

Covert Channels and Anonymizing Networks

Ira S. Moskowitz
Center for High Assurance Computer Systems
Naval Research Laboratory
Washington, DC 20375
moskowitz@itd.nrl.navy.mil

Daniel P. Crepeau
Transmission Technology Branch
Naval Research Laboratory
Washington, DC 20375
dcrepeau@itd.nrl.navy.mil

Richard E. Newman
University of Florida
CISE Department
Gainesville, FL 32611-6120
nemo@cise.ufl.edu

Allen R. Miller
Private Consultant
Washington, DC

ABSTRACT

There have long been threads of investigation into covert channels, and threads of investigation into anonymity, but these two closely related areas of information hiding have not been directly associated. This paper represents an initial inquiry into the relationship between covert channel capacity and anonymity, and poses more questions than it answers. Even this preliminary work has proven difficult, but in this investigation lies the hope of a deeper understanding of the nature of both areas. MIXes have been used for anonymity, where the concern is shielding the identity of the sender or the receiver of a message, or both. In contrast to traffic analysis prevention methods which conceal larger traffic patterns, we are concerned with how much information a sender to a MIX can leak to an eavesdropping outsider, despite the concealment efforts of MIXes acting as firewalls.

Categories and Subject Descriptors

H.1.1 [Models and Principles]: Systems and Information Theory—*Information theory*

General Terms

Theory

Keywords

anonymity, MIX, covert channel, information theory

1. INTRODUCTION

In this paper we discuss a particular covert channel that exists in an anonymizing network. We discuss how *less than*

perfect anonymity can inadvertently introduce covert communication channels. We do not discuss “fixes” to the covert channel problem as has been done in traffic analysis of network communications [16, 17, 26, 27, 28]. Rather, our interest is in measuring the covert channel capacity. These results can assist in bounds for covert channels, and lead one to consider different, or modified, design scenarios. Note that even though some may consider studying covert channels as being overly paranoid, covert channels should not be ignored [13] (a good starting place for the reader unfamiliar with covert channels).

We present some simplified scenarios as a first step in this analysis. Unfortunately, the mathematical details of the results showcased in this paper are quite complicated and detailed. Therefore, in the interest of writing a proceedings size paper, we have delegated the lengthier mathematical details to the internal (publicly available) tech report [14]. We have included the mathematical and information theoretic details for the simpler cases in this paper, in the hopes of giving the reader a taste for the more complex cases. We thank a reviewer for pointing out [1, 6, 11], where some informal studies of covert channels and anonymity were discussed.

There is always one special transmitting node in a network called *Alice*. Alice and possibly other transmitters have legitimate business transmitting messages to a set of Receivers $\{R_i | i = 1, 2, \dots, M\}$. These transmitters act completely independently of one another, and have no direct knowledge of each other’s recent transmission behavior. Alice may have some general knowledge of the long-term traffic levels produced by the other transmitters, e.g., the number of other transmitters and their probabilistic behavior, which can allow Alice to write a code that can improve the covert communication channel’s data rate. She cannot, however, perform short-term adaptation to their behavior. Our simplified communication is one-way (transmitters are never receivers). We also assume that there is a clock, and that transmissions only occur in the unit interval of time called a *tick*. Any subset of transmitters can each either send a single message to a single receiver in a tick, or not send a message at all. Each transmitter in a tick can send to a different receiver, and two or more transmitters may send to

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

WPES’03, October 30, 2003, Washington, DC, USA.
Copyright 2003 ACM 1-58113-776-1/03/0010 ...\$5.00.

the same receiver in the same tick. All messages’ contents are encrypted end-to-end.

There is also an eavesdropper on the network called *Eve*. Since all transmissions are encrypted, they appear to the eavesdropper *Eve* as having indistinguishable content. *Eve* may be either a global passive adversary (GPA), with the ability to see link traffic on every link in the network, or a restricted passive adversary (RPA), with the ability to observe traffic only on certain links.

Alice is not allowed any direct communication with *Eve*. However, Alice can influence what *Eve* sees on the network. We study network scenarios that attempt to achieve a degree of anonymity with respect to the network communication. That is, the networks are designed with various anonymity devices to prevent *Eve* from learning who is sending a message to whom. Even if a certain degree of anonymity is achieved, it still may be possible for Alice to communicate covertly with *Eve*. Note anonymous communication networks were *not* designed with this covert channel threat in mind. Our study of these anonymity networks caused us to realize that even in what appears to be a benign form of communication, information may still leak out of the network. This may cause the system designer to rethink and/or modify their ideas.

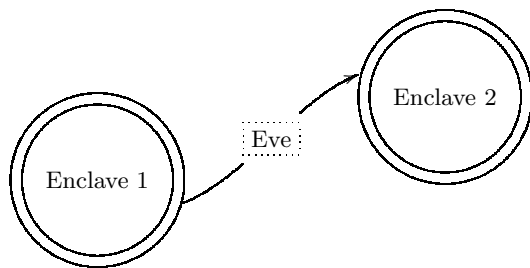


Figure 1: Restricted Passive Adversary Model.

The main thrust of this paper is to analyze the situation where there are two enclaves, communication between them is encrypted, and packets are sent only from the first enclave (which contains Alice) to the second (Fig. 1). *Eve* is able to monitor the communication from the first enclave to the second. Anonymity is “achieved” in that an eavesdropper such as *Eve* (as RPA) does not “know” who is sending a message (that is hidden inside of the first enclave) nor who is receiving the message (this can only be known if one is interior to the second enclave). *Eve* is only allowed to know how many messages per tick travel from the first enclave to the second. Nonetheless, Alice attempts to communicate covertly with *Eve*.

This paper analyzes the covert communication channel from Alice to *Eve*. We show that even if anonymity is taken into consideration with respect to system design, covert channels may remain. As a baseline, we first consider situations in which no attempt at anonymity has been made (only encryption of the messages, so that they all appear to be identical to an eavesdropper). Later, we will consider covert channel capacity in networks with the stronger anonymity controls just described.

2. BASE SCENARIO — NO ANONYMITY

One transmitter

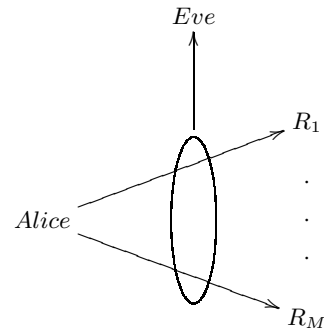


Figure 2: Global Passive Adversary Model.

