MANAGEMENT SCIENCE
Vol. 35, No. 2, February 1989
Printed in U.S.A.

MODELS FOR MANAGING SECRETS*

J. MICHAEL STEELE

Program in Statistics and Operations Research, School of Engineering and Applied Science, Princeton University, Princeton, New Jersey 08544

Some elementary probability models are given that help articulate the trade-offs involved in managing secrets. One of these models captures the notion that the likelihood of disclosing a secret increases as the square of the number of people who are aware of it. Also, several classes of countermeasures are examined to determine their ability to increase the time window during which a secret can be held. One model suggests that disinformation countermeasures offer a promising technique for conserving secrets.

(CONFIDENTIALITY; DISINFORMATION COUNTERMEASURES; CLIQUE COMMUNICATION MODEL; SECRET DISCLOSURE)

1. Introduction

In Tom Clancy's *Hunt for Red October*, Vice Admiral James Greer says, "The likelihood of a secret's being blown is proportional to the *square* of the number of people who're in on it." The models given here are motivated in part by the desire to understand the analytic validity of Admiral Greer's theorem. More broadly, they aim to provide quantitative insight into the factors that influence the disclosure of secrets.

One of the models introduced here confirms Greer's square law, and, while the recapture of conventional wisdom provides some reassurance, the square law *per se* should not be over-valued. The real benefit of any honest secrecy model has to be the quality of guidance it provides about secret disclosure under conditions that can be influenced by management. Certainly, one quantity that should be under management control is the number of people who are initially informed of the secret, but even simple models should allow additional possibilities. Some practical examples include (1) the level of security of communication channels, (2) the use of protocols that govern communication between people who are in on the secret, and possibly (3) the protections provided by disinformation countermeasures.

Before giving specific models, some clarifications should be made of the words *person* and *secret*. There are many secrets of importance besides those that influence national security. More mundane secrets, like those that influence stock prices, hold sufficient value to society to justify careful modeling. When the notion of a *secret* is properly framed and generously interpreted, secrets can be seen as an inevitable part of life.

Who holds these secrets whose disclosure we want to model? Anthropomorphic language overemphasizes secrets held directly by individuals, although many significant secrets are held electronically. Still, managers with responsibility for secure information are well aware of this, and one does no harm by talking about *people* who hold secrets. Nevertheless, before endorsing or rejecting any model of secret disclosure, attention must be given to the substantial variety of systems that hold secrets.

The next two sections illustrate components that help build models that capture the analytical essence of the social and technological systems that hold secrets. Each of these sections provides an elementary probability model for secret disclosure that is analytically tractable yet intuitively reasonable.

* Accepted by Linda Green; received February 10, 1987. This paper has been with the author 4 months for 2 revisions.

0025-1909/89/3502/0240$01.25

The fourth, fifth, and sixth sections take on the task of modeling countermeasures to increase security. The first discussion of countermeasures focuses on the benefits of secure communication channels. One appealing aspect of the secure channel countermeasures is that it joins secrecy modeling with some topics of classical optimization.

§5 explores the possibility of increasing the window of secret security by a mechanism called *slices*. Sometimes a secret can be separated into parts in such a way that until the adversary learns all of the parts of the secret no critical disclosure takes place. A simple example is a 30-digit code that is divided into three ten-digit codes. Some loss is incurred when two of the three parts are disclosed, but, as a practical matter, the secret is secure until the third set of ten digits is revealed.

In §6, a disinformation countermeasure is examined to see how it can increase the window of secret conservation. Although the analysis must leave many issues unexplored, a clear sense emerges of substantial benefits that can be gained.

The final section engages the trade-offs inherent in secret keeping models and countermeasure techniques. Cautions are given there concerning the application of all our models, but one critical caution cannot be deferred. The focus in this exposition is on *components* of secret models, but secrets of real value require models that incorporate many components acting in concert. It would be inappropriate, or even disastrous, to rely on simple components without tempering one's judgement with all the additional wisdom that can be found.

