le h(\gamma_1)$. A direct computation gives for the derivative of h:

$$h'(\gamma)\left(1 - \lambda q(1 - p^-\gamma - p^+(1 - \gamma))e^{-\lambda q h(\gamma)}\right) = (p^- - p^+)e^{-\lambda q h(\gamma)},$$

so that we have $1 - \lambda q(1 - p^-\gamma - p^+(1 - \gamma))e^{-\lambda q h(\gamma)} > 0$. Then for the second derivative of h, we get

$$\begin{aligned} & h''(\gamma)\left(1 - \lambda q(1 - p^-\gamma - p^+(1 - \gamma))e^{-\lambda q h(\gamma)}\right) \\ = \quad & -2\lambda q h'(\gamma)(p^- - p^+)e^{-\lambda q h(\gamma)} \\ & -(\lambda q h'(\gamma))^2(1 - p^-\gamma + p^+(1 - \gamma))e^{-\lambda q h(\gamma)}, \end{aligned}$$

hence $h''(\gamma) \le 0$ and concavity follows. □

We recall here the definition of the two depth-d RTPs in the proof of 4: let $\partial T(\lambda, d)$ denote the leaves of $T(\lambda, d)$. Define $L_i^{(d)} = \chi_i$ and $U_i^{(d)} = 1$ for $i \in \partial T(\lambda, d)$. Then we use (10) recursively to define $(L_i^{(d)}, i \in T(\lambda, d))$ and $(U_i^{(d)}, i \in T(\lambda, d))$.

LEMMA 2. *Assume the sequence χ_i is a sequence of independent Bernoulli random variables with parameter $\kappa(\gamma)$. Then both $L_\emptyset^{(d)}$ and $U_\emptyset^{(d)}$ converge in distribution as $d \to \infty$ to a Bernoulli random variable with parameter $h(\gamma)$ defined in Proposition 2.*

PROOF. This lemma follows from an easy induction on d. First consider $d = 1$, then we have

$$\begin{aligned} 1 - h^{(1)} &= \mathbb{P}(L_\emptyset^{(1)} = 0) \\ &= \mathbb{P}((1 - \chi_\emptyset)\prod_{k=1}^{N}(1 - B_k\chi_k) = 1) \\ &= (1 - \kappa(\gamma))\sum_{n=0}^{\infty} e^{-\lambda}\frac{\lambda^n}{n!}(1 - q\kappa(\gamma))^n \\ &= (1 - \kappa(\gamma))e^{-\lambda q \kappa(\gamma)}. \end{aligned}$$

Hence we have $h^{(1)} = f(\kappa(\gamma), \gamma)$. It is easy to see that for $d \ge 1$ if $h^{(d)} = 1 - \mathbb{P}(L_\emptyset^{(d)} = 0)$, then we have $h^{(d)} = f^d(\kappa(\gamma), \gamma)$, where the composition of the function f is in the first variable, γ being fixed.

Similarly if $g^{(d)} = 1-\mathbb{P}(U_\emptyset^{(d)} = 0)$, we have $g^{(d)} = f^d(1, \gamma)$. Then Lemma 1 implies that $h^{(d)}, g^{(d)} \to h(\gamma)$ and the lemma follows. □