Alice is the only transmitter, and there are M possible receivers. *Eve* has knowledge of the network traffic (*Eve* is a GPA — see Figure 2). The only properties that *Eve* can discern from a message is its source (trivially Alice) and its destination. Alice can use that fact to send information covertly to *Eve*. In this simplistic scenario *Eve* can see if Alice is sending a message, and if Alice is sending a message *Eve* can determine for which receiver the message is meant. This gives Alice the ability to signal *Eve* with an alphabet of $M+1$ symbols: M symbols for the M different receivers, and one symbol (“0”) for the choice of not sending a message.

Since nothing is able to interfere with Alice’s transmission, we have a noiseless discrete memoryless channel (DMC) modeling the covert channel, whose capacity is $\log(M+1)$ bits per tick.¹

Several transmitters

Now, if there are other transmitters aside from Alice, but their transmissions to any of the M receivers do not affect Alice’s transmissions, then the covert channel from Alice to *Eve* is as above. This would be the case if the links into a receiver can handle all of the traffic meant for them. Of course, if the link capacity into a transmitter *does* affect the number of receivable transmissions then that introduces noise into the channel and the capacity is obviously less than $\log(M+1)$. This is a course of research worth pursuit.

Anonymity discussion

In the above scenario Alice can obviously leak considerable information to *Eve*. This is no secret to the anonymity community, *e.g.*, [2, 3, 4, 5, 8, 18, 19, 22, 23] (while the preceding list is only a representative sample of papers/URLs on the topic, these papers relate particularly well to what we discuss in this paper). However, in the past the concerns have focused on retaining or regaining anonymity. It is the “anonymity lost” that we exploit for covert communication. If there were “*perfect*” anonymity,² then we would not expect to find a covert channel.

¹All logarithms are base 2, the units of capacity are bits per tick.

²We intentionally leave the notion of perfect anonymity as fuzzy in this paper. We ponder the somewhat circular question: If we did have perfect anonymity, how could we

To provide anonymity, transmissions from a transmitter are often first sent to an intermediary, such as a MIX [5] or an onion router [18], before they are forwarded to the receiver. This has the effect of hiding where the message is going. Thus, these intermediaries serve to anonymize the transmission. Of course, Eve still knows the set of those who receive a message, and she also knows the set of those who sent a message, but she does not know who sent a message to whom. It is interesting that, even when we seem to have “good” statistical anonymity, Alice may still non-trivially be able to communicate covertly with Eve.

The use of a MIX alone does not prevent Alice from covert communication with Eve. In fact there are two possible situations when Alice is the only transmitter.

1. Alice signals Eve by sending or not sending a message. A MIX alone does nothing to prevent Eve from learning this information (this is not what a MIX is designed to do). We discuss this further at the beginning of the next section. Therefore Alice has a noiseless channel to Eve, with capacity = 1.
2. Alice signals Eve by sending a message to any one of M different receivers. Eve simply sees where messages are going when they leave the MIX (a concern well-known to MIX designers). This allows a covert channel with a capacity of $\log(M + 1)$. If there are other users, their behavior affects what Eve is receiving and the capacity is then less than $\log(M + 1)$.

We will not study the latter situation in this paper, because we do not use pure MIXes. Instead, we use MIXes acting as firewalls.

3. SCENARIO 2: INDISTINGUISHABLE RECEIVERS- 2 MIX-FIREWALLS

Consider the situation in which every message goes into the anonymizing intermediary referred to as a MIX [5]. The MIX has the effect of hiding the “linking” knowledge of which transmission is sent to which receiver. In other words, Eve knows who is transmitting and who is receiving, but in general, Eve does not know which transmitter is sending to which receiver. This assumes that Eve is a GPA. Of course, if only one transmitter is operating then the MIX hides nothing. In other words the MIX gives statistical anonymity. The amount of anonymity has been measured as the log of the number of transmitters (*anonymity set size*), sometimes in conjunction with probabilistic behavior (e.g., [3, 4, 5, 8, 23]).

The main concern of this paper is not with measuring anonymity, rather it is the amount of covert information that may be leaked through less than perfect anonymity. However, we do note the very important observation from our study: *the ability to covertly communicate arises due to a lack of anonymity*. As the number of transmitters goes up and as the transmitters behave in a “uniform (equi-probabilistic) manner,” the anonymity increases and we will show that the covert channel capacity diminishes.

have covert communication? We thank P. Syverson for his thoughts.

For Scenario 2 we assume that there are transmitters Alice and Clueless_{*i*}, $i = 1, \dots, N$. The N Clueless_{*i*} transmitters behave independently of each other and of Alice, and they all have the same time-invariant probabilistic behavior. Throughout this paper we assume that Alice acts independently of the Clueless_{*i*}. Alice and the Clueless_{*i*} are hidden from Eve. They submit their messages to a MIX that also functions as a firewall. This first *MIX-firewall* acts as an exit point. This MIX-firewall sends its encrypted messages to a second MIX-firewall that is an entrance to a second hidden (from Eve) enclave. We further assume that Eve only has knowledge of how many messages come out of the first MIX-firewall per tick, and Eve does not know to whom the messages are going. Thus Eve is an RPA. The situation is described by the following diagram (Figure 3). This situa-

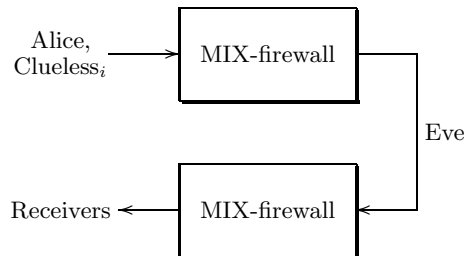


Figure 3: MIX-firewalls with Restricted Passive Adversary.

tion is realistic³ if the MIXes are acting as (first) firewall exit and (second) entrance points, or if the MIXes are onion-type routers acting as firewalls. Therefore, the only knowledge that Eve can get by eavesdropping is the number of messages per tick passing between the two MIX-firewalls. In other words, every tick, Eve observes the number of packets leaving the MIX-firewall and “receives” some number from the set $\{0, 1, \dots, N + 1\}$.

Therefore the only quantity observable by Eve that Alice can affect, per tick, is the number of messages that Eve counts. This covert channel is a discrete memoryless channel *with* noise since the Clueless_{*i*} randomly affect the output. Shannon’s information theory [24] tell us how useful the channel is.