2. The Simplest Model: Clique Communication

The first and simplest model studied here is the Clique Communication Model (or CCM). The context for the CCM (and the other models that follow) is a social environment where a secret is created and is eventually disclosed through some natural feature of the evolving environment.

The variable of fundamental interest is the length of time T between the creation of the secret and the moment it is disclosed. Under the CCM, the expected value of T will be shown to be

$$E(T) = 2/(n(n-1)\lambda p). \tag{2.1}$$

Here, n denotes the number of people ($n \geq 2$) who are initially aware of the secret, λ is a parameter that measures the rate of communication between any pair of people who are in on the secret, and p is the probability that any given communication is compromised.

Some feeling for the behavior of $E(T)$ in the CCM is gained by considering a typical situation where $\lambda = 2/\text{day}$ and $p = 0.01$. The brief table (Table 1) of $E(T)$ as a function of n shows the benefits of keeping n to a bare minimum.

Constructing Table 1 is straightforward using formula (2.1), yet one may still find it surprising that reasonably secure channels ($p = 0.01$) and modestly infrequent communication ($\lambda = 2/\text{day}$) lead to an expected disclosure time of just a bit over two days when seven people start out with knowledge of the secret.

Before investigating any credibility in the expression for $E(T)$, one has to examine the formal description of the CCM. The essence of the CCM is that it is a communication model. It rests on the premise that if a secret must be known to n people, one can

TABLE 1

Expected Numbers of Days Until Disclosure

Number of People in on the Secret	2	3	4	5	6	7
Expected Number of Days until Disclosure	50	16.6	8.3	5	3.3	2.4

reasonably assume that the people know the secret so that they can work together in some way. Of the many ways to tease out the consequences of this premise, the simplest suggestion is that each of the n people is in regular communication with each other and with no one else. This motivates the name of Clique Communication Model.

To specify the stochastic behavior of the CCM, we need a model for the communication between two people who are in on the secret. We begin by associating a Poisson process $N_{ij}(t)$ to each pair of individuals i and j so that $N_{ij}(t)$ counts the number of communications between person i and j in the time period $[0, t]$. Also, in the absence of more detailed information, one may as well assume that the $n(n-1)/2$ Poisson processes $N_{ij}(t)$, $1 \le i < j \le n$, are independent and have the same intensity λ. If $N(t)$ denotes the number of communications that take place in the clique in the time period $[0, t]$, then

$$N(t) = \sum_{1 \le i < j \le n} N_{ij}(t) \tag{2.2}$$

and $N(t)$ is again a Poisson process. Moreover, from the assumptions made about the N_{ij}, we see that $N(t)$ has a constant intensity μ and $\mu = \lambda n(n-1)/2$.

We now need a mechanism to disclose the secrets, and we make the basic assumption that disclosures take place only through a compromised communication. One way to incorporate this notion into our model is to associate a random variable X_{ijk} to the kth communication between individuals i and j such that the random variables X_{ijk} are independent and satisfy

$$P(X_{ijk} = 1) = p = 1 - P(X_{ijk} = 0).$$

The final assumption we make in the CCM is that the variables $\{X_{ijk}\}$ are independent of the processes $\{N_{ij}\}$, and of course the interpretation of X_{ijk} is that $X_{ijk} = 1$ if and only if the kth communication between i and j is compromised.

This completes the formal definition of the CCM, and the first task is to develop information about T, the time until disclosure. If $\hat{N}$ denotes the number of compromised communications, then by the independence hypotheses and the definition of p we see that $\hat{N}$ is again a Poisson process. This new process has constant intensity $p\mu$, and, from the definition of T, we have an expression for tail probability

$$P(T > t) = P(\hat{N}(t) = 0) = e^{-p\mu t}, \qquad t \ge 0. \tag{2.3}$$

By integrating $P(T > t)$ over $0 < t < \infty$, we find the expression for $E(T)$ cited in formula (2.1).