Let us go back to the base scenario; here we stated that the capacity is obviously $\log(M + 1)$. How do we know that some other exploitation of the base scenario will not give us a higher capacity? The reason is that there are at most $M + 1$ symbols in whatever exploitation we use, and if the channel is noiseless we have maximized the capacity (this is related to the maximum entropy as discussed in [15].) For Scenario 2 capacity cannot be explained so easily and is the major study of this paper.

Keep in mind that for Scenario 2 it does not matter if there is one receiver or there are one hundred and one receivers. Eve can only count, and Alice or Clueless_{*i*} can only send

³Consider the case of packets from one LAN/enclave being sent to another LAN/enclave using IPSEC tunneling [10]. In this case, an eavesdropper can only count the number of outgoing messages destined for the receiving enclave. What goes on inside each LAN/enclave is hidden from an eavesdropper. If UDP with no application level ACKs is employed, communication is only one-way [20].

one message per tick. Therefore the number of receivers does not matter. It is only important that there is at least one receiver.

We break Scenario 2 down into four cases: 2.0, 2.1, 2.2, and 2.3. Case 2.3 is the general form of Scenario 2 and the first three are simplified special cases.

3.1 Two special cases of Scenario 2: — Alice alone, and with and one additional transmitter

Case 2.0 — Alice

This is the case where $N = 0$. Alice is the only transmitter. Alice sends either 0 (by not sending a message) or 0^c (by sending a message). Eve receives either $e_0 = 0$ (Alice did nothing) or $e_1 = 1$ (Alice sent a message to a receiver). The capacity of this noiseless covert channel is 1.

Note though the capacity is the maximum, over the probability x for Alice inputting a 0, of the mutual information $I(E, A)$. A is the distribution for Alice described by x , and E is the distribution for Eve. Since there is no noise, I is simply the entropy $H(E)$ describing Eve (which is maximized to 1 when $x = .5$).

$$I(E, A) = H(E) = -x \log x - (1 - x) \log(1 - x).$$

These terms are made precise later in this section.

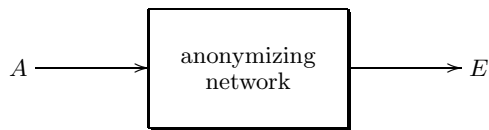
Case 2.1 — Alice and one additional transmitter (Clueless)

In this case $N = 1$. Therefore, Eve receives:

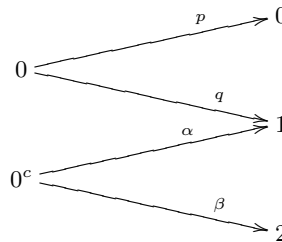
- 0 if neither Alice nor Clueless transmit;
- 1 if Alice does not transmit and Clueless does transmit, or Clueless transmits and Alice does not; or
- 2 if both Alice and Clueless transmit.

A is the input random variable describing Alice, and E is the output random variable describing Eve. Clueless contributes to the noise, but is not modeled as an input. Alice communicates with Eve via the covert channel. The input symbols for the channel are 0, which signifies that Alice is not transmitting a message to any receiver, and 0^c , which signifies that Alice is transmitting a message to some receiver (keep in mind that Alice is oblivious to the other transmitters).⁴

⁴At this point we caution the reader not to confuse Alice transmitting a message to a receiver R_i , and Alice communicating to Eve via the covert channel. Eve is *not* the receiver R_i in the sense of Alice or Clueless transmitting a message. Eve receives symbols via the covert channel from Alice. There are two different communication paths that must be kept separate. One is the legitimate network communication that the anonymizing device attempts to keep unknown. The other is the covert communication that Alice has to Eve. A way to stop the covert communication would be for the anonymizing device to pad [15, 16, 17, 26, 27] messages so that it would appear to Eve that both Alice and Clueless are transmitting a message. This inefficiency might be tolerated in such an ideal situation as Case 2.1, but such a strategy must be called into question when it comes to real traffic. In Case 2.1 the anonymizing effect is done by a MIX-firewall, which does not *a priori* pad. Of course, before advocating traffic padding one should be fully aware of the threat that the padding is intended to stop. Failure to understand the threat first is inadvisable since padding comes at the pragmatic costs of efficiency and proper network resource utilization.



(a) Channel block diagram



(b) Channel transition diagram

Figure 4: Channel model for Case 2.1

Part (b) of Fig. 4 shows the output symbols corresponding to the three states E might perceive. Let us consider the channel matrix.

$$M_{2.1} = \begin{matrix} & \begin{matrix} 0 & 1 & 2 \end{matrix} \\ \begin{matrix} 0 \\ 0^c \end{matrix} & \begin{pmatrix} p & q & 0 \\ 0 & \alpha & \beta \end{pmatrix} \end{matrix}$$

The 2×3 channel matrix $M_{2.1}[i, j]$ represents the conditional probability of Eve receiving the symbol j when Alice sends the symbol i . It follows that $p = \alpha$, and thus it trivially follows that $q = \beta$.

So our channel matrix simplifies to:

$$\begin{matrix} & \begin{matrix} 0 & 1 & 2 \end{matrix} \\ \begin{matrix} 0 \\ 0^c \end{matrix} & \begin{pmatrix} p & q & 0 \\ 0 & p & q \end{pmatrix} \end{matrix}$$

The probability that Alice sends a 0 is $P(A = 0) = x$, and therefore $P(A = 0^c) = 1 - x$. The term x is the only term that can be varied to achieve capacity. Here is where Alice may use knowledge of long-term transmission characteristics of the other transmitters, as well as how many other transmitters there are, to change her (long-term) behavior. As with other studies of covert channels [13] we are not concerned with source coding/decoding issues [24]. Our concern is the limits on how well a transmitter can “optimize” its bit rate to a receiver, given that a channel is noisy. Given a discrete random variable X , taking on the values $x_i, i = 1, \dots, n_X$, the entropy of X is:

$$H(X) = - \sum_{i=1}^{n_X} p(x_i) \log p(x_i) .$$

We use $p(x_i)$ as a shorthand notation for $P(X = x_i)$. Given two such discrete random variables X and Y we define the conditional entropy (equivocation) to be:

$$H(X|Y) = - \sum_{i=1}^{n_Y} p(y_i) \sum_{j=1}^{n_X} p(x_j|y_i) \log p(x_j|y_i) .$$

Given two such random variables we define the mutual in-

formation between them to be:

$$I(X, Y) = H(X) - H(X|Y).$$

Note that $H(X) - H(X|Y) = H(Y) - H(Y|X)$, so we see that $I(X, Y) = I(Y, X)$.

For a DMC whose transmitter random variable is X , and whose receiver random variable is Y , we define the *channel capacity* [24] to be:

$$C = \max_X I(X, Y),$$

where the maximization is over all possible distribution values $p(x_i)$ (that is, the $p(x_i)$ are all non-negative and sum to one).

For us, the capacity of the covert channel between Alice and Eve is

$$C = \max_x \{H(E) - H(E|A)\}.$$

Given the above channel matrix we have:

$$H(E) = -\{px \log px + [qx + p(1-x)] \log [qx + p(1-x)] + q(1-x) \log q(1-x)\}.$$

$$\text{and } H(E|A) = -\sum_{i=0}^1 p(a_i) \sum_{j=0}^2 p(e_j|a_i) \log p(e_j|a_i) = h(p).$$

Where $h(p)$ denotes the function $-p \log p - (1-p) \log(1-p)$. Thus,

$$C = \max_x \left\{ -\left(px \log px + [qx + p(1-x)] \log [qx + p(1-x)] + q(1-x) \log q(1-x) \right) - h(p) \right\}.$$

We cannot analytically find the x that maximizes the mutual information, even doing the standard trick of setting the derivative of the mutual information to zero. However, we numerically show our results in Figure 5.

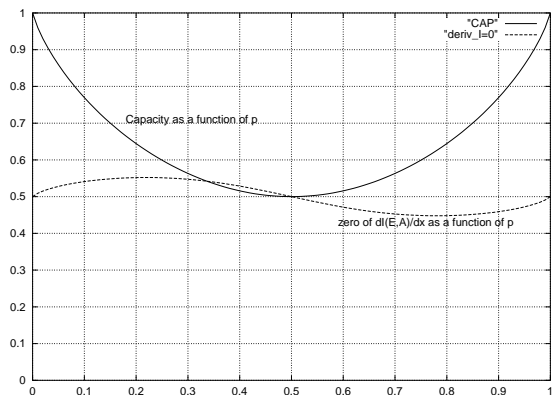


Figure 5: Plots of covert channel capacity as a function of p , and of the x value that maximizes the mutual information as a function of p .

We see in Figure 5 certain symmetries. The capacity graph is symmetric about $p = .5$, and the graph of the x that achieves capacity is skew-symmetric about $p = .5$.

Consider the two situations where $p = \epsilon$, and where $p = 1 - \epsilon$; in both situations $0 \leq \epsilon \leq .5$. Let x_ϵ be the probability for the input symbol 0 that achieves capacity in the first situation, and let $x_{1-\epsilon}$ be the probability that achieves capacity for the second situation. For the first situation we have that $1 - x_\epsilon$ is the capacity achieving probability for the output symbol 0^c , and similarly for the second situation $1 - x_{1-\epsilon}$ is the capacity achieving probability for the output symbol 0^c . Physically the two situations are “the same” if we reverse the roles of the outputs symbols 0 and 2. Therefore $x_\epsilon = 1 - x_{1-\epsilon}$. Writing x_ϵ as $x_\epsilon = \frac{1}{2} + \Delta$, we see that $x_{1-\epsilon} = \frac{1}{2} - \Delta$; this is what the lower dotted plot shows in Figure 5 ($\epsilon = 1/2 \Rightarrow \Delta = 0$).

OBSERVATION 1. *In conditions of very little extra traffic, or very high extra traffic, the covert channel from Alice to Eve has higher capacity.*

OBSERVATION 2. *The capacity $C(p)$, as a function of p is strictly bounded below by $C(.5)$, and $C(.5)$ is achieved when the mutual information is evaluated at $x = .5$.*

It is obvious that very little extra traffic corresponds to very little noise. At first glance though, it seems counterintuitive that heavy traffic also corresponds to a small amount of noise. This is because the high traffic is used as a baseline against which to signal. This is analogous to transmission of bits over a channel where the bit error rate (BER) P_e is greater than 1/2. In this case, the capacity of the channel is the same as that of a channel with BER of $1 - P_e$, by first inverting all the bits. It is the in-between situations that negatively affect the signaling ability of Alice. But, even in the noisiest case (i.e., where $p = .5$) Alice can still transmit with a capacity of a half bit per tick.

Note that we can never guarantee error-free transmission, no matter how we group the output symbols. In fact, it is possible that the outputs will always be the symbol 1 (of course the probability of this quickly approaches zero, as the number of transmissions goes up). So this covert channel has a *zero-error capacity* [25] of zero. Capacity is a useful measure of a communication channel if the assumption is that the transmitter can transmit a large number of times. With a large number of transmissions, an error-correcting code can be utilized so as to achieve a rate close to capacity. If the transmitter only transmits a small number of transmissions, then using the capacity alone can be misleading.

3.2 Case 2.2—Alice and two additional transmitters ($N = 2$)

This is similar to Case 2.1, the difference being that we have three possible transmitters, A (random variable as before) for Alice, who is attempting to communicate covertly with E (random variable as before) for Eve, and two other benign “clueless” transmitters. Since the MIX-firewalls only allow Eve to count the number of outgoing messages, our covert channel has four possible output symbols (the inputs are as before 0, for Alice not sending a message, and 0^c , if Alice does send a message). The outputs are:

- 0 — No one sends a message;

- 1 — Alice sends a message, and neither Clueless_i send a message; or, Alice does not send a message, and one, and only one, Clueless_i sends a message;
- 2 — Alice sends a message and one, and only one, Clueless_i sends a message; or, Alice does not send a message and both Clueless_i send a message;
- 3 — Alice, Clueless₁, and Clueless₂ all send a message.

As stated earlier we assume that Clueless₁ and Clueless₂ act independently of each other (and Alice is independent of them). Therefore, if, as before, p is the probability of a clueless transmitter (Clueless₁ or Clueless₂) not sending a message into the MIX-firewall, and $q = 1 - p$ is the probability of a clueless transmitter sending a message, the conditional probabilities of E given Alice sending 0 are show in the covert channel diagram and channel matrix in Figure 6.