What does this formula for $E(T)$ say about the management of secrets? The most salient feature of this model is the built-in fact that the expected window of secrecy decays quadratically with the number of people who are in on the secret. But the expression for $E(T)$ also reflects the quantitative roles of λ and p. If we let $\rho = \lambda p$, then in sympathy with the usual terminology of survival analysis (Miller 1983) the parameter ρ is properly called the *hazard rate*. After making the wisest available choice for the value for n, the efforts of management are well invested on making ρ as small as possible.

If controlling the compromise probability p is impossible, or too expensive, or too hard on working conditions, the CCM says the only way to lower the hazard rate is to lower the communication rate λ. Mastering λ (and more generally p_{ij} and λ_{ij}) turns out to be a fruitful principle, and a later section shows how recombination of the p_{ij} can help create more secure communication.

There are several oversimplifications in the CCM. First, why should we assume that all the intensities of the processes $N_{ij}(t)$ are equal? If we let λ_{ij} denote the intensity of the $N_{ij}(t)$ process and let λ be the average of the λ_{ij}, then again we have that $\hat{N}(t)$ as defined by equation (2.2) is a Poisson process with intensity given by

$$p\mu = p \sum_{i<j} \lambda_{ij} = p\binom{n}{2}\lambda. \tag{2.4}$$

The derivation of the formula for the expectation of $E(T)$ can be derived just as before.

On the other hand, if one accepts constancy of the communications rate $\lambda_{ij} = \lambda$, then the probability p of compromised communication can be permitted to depend on i and j. Parallel with the case of (i, j)-dependent communication rates, we can let p denote the average of the p_{ij}, and we again recapture the exact expression (2.1) for $E(T)$. This persistence finally fails if one insists on varying both communication rates and compromise probabilities, but one still gets a succinct expression for the expected value of T,

$$ET = \Big(\sum_{1 \le i < j \le n} p_{ij}\lambda_{ij}\Big)^{-1}. \tag{2.5}$$

Naturally, this expression reduces to (2.2) for $p_{ij} = p$ and $\lambda_{ij} = \lambda$; and, although (2.5) is formally more general than the modest formula (2.2), the basic CCM is probably a more valuable tool for providing defensible advice in applications. For the expression (2.5) to provide honest guidance in a *bona fide* secrecy modeling context, one needs values for p_{ij} and λ_{ij}, and such detailed information is not likely to be available in many situations.

Still, there are useful inferences to be drawn from the formula for ET given by (2.5). In particular, (2.4) supports the intuitive fact that to maximize ET one should minimize the rate of communication across the least secure channels. The analysis of §4 systematizes this observation and uses (2.5) to frame an optimization criterion for communication over insecure networks.

We have just seen that the assumptions of the MCC concerning p_{ij} and λ_{ij} can be modified considerably, but it is not possible to drop our other assumptions and still retain simplicity. One can build more complex models with dependencies that are built into the processes $\{N_{ij}\}$ or the disclosure indicators $\{X_{ijk}\}$, but the resulting models are not likely to yield explicit formulas. In such cases one can proceed only by simulation. Here we take the point of view that such analyses are best deferred until one is confident that the simplest analytic models are well understood, both for their shortcomings and their strengths.

3. Birth Process Models

The Clique Communication Model is so natural one needs to fight not to become wedded to the notion that any secret disclosure model has to grow out of a communication model. Only by breaking radically from the CCM, can one develop confidence that the general ground of disclosure modeling is covered. Our second model, the Birth Process Model (or BPM), begins by acknowledging that in some situations the sharing of secrets between co-workers is almost inevitable. Obviously, in many situations there are innocent diffusions of the secret, and not every co-worker disclosure should be viewed as a critical breach. The widening circle of individuals who know the secret is thus analogous to a collection of organisms with a communicable infection, and the diffusion of the secret is much like the spread of an epidemic.

The key distinction between the CCM and the BPM is that the BPM acknowledges that simply enlarging the pool of people who know the secret does not automatically constitute disclosure. We could continue to consider disclosure by compromised communication, but to put as much distance as possible between the BPM and the CCM, we propose a new mechanism. We will assume the secret is eventually communicated to a *leaker*, *i.e.* a person who leaks the secret immediately after learning it.