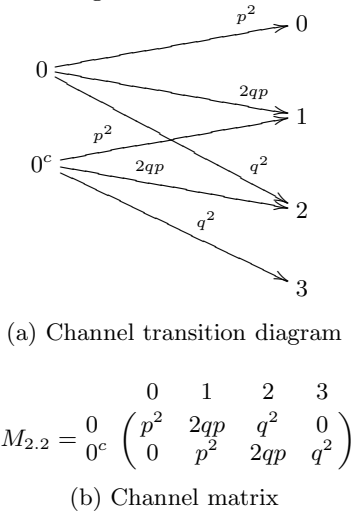


Figure 6: Channel for Case 2.2.

We can easily observe that the zero-error capacity is zero because the output symbols 1 and 2 can both be received if 0 or 0^c is transmitted. Therefore there is always some statistical error in what is received. This is similar to Case 2.1. For capacity itself, after some numerical calculation we plot the capacity in Fig. 7.

Except for the boundary values, the capacity is always less for a given p with three transmitters (two clueless) than with two (one clueless). This is not surprising, the extra clueless transmitter means extra noise. Note that the noisiest case is when $p = .5$, which again acts as a lower bound.

Unfortunately we cannot derive closed form solutions even for these simple cases. Therefore, it seems unlikely that we can derive a closed form for the general case of N clueless transmitters in addition to Alice. Of course, we could still derive the capacity numerically. However, we are able to obtain some bounding results.

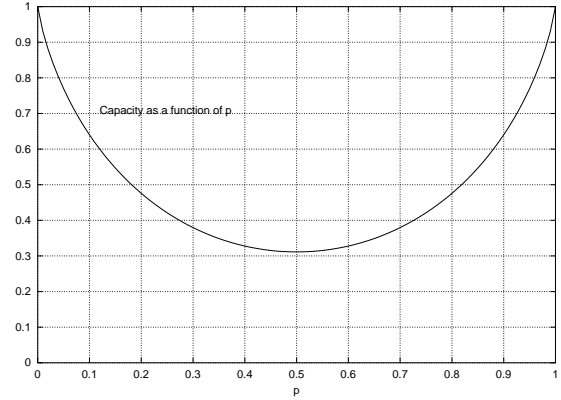


Figure 7: Capacity as a function of p for Alice with two additional transmitters.

3.3 Case 2.3—Alice and N additional transmitters

Case 2.3 is the general form of Scenario 2, see Figure 8. Now⁵ we imagine that there are $N + 1$ transmitters, Alice is one of them, and the other N are all independently identical clueless transmitters. That is, there are transmitters Clueless₁, Clueless₂, ..., Clueless _{N} . Again, Eve can only see how many messages are leaving the first MIX-firewall headed for the second MIX-firewall. Therefore Eve can determine if there are $0, 1, \dots, N + 1$ messages leaving the firewall. That is all Eve can determine. Therefore, there are still the two input symbols $a_0 = 0$ and $a_1 = 0^c$, but we have $N + 2$ output symbols. The probability that Clueless _{i} does not send a message is still p , and that it does send a message is $q = 1 - p$. Now, calculate the channel matrix. Keep in mind that Alice acts independently of the Clueless _{i} .

Alice sends a 0.

- For Eve to receive e_k (that is $E = k$), $0 \leq k \leq N$ we need k of the clueless transmitters to send a message, and $N - k$ not to send a message. Therefore,

$$p(e_k | A = 0) = \binom{N}{k} p^{N-k} q^k, \quad 0 \leq k \leq N.$$

- $p(e_{N+1} | A = 0) = 0$.

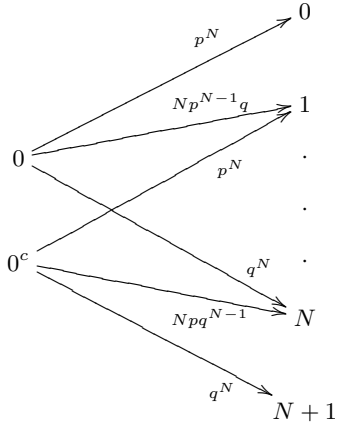
Alice sends a 0^c .

- $p(e_0 | A = 0^c) = 0$, since the event never happens.
- For Eve to receive e_k (that is $E = k$), $1 \leq k \leq N + 1$ we need $k - 1$ of the clueless transmitters to send a message, and $N - k + 1$ not to send a message.

$$p(e_k | A = 0^c) = \binom{N}{k-1} p^{N-k+1} q^{k-1}, \quad 1 \leq k \leq N + 1.$$

We delegate to the appendix the outline of the following important results (the full details and proofs are in [14]).

⁵One could relax the assumption that all the Clueless _{i} have identical and independent behavior.



(a) Channel transition diagram

The channel matrix $M_{3,N}$ is

$$\begin{matrix} & 0 & 1 & 2 & \dots & N & N+1 \\
 \begin{matrix} 0 \\ 0^c \end{matrix} & \begin{pmatrix} p^N & Np^{N-1}q & \binom{N}{2}p^{N-2}q^2 & \dots & q^N & 0 \\ 0 & p^N & Npq^{N-1} & \dots & Npq^{N-1} & q^N \end{pmatrix}
 \end{matrix}$$

(b) Channel matrix

Figure 8: Channel for Case 2.3, the general case of N clueless users.

- For any p , $C(p)$ is strictly bounded below by $C(.5)$.
- As the number of clueless transmitters goes to infinity, $C(.5)$ goes to zero.
- $C(p)$ is a continuous function of p .

4. COMMENTS, GENERALIZATIONS & FUTURE WORK

We first note that despite the obfuscation provided by MIX-firewalls, and the attendant noise introduced by other transmitters, Alice is still able to transmit information to Eve. At this point, we recall our earlier observations and add to them below.

1. In conditions of very little extra traffic, or very high extra traffic, the covert channel from Alice to Eve has higher capacity.
2. The capacity $C(p)$, as a function of p is strictly bounded below by $C(.5)$, and $C(.5)$ is achieved when the mutual information is evaluated at $x = .5$ (of course $p = .5$ also in this situation).
3. The capacity $C(p)$, as a function of p is strictly bounded below by a function that decreases monotonically to zero as the number of transmitters increases, but is never zero.
4. The bias in the code used by Alice to achieve the optimum data rate on the channel is not always $x = 0.5$, but it is never far from 0.5, and our preliminary experimental results indicate that the difference in capacity is minor.