Let $N(t)$ denote the number of people who know the secret at time t. It is reasonable to assume that during the small time period h each informed person tells the secret to an uninformed person with probability $\lambda h + o(h)$. From this assumption and the hypothesis that $N(t)$ is a counting process with independent increments, it follows that $N(t)$ is a pure birth process with parameters $\lambda_n = n\lambda$ for $n \geq N$. If N people initially know the secret, then the probability $P_n(t)$ that exactly n people know the secret at time t is given by Feller (1968):

$$P_n(t) = \binom{n-1}{N-1} e^{-\lambda N t}(1 - e^{-\lambda t})^{n-N}, \tag{3.1}$$

for all $n \geq N$, and, of course, $P_n(t) = 0$ for $n < N$. We can now easily calculate $E(T)$, the expected time until disclosure. First, we note

$$P(T > t) = \sum_{n=N}^{\infty} P(T > t, N(t) = n) = \sum_{n=N}^{\infty} P_n(t)(1-p)^n \tag{3.2}$$

where p is the probability that any given person in the system is a leaker. The integral of $P_n(t)$ simplifies to a beta integral so integration of (3.2) leads to a simple expression for $E(T)$:

$$E(T) = \sum_{n=N}^{\infty} (1-p)^n/(\lambda n) = \lambda^{-1}\{\log(1/p) - \sum_{n=1}^{N-1} (1-p)^n/n\}. \tag{3.3}$$

This formula is not as succinct as that obtained under the CCM, but it still contributes to our intuition. First, we see $E(T)$ decreases monotonely with N, the number of people originally in on the secret. The rate of decrease does not confirm Greer's square law, but this lack of confirmation is effectively good news since we wanted our second model to be genuinely different from the CCM. Another reassuring feature of (3.3) is that the slower the secret innocently diffuses, the longer the secret can be kept. In fact, $E(T)$ can be made arbitrarily large provided the diffusion rate λ of the secret can be made arbitrarily small. Finally, our intuition about the quantitative dependence of $E(T)$ on N in the BPM can be reinforced by considering Table 2 where $E(T)$ is evaluated for $p = 0.01$, $\lambda = 0.2$/day.

The BPM just analyzed can be embedded in a larger class by taking the parameters of the birth process to be $\lambda_n = (\lambda n)^\alpha$ instead of just $\lambda_n = \lambda n$. For $0 < \alpha \leq 1$ this parametrized birth model (BPM) is physically reasonable in the sense that $P(N(t) < \infty) = 1$ for all $0 < t < \infty$; moreover, the choice of α can accommodate useful changes in the modeling of the secret's diffusion. For example, if n is large it is feasible that it might become less likely for an individual to spread the secret because it is already widely known. This phenomenon can be accommodated by the parametrized birth model, though it would lie outside the scope of the BPM. The parametrized birth model is more specialized than the CCM or PBM and it will not be pursued further, but the flexibility it suggests is valuable enough to bare keeping in mind.

Before leaving the BPM and PBM, some comparisons should be made between the material of this section and related models for rumors and epidemics. In Moon (1972), Boyd and Steele (1979), and Haddad, Roy, and Schäffer (1987), stochastic rumor models were analyzed where the focus was on the number of communications in a network before a piece of information had been communicated to *all* the individuals in the network.

TABLE 2

Expected Number of Days to Disclosure under BPM

N	2	3	4	5	10	20	30
$E(T)$	18.0	15.6	14.0	12.8	9.3	6.2	4.6

The models of this section thus have more in common with the epidemic models of rumor spread that were studied in Rapoport and Rebhun (1952) and Daley and Kendall (1965) (cf. Chapters 9 and 10 of Bartholomew 1973). A key distinction between these models and BPM and PBM rests in the fact that under the BPM and PBM the critical variable is the time until a *specific subset* of the population obtains the information.

4. Simplest Countermeasure: Cautious Channels

One benefit of having models for secret disclosure is that they provide a context for analyzing countermeasures. The first one we explore rests on the systematic use of the more secure channels in the CCM. We recall that in the general CCM the communication rates λ_{ij} and the compromise probabilities p_{ij} depend on i and j, and the time until disclosure T satisfies

$$E(T) = (\sum_{i,j} p_{ij}\lambda_{ij})^{-1}. \tag{4.1}$$

Now, if the compromise probabilities satisfy $p_{ik} + p_{kj} < p_{ij}$ for some triple of distinct integers i, j and k, it is easy to increase the expected length of the window of security in the CCM; one just replaces any required call from i to j by a call routed through k. This single-step improvement carries the seed of a general method. The key to this method rests in viewing the compromise probability matrix (p_{ij}) as a matrix of edge costs on the complete directed graph with vertex set $\{1, 2, \ldots, n\}$, i.e. a weight of p_{ij} is assigned to the directed edge (i,j) from i to j. We begin by calculating the minimal cost path between i and j for each of the distinct pairs i and j. For small n this is easily obtained by inspection and for larger n one can use Dijkstra's algorithm (cf. Reingold, Nievergelt, and Deo 1977, pp. 341–346). Now, all communications between i and j are routed along the minimum weighted paths.

To weigh the benefits of this protocol, we consider the five-person network specified in Figure 1. For each pair of distinct vertices, we label the associated edge with p_{ij}, and for the sake of this example, we suppose $p_{ij} = p_{ji}$. Also, assuming $\lambda_{ij} = \lambda = 2$ calls/day, we find unrouted communication (4.1) lends to $E(T) = 16.6$ days. If we route calls along the paths that attain minimum cost p^*_{ij}, the weights on the edges of the network of Figure 1 can be replaced by the values in parentheses; and, for the time T^* until disclosure under the re-routing protocol, we find $E(T^*) = 25$ days. Thus, by simply re-routing calls

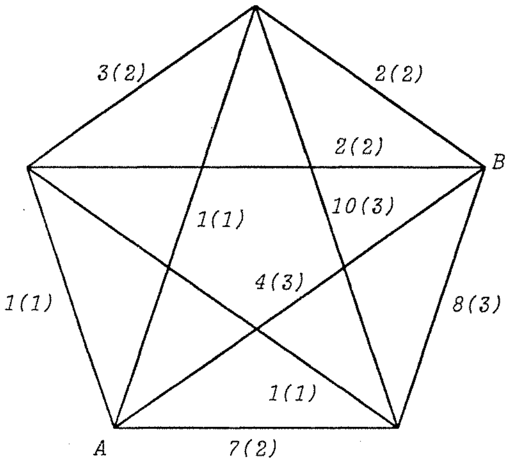

FIGURE 1. Communication Network with Compromise Probabilities (in 1/1000ths), Numbers in Parentheses Correspond to Shortest Paths.

along paths that minimize total probability weight, we provide a 50% improvement in an expected secrecy window.

5. Countermeasures by the Slice

One can sometimes divide a secret into parts so that all of the parts have to be learned by the adversary before there is a critical breach of the secret. The nature of the secret determines when such a decomposition is reasonable, and the possibility can exist just as easily in the context of the CCM, BPM, or the PBM. For simplicity, we will focus mainly on the CCM.

As before, we let $\hat{N}(t)$ denote the number of disclosures that have been made up to time t, but in this case *disclosure* means the disclosure of just one slice of the secret. At any time of increase of $\hat{N}(t)$, we will suppose that the slice which the adversary obtains is a random selection (with replacement) from the set $\{s_1, s_2, \ldots, s_k\}$ of all slices.

We originally motivated the CCM as a model of a clique of people working together, and such a motivation suggests that slicing the secret into k pieces has to lead to an increase in the communication rates λ_{ij}. A conservative way to increase the λ_{ij} would be multiplication by a factor of k, since one always has the option of communicating the slices one after another.