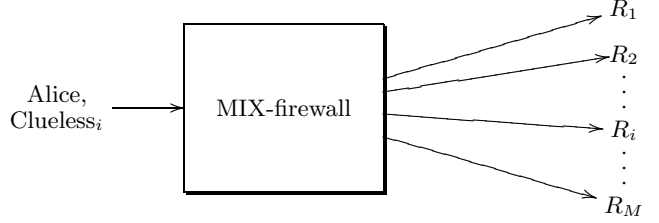


Figure 9: Exit firewall only

This last observation agrees with [12], which presents the general result that in DMCs, mutual information bit rates obtained by using $x = .5$ is no less than 94.21% of the channel capacity. Even if Alice has no knowledge of the probabilistic behavior of the other transmitters, her data rate will not be too far from optimal if she uses an unbiased code. (Note, however, that the coding rate is very much dependent on knowledge of the number of other transmitters and their behavior.)

In future work we will also analyze the situation where we have only an exit point MIX-firewall as shown in Figure 9.

We have M receivers denoted R_1, \dots, R_M . Eve still does not know directly who sent a message, but Eve does know where messages are going. This increases the capacity of the covert channel. Alice now instead of just sending 0 or 0^c can send: 0 (not transmitting); 1 (message to the first receiver), \dots , i (message to the i th receiver, \dots , M (message to the M th receiver). The greatest the capacity can be is $\log(M + 1)$. Of course if $M = 1$ the situation reduces to Scenario 2.

(See [14] for other related scenarios.)

Other areas begging for further investigation include scenarios in which there is limited network capacity (on links or aggregate), whether or not there is anonymity. We are currently investigating this using the model in which at most B messages can be sent through the network (as output from a sender of as output of a MIX-firewall) in a given tick, and if there are more than B messages awaiting transmission, B of them are chosen at random for delivery. This may relate the work to more sophisticated MIX models, such as pool MIXes, which is also desirable.

A deeper issue raised in this preliminary paper is that of the relationship between anonymity and covert channel capacity (fixing the other factors that affect capacity). It seems evident that as system level anonymity increases in the simple models shown here (i.e., the number of potential senders increases), the minimum capacity decreases to zero. However, as the probability that a Clueless sender transmits in a given tick increases, the expected number of actual senders in a given time tick also increases, hence the anonymity increases, but the capacity of the covert channel increases once this probability exceeds 0.5. The relationships are not simple, but their discovery has the potential to increase our understanding of fundamental aspects of network design.

5. ACKNOWLEDGEMENTS

We are grateful to Paul Syverson for his discussions about anonymity, to LiWu Chang for his assistance with the mathematical results, and also a special thanks to Gerard All-

wein for his technical expertise. We thank the anonymous reviewers for their helpful comments. This paper is US Government Work. Research supported by the Office of Naval Research.

6. REFERENCES

- [1] Dakshi Agrawal, Dogan Kesdogan, and Stefan Penz. Probabilistic treatment of MIXes to hamper traffic analysis. In *IEEE Symposium on Security and Privacy*, pages 16–27, Oakland, California, May 2003.
- [2] The anonymizer. <http://www.anonymizer.com/>.
- [3] Oliver Berthold, Hannes Federrath, and Stefan Köpsell. Web MIXes: A system for anonymous and unobservable internet access. In Hannes Federrath, editor, *Designing Privacy Enhancing Technologies: Design Issues in Anonymity and Observability*, pages 115–129. Springer-Verlag, LNCS 2009, July 2000.
- [4] Oliver Berthold, Andreas Pfitzmann, and Ronny Standke. The disadvantages of free MIX routes and how to overcome them. In Hannes Federrath, editor, *Designing Privacy Enhancing Technologies: Design Issues in Anonymity and Observability*, pages 27–45. Springer-Verlag, LNCS 2009, July 2000.
- [5] David Chaum. Untraceable electronic mail, return addresses, and digital pseudonyms. *Communications of the ACM*, 24(2):84–88, 1981.
- [6] Richard Clayton, George Danezis, and Markus G. Kuhn. Real world patterns of failure in anonymity systems. In Ira S. Moskowitz, editor, *Information Hiding, 4th International Workshop (IH 2001)*, pages 230–244. Springer-Verlag, LNCS 2137, 2001.
- [7] Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory*. Wiley, 1991.
- [8] Claudia Díaz, Stefaan Seys, Joris Claessens, and Bart Preneel. Towards measuring anonymity. In Paul Syverson and Roger Dingledine, editors, *Privacy Enhancing Technologies (PET 2002)*. Springer-Verlag, LNCS 2482, April 2002.
- [9] Robert G. Gallager. *Information Theory and Reliable Communication*. Wiley, 1968.
- [10] S. Kent and R. Atkinson. Security architecture for the Internet Protocol, 1998.
- [11] Dogan Kesdogan, Dakshi Agrawal, and Stefan Penz. Limits of anonymity in open environments. In F.A.P. Petitcolas, editor, *Information Hiding, 5th International Workshop (IH 2002)*, pages 53–69. Springer-Verlag, LNCS 2578, 2002.
- [12] E.E. Majani and H. Rumsey. Two results on binary input discrete memoryless channels. In *IEEE International Symposium on Information Theory*, page 104, June 1991.
- [13] Ira S. Moskowitz and Myong H. Kang. Covert channels — here to stay? In *Proc. COMPASS'94*, pages 235–243, Gaithersburg, MD, June 27- July 1 1994. IEEE Press.
- [14] Ira S. Moskowitz, Richard E. Newman, Daniel P. Crepeau, and Allen R. Miller. A detailed mathematical analysis of a class of covert channels arising in certain anonymizing networks. In *NRL Memorandum Report, NRL/MR/5540-03-8691*, 2003. <http://chacs.nrl.navy.mil/publications>
- [15] Richard E. Newman, Ira S. Moskowitz, Paul Syverson, and Andrei Serjantov. Metrics for traffic analysis prevention. In *PET 2003*, Dresden, March 2003.
- [16] R. E. Newman-Wolfe and B. R. Venkatraman. High level prevention of traffic analysis. In *Proc. IEEE/ACM Seventh Annual Computer Security Applications Conference*, pages 102–109, San Antonio, TX, Dec 2-6 1991. IEEE CS Press.
- [17] R. E. Newman-Wolfe and B. R. Venkatraman. Performance analysis of a method for high level prevention of traffic analysis. In *Proc. IEEE/ACM Eighth Annual Computer Security Applications Conference*, pages 123–130, San Antonio, TX, Nov 30-Dec 4 1992. IEEE CS Press.
- [18] Onion routing home page. <http://www.onion-router.net>.
- [19] Andreas Pfitzmann and Marit Köhntopp. Anonymity, unobservability and pseudonymity — a proposal for terminology. In Hannes Federrath, editor, *Designing Privacy Enhancing Technologies: Design Issues in Anonymity and Observability*, pages 1–9. Springer-Verlag, LNCS 2009, July 2000.
- [20] J. Postel. User Datagram Protocol, 1980.
- [21] A.P. Prudnikov, Yu. A. Brychkov, and O.I. Marichev. *Integrals and Series, Volume 1*. Gordon and Breach, 1986.
- [22] Michael K. Reiter and Aviel D. Rubin. Crowds: anonymity for web transactions. *ACM Transactions on Information and System Security*, 1(1):66–92, 1998.
- [23] Andrei Serjantov and George Danezis. Towards an information theoretic metric for anonymity. In Paul Syverson and Roger Dingledine, editors, *Privacy Enhancing Technologies (PET 2002)*. Springer-Verlag, LNCS 2482, April 2002.
- [24] Claude E. Shannon. The mathematical theory of communication. *Bell Systems Technical Journal*, 30:50–64, 1948.
- [25] Claude E. Shannon. The zero error capacity of a noisy channel. *IRE Trans. on Information Theory*, Vol. IT-2:S8–S19, September 1956.
- [26] B. R. Venkatraman and R. E. Newman-Wolfe. Transmission schedules to prevent traffic analysis. In *Proc. IEEE/ACM Ninth Annual Computer Security Applications Conference*, pages 108–115, Orlando, FL, December 6-10 1993. IEEE CS Press.
- [27] B. R. Venkatraman and R. E. Newman-Wolfe. Performance analysis of a method for high level prevention of traffic analysis using measurements from a campus network. In *Proc. IEEE/ACM Tenth Annual Computer Security Applications Conference*, pages 288–297, Orlando, FL, December 5-9 1994. IEEE CS Press.
- [28] B. R. Venkatraman and R. E. Newman-Wolfe. Capacity estimation and auditability of network covert channels. In *Proc. IEEE Symposium on Security and Privacy*, pages 186–198, Oakland, CA, May 8-10 1995. IEEE CS Press.

APPENDIX

A. APPENDIX

Now we show that $C(.5)$ is a strict lower bound for $C(p)$, and that as the number of clueless transmitters goes to infinity that $C(.5)$ goes to zero. We also discuss a continuity result for $C(p)$. Now we continue with the general case 2.3.

Since $p(e_k) = p(e_k|A=0)P(A=0) + p(e_k|A=0^c)P(A=0^c)$, we have that

$$\begin{aligned} p(e_0) &= xp^N, \\ p(e_k) &= x \binom{N}{k} p^{N-k} q^k + \\ &\quad (1-x) \binom{N}{k-1} p^{N-k+1} q^{k-1}, \quad 1 \leq k \leq N \\ p(e_{N+1}) &= (1-x)q^N. \end{aligned}$$

The **mutual information** is

$$\begin{aligned} I(E, A) &= - \left\{ xp^N \log xp^N + \right. \\ &\quad \sum_{k=1}^N \left[x \binom{N}{k} p^{N-k} q^k + \right. \\ &\quad \quad \left. (1-x) \binom{N}{k-1} p^{N-k+1} q^{k-1} \right] \times \\ &\quad \log \left[x \binom{N}{k} p^{N-k} q^k + \right. \\ &\quad \quad \left. (1-x) \binom{N}{k-1} p^{N-k+1} q^{k-1} \right] + \\ &\quad \left. (1-x)q^N \log(1-x)q^N \right\} \\ &\quad + \sum_{l=0}^N \left[\binom{N}{l} p^{N-l} q^l \right] \log \left[\binom{N}{l} p^{N-l} q^l \right] \end{aligned}$$

(For Case 2.1 (one Clueless in addition to Alice) and for Case 2.2 (two clueless in addition to Alice) we discussed the symmetry about $p = .5$ informally.)

THEOREM 1. $I(E, A)|_{x,p} = I(E, A)|_{1-x,q}$

PROOF: See [14]

We will need the following in the rest of the appendix so we will consider $I(E, A)|_{p=.5} = H(E)_{p=.5} - H(E|A)_{p=.5}$ now.

Consider the entropy of E evaluated when $p = \frac{1}{2}$.

$$\begin{aligned} H(E)|_{p=.5} &= - \left\{ x \left(\frac{1}{2} \right)^N \log x \left(\frac{1}{2} \right)^N + \right. \\ &\quad \sum_{k=1}^N \left[x \binom{N}{k} \left(\frac{1}{2} \right)^N + \right. \\ &\quad \quad \left. (1-x) \binom{N}{k-1} \left(\frac{1}{2} \right)^N \right] \times \\ &\quad \log \left[x \binom{N}{k} \left(\frac{1}{2} \right)^N + \right. \\ &\quad \quad \left. (1-x) \binom{N}{k-1} \left(\frac{1}{2} \right)^N \right] + \\ &\quad \left. (1-x) \left(\frac{1}{2} \right)^N \log(1-x) \left(\frac{1}{2} \right)^N \right\} \end{aligned}$$

Consider the conditional entropy when $p = \frac{1}{2}$.

$$H(E|A)|_{p=.5} = N - \left(\frac{1}{2} \right)^N \sum_{l=0}^N \binom{N}{l} \log \binom{N}{l}$$

Note that $H(E|A)|_{p=.5}$ is independent of x . Keep in mind that we may express the mutual information evaluated at (x', p') by the slightly overloaded notation $I(E, A)|_{x=x', p=p'}$. Of course $I(E, A)|_{p=p'}$ is simply just a function of x , and $I(E, A)|_{x=x'}$ is a function of p .

DEFINITION 1. We say that an arbitrary (real valued) function is not locally-constant iff for all x with $f(x)$ defined at x , and for every $\delta > 0$, there exists an x' such that $d(x', x) < \delta$ (i.e., x' in the neighborhood of x) with $f(x') \neq f(x)$.

That is, for no neighborhood, no matter how small, is the function constant.

DEFINITION 2. We say that a function $f : [0, 1] \rightarrow \mathfrak{R}$ is symmetric about $x = .5$, iff $f(x) = f(1-x)$.