Assuming the λ_{ij} are replaced by $k\lambda_{ij}$, the process $\hat{N}(t)$ is easily seen to be a Poisson process with parameter $\mu' = k \sum \lambda_{ij} p_{ij}$. If we let $\hat{N}_i(t)$ denote the number of times slice s_i is disclosed by time t, then the k Poisson processes $\hat{N}_i(t)$, $1 \le i \le k$, are all independent and share the same intensity parameter $\mu'/k = \sum \lambda_{ij} p_{ij} = \mu$. If $T_i = \min\{t\colon \hat{N}_i(t) \ge 1\}$ denotes the time the adversary first obtains the ith slice, we see $T = \max_{1 \le i \le n} T_i$ is the first time at which all k slices of the secret have been disclosed. From the independence of the T_i, we see the expectation of T is

$$E(T) = \int_0^\infty \{1 - (1 - e^{-\mu t})^k\}\, dt = \mu^{-1} H_k \tag{5.1}$$

where H_k is the harmonic sum $1 + 2^{-1} + 3^{-1} + \cdots + k^{-1}$. An approximation for H_k that is accurate even for small k is $\ln(7k/4)$, so we have a simple approximation for $E(T)$:

$$E(T) \cong \mu^{-1} \ln(7k/4) = \tilde{E}(T). \tag{5.2}$$

Many practical secrets fail to permit decomposition into a large number of slices, so the most relevant values of k are the small ones. Table 3 provides the values of $E(T)$ and its approximation $\tilde{E}(T)$ for $1 \le k \le 5$.

From Table 3 we see that breaking a secret into four parts can more than double the expected length of the window of security. Moreover, this increase does not require any change in the basic technology as reflected in λ or p. One cannot ignore the potential importance of unmodeled costs, but the prospect of doubling the expected secrecy window is too attractive to ignore without reason. Many planners are well served by asking themselves if their secret might be sensibly decomposed into parts.

TABLE 3
Expected Time Until Disclosure

k	1	2	3	4	5
$E(T)$	μ^{-1}	$1.5\mu^{-1}$	$1.83\mu^{-1}$	$2.08\mu^{-1}$	$2.28\mu^{-1}$
$\tilde{E}(T)$	$1.01\mu^{-1}$	$1.50\mu^{-1}$	$1.83\mu^{-1}$	$2.08\mu^{-1}$	$2.28\mu^{-1}$

6. Slyer Countermeasures: Disinformation

It is as foolish to study secrecy without disinformation as it is to study poker without bluffing. But, even though disinformation countermeasures cannot be ignored, we cannot now embark on this systematic analysis. For one thing, such an analysis requires game theoretic tools that are remote from the elementary probability models we have counted on so far.

Instead of pursuing a full theory of disinformation we will develop a variant on the CCM that illustrates some essential features of a disinformation countermeasure. This particular method has the benefits of (1) being simple enough to be used and (2) providing a dramatic improvement in the expected size of the security window of the basic CCM.

In any model where some communications contain false information, the secret preservers need a method for knowing which communications are *bona fide*. The method we explore is driven by the principle "What I tell you twice is true." Formally, the protocol is specified by two rules:

(1) When a *bona fide* communication is to be sent, two messages are initiated simultaneously on different communication channels. If these messages confirm each other the communication is accepted as valid, otherwise the communication is ignored.

(2) At a rate that is much greater than the rate at which *bona fide* communications are sent, one sends pairs of messages simultaneously on different channels such that the two messages are *not* confirmatory.

How does this protocol influence the adversary's ability to discover the secret? Since the number of voided communications is large compared to the number of *bona fide* communications, the adversary cannot trust any single isolated intercepted message as being a part of a legitimate communication. The adversary has to wait until two simultaneous and corroborative messages are intercepted. Parallel to the CCM, the total demand for *bona fide* communications is modeled by a Poisson process with parameter $\lambda n(n-1)/2$, and p denotes the probability that any given communication will be compromised. Since the probability that both parts of the pair of messages that make up a *bona fide* communication are intercepted is p^2, the process $\hat{N}(t)$ that counts the number of valid communications that are intercepted by the adversary is a Poisson process with parameter $p^2\lambda n(n-1)/2$, and the time T until the secret is disclosed satisfies

$$E(T) = 2/p^2\lambda n(n-1). \tag{6.1}$$

The simplicity of the disinformation protocol and the straightforward derivation of formula (6.1) should not prejudice the practical implications of (6.1). For example, if we use the values $p = 0.01$ and $\lambda = 2/\text{day}$ as before, formula (6.1) says the expected length of the security window under the disinformation protocol is 100 times larger than it was under the basic CCM.