OBSERVATION 3. If $f(x)$ is symmetric about $x = .5$ and it is concave down (convex up) then $f(.5)$ is a maximum (minimum) value. Further, if $f(x)$ is not locally-constant then $.5$ is the only such critical point.

THEOREM 2. $I(E, A)|_{p=.5}$ is symmetric about $x = .5$.

PROOF: By Thm. 1, $I(E, A)|_{x,.5} = I(E, A)|_{1-x,.5}$.

THEOREM 3. $C(.5) = I(E, A)|_{x=.5, p=.5}$.

PROOF: By Theorem 2, we know that $I(E, A)|_{p=.5}$ is symmetric about $x = .5$, and [9][Thm. 4.4.2]&[7][Thm.2.7.4] show that $I(E, A)|_{p=.5}$ (and in general $I(E, A)$ for fixed p) is concave down. Therefore, from Observation 1, $I(E, A)|_{p=.5}$ obtains its maximum value when $x = .5$. Since capacity, when $p = .5$, is the maximum of $I(E, A)|_{p=.5}$, we are done.

THEOREM 4. $C(p) \geq I(E, A)|_{x=.5, p=.5}$.

PROOF: By definition $C(p) \geq I(E, A)|_{x=.5}$, since capacity is the maximum of the mutual information. For x fixed, $I(E, A)|_x$ is a convex up function of p (see [9][Thm. 4.4.2]

and [7][Thm.2.7.4]). By Thm. 1 we see that $I(E, A)|_{x=.5}$ is symmetric about $p = .5$. By Observation 3 we see that $I(E, A)|_{x=.5} \geq I(E, A)|_{x=.5, p=.5}$.

This allows us to use $I(E, A)|_{x=.5, p=.5}$ (simple single value) as a lower bound for the covert channel capacity.

COROLLARY 1. $C(p) \geq C(.5)$

PROOF: Apply Theorems 3 and 4 together.

THEOREM 5. $C(p) = C(1 - p)$ and if x_p is the unique x such that $C(p) = I(E, A)|_{x_p, p}$, then $x_{1-p} = 1 - x_p$.

PROOF: This trivially follows from Thm. 1 and the uniqueness (follows from the concavity properties and the fact that the mutual information is not locally-constant) of the critical x value.

Let us now use these results to bound capacity from below. After many calculations and simplifications [14] we obtain

$$C(.5) = 1 - \left(\frac{1}{2}\right)^N \sum_{k=0}^N \left\{ \frac{1}{2} \binom{N+1}{k} \log \binom{N+1}{k} - \binom{N}{k} \log \binom{N}{k} \right\}. \quad (1)$$

We show some numerical results for C .

N	$C(.5)$	N	$C(.5)$
1	0.500000	13	0.053593
2	0.311278	14	0.049873
3	0.219361	15	0.046638
4	0.167553	16	0.043799
5	0.135170	17	0.041287
6	0.113278	18	0.039048
7	0.097558	19	0.037039
8	0.085730	20	0.035228
9	0.076502	21	0.033586
10	0.069092	22	0.032090
11	0.063007	23	0.030722
12	0.057917	24	0.029466
		25	0.028309

$C(.5) =$ **lower capacity bounds for all p , $N = 1, \dots, 25$**

Note that in the general circumstances of Case 2.3, if $p = 0$ (or similarly $q = 0$), we have a noiseless channel and the capacity is one, which is achieved when $x = .5$. So we see that 1 is a tight upper bound for the capacity. Therefore we have the following result:

For Alice and $N(N > 0)$ transmitters: $C(.5) \leq C(p) \leq 1$

and bounds ON $C(p)$ are tight. Of course keep in mind the result from Case 2.0:

For Alice and no additional transmitters: Capacity = 1.

As N grows so does the noise. Therefore, we see that the capacity is non-increasing. We are interested in the lower bound $C(.5)$. We have numerically calculated $C(.5)$ to $N = 7750$ and have shown that $C(.5)$ is monotonically decreasing to zero (for $N=7750$, $C(.5) = .000093$). We can (but do not since it is many pages in length) analytically show $C(.5)$ is monotonic decreasing. That is not surprising since increasing the number of clueless users increases the noise, but it

is surprising that it is so difficult to show that $C(.5)$ goes to zero as N goes to infinity. Below we discuss that fact, leaving the interesting and subtle details to [14].

From Eq. 1 we can express $C(.5)$ as

$$C(.5) = 1 - \left(\frac{1}{2}\right)^N S(N),$$

where

$$S(N) \triangleq \sum_{k=0}^N \left\{ \frac{1}{2} \binom{N+1}{k} \log \binom{N+1}{k} - \binom{N}{k} \log \binom{N}{k} \right\}.$$

$$\text{THEOREM 6. } S(N) = 2^N \log(N+1) - \sum_{k=0}^N \binom{N}{k} \log(k+1)$$

PROOF: Not shown, basically involves combinatorial identities.

Keep in mind our goal is to study the behavior of $C(.5)$ as $N \rightarrow \infty$. However, first we need a technical lemma.

LEMMA 1. $\sum_{k=1}^N \binom{N}{k} k^p = 2^{N-p} Q_p(N)$, for $p < N$, where $Q_p(N)$ is a monic polynomial in N of degree p .

PROOF: Follows from [21, Formulas 1,2,7,8,9,10 p. 608].

THEOREM 7. $\lim_{N \rightarrow \infty} C(.5) = 0$.

PROOF: The proof is asymptotic in nature, but follows by applying Lemma 1 to Thm. 6.

A.1 Continuity

For Scenario 2 we wished to say that capacity was a continuous function of p . We thought that we could just use some standard information-theoretic result. Unfortunately, we could not find such a result. We do not think that it would be too hard to argue from the various concavity properties of mutual information that $C(p)$ is a continuous function (of p). However, we decided to present a more general result which relies on the following theorem.

THEOREM 8. Let $F(x, p)$ be a continuous function⁶ defined on $[0, 1] \times U$, U an arbitrary subset of the reals, and assume that for each fixed p , $F(x, p)$ achieves a maximum denoted as $\Gamma(p)$. Then $\Gamma(p)$ is a continuous function of p .

PROOF: Not shown — standard analysis result using compactness arguments.

We believe that continuity results such as these are important, but they seem to be overlooked in the literature. Note we can replace the closed interval $[0, 1]$ by any compact subset of the reals.

⁶Of course in this paper all functions are real valued.