7. Summary and Conclusions

Three secret disclosure models and three classes of countermeasures have been introduced and analyzed. The Clique Communication Model that was introduced first has the benefit of an explicit square law. The CCM served to introduce the communication rates λ_{ij} and compromise probabilities p_{ij} that help in the effort to understand countermeasures.

The Birth Process Model (BPM) given after the CCM introduced the useful notion of innocent diffusions of a secret in addition to critical leaks. The Parametrized Birth Model (PBM) introduced as a variant of the BPM has the benefit of permitting a change in the rate of spread of innocent diffusions as the number of people who are in on the secret becomes larger.

A model for secrecy preservation creates a context for modeling countermeasures, and the three countermeasures explored here add considerably to the richness of our models. The method of Cautious Channels connects secrecy models with a classical problem of operations research and isolates a technique that improves secrecy without requiring technological innovation. Next, the technique of slicing is shown to provide an increase in the expected length of the window of secrecy from 50% to 128%, depending on whether two to five slices are feasible. Such an improvement seems impressive, but reconsideration is required when the disinformation method suggests that one can sometimes increase the expected window of security by a factor of 100.

All the models and countermeasures considered here make suppositions that are sometimes profoundly false. One of our unstated but persistent assumptions is that the adversary is passive. In our models the adversary waits and listens but never takes action to lure us into compromise. One should not forget that real adversaries also have active strategies. Some of these, like the "Trojan horse" analyzed in Karger (1987), are known to be powerfully destructive. Still, many systems are primarily subject to passive threats, and the logical first step is to model such systems as simply as possible.[1]

[1] The second identity of (5.1) and the approximation (5.2) were kindly suggested by a referee. Also, the model given in §3 has been substantially influenced by the referees' comments.

Research supported in part by National Science Foundation Grant #DMS-8414069.

References

BARTHOLOMEW, D. G., *Stochastic Models for Social Processes*, 2nd ed., John Wiley and Sons, New York, 1973.

BOYD, D. W. AND J. M. STEELE, "Random Exchanges of Information," *J. Appl. Probab.*, 16 (1979), 657–661.

CLANCY, T., *The Hunt for Red October*, Naval Institute Press, Annapolis, MD. (also published by Berkeley Books, New York), 1984.

DALEY, D. J. AND D. G. KENDALL, "Stochastic Rumours," *J. Inst. Math. Appl.*, 1 (1965), 42–55.

FELLER, W., *Introduction to Probability Theory and Its Applications*, Second Edition, John Wiley and Sons, New York, 1968.

HADDAD, R. W., S. ROY AND A. SCHÄFFER, "On Gossiping with Faculty Telephone Lines," *SIAM J. Algebraic Discrete Methods*, 8, 3 (1987), 439–445.

KARGER, P. A., "Limiting the Damage Potential of Discretionary Trojan Horses," 1987 *IEEE Sympos. Security and Privacy*, 32–37, IEEE Computer Society Press, Los Angeles, CA, 1987.

MILLER, R. G., JR., *Survival Analysis*, John Wiley and Sons, New York, 1981.

MOON, J. W., "Random Exchanges of Information," *Nieuw. Arch. Wisk.*, 20, 20 (1972), 246–249.

RAPOPORT, A. AND L. I. REBHUN, "On the Mathematical Theory of Rumour Spread," *Bull. Math. Biophysics*, 14 (1952), 375–383.

REINGOLD, E. M., J. NIEVERGELT AND N. DEO, *Combinatorial Algorithms: Theory and Practice*, Prentice-Hall, Inc., Englewood Cliffs, NJ, 1977.