PERFECT SECRECY AND ADVERSARIAL INDISTINGUISHABILITY

BURTON ROSENBERG
UNIVERSITY OF MIAMI

Contents

1. Perfect Secrecy

Claude Shannon introduced an entropy model for information, and applied it to secrecy in communications. It supposes a source of information, Alice, which chooses among a set of possible messages. There is associated with this choice a likelihood that Alice would chose a particular message. Symbols are then sent across a channel to Bob. These symbols should refine Bob's likelihood function, emphasizing the likelihood of Alice's chosen message. Alice and Bob share a secret key, but this key is not shared with the eavesdropper Eve. Eve sees the symbols on the channel, and understands as well the likelihood by which Alice chooses messages. However, because Eve does not share the secret key, Eve should find no use for these symbols. Her likelihood function should not be refined.

Alice's likelihood is represented as a probability distribution over a message space. The messages space M is assumed finite. A probability distribution $P(M)$ is a map from M to $[0, 1]$, satisfying the axioms of a probability distribution; but might be better to think of $P(M)$ as a map from events in M, that is, subsets of M, to $[0, 1]$. Events are things we can learn about the message, such as "the event that the

Date: January 27, 2018. *Updated:* October 3, 2025.

message contains a vowel". Generally, for every message m, the event "the message is m" is an admissible event, and so there is no difference between $P : M \to [0,1]$ and $P : Pwr(M) \to [0,1]$.

Complete uncertainty on Alice's choice corresponds to the uniform distribution: $P(M = m) = 1/|M|$. In this case, Bob will have no preferred message that he can act on in advance of any symbols. Complete certainty corresponds to,

$$P(M = m) = \begin{cases} 1 & \text{if } m = m^* \\ 0 & \text{else} \end{cases}$$

In this case, each time Alice picks a message, that message is m^*, and it is even unnecessary that she sends symbols. Bob can act in advance on the knowledge that when Alice chooses, she will choose m^*.

If the symbols placed on the channel are from the space C, the ciphertext, we wish that Bob learns from this symbol. Consequently the probability is updated. However, as Eve learns nothing, the probability on M conditioned on C should be unchanged. This is the Shannon definition of Perfect Secrecy:

Definition 1.1. An encryption scheme as *Perfect Secrecy* if for every probability distribution $P(M)$ and for every $c \in C$, the probability distribution $P(M \mid c)$ is the same as the a priori likelihood distribution $P(M)$.

1.1. **A Perfectly Secret Cipher.** The *Vernam Cipher*, or *One Time Pad*, is an example of a perfectly secret cipher. It works on a message space of bits, and the key is a stream of bits matching the length of the message. We discuss the case of a one bit message.

Alice and Bob flip a fair coin (or a coin of bias β). The result $k \in \{0,1\}$ is their secret key. Given the message $m \in \{0,1\}$, Alice forms ciphertext $c = k \oplus m$. Bob receives c and recovers m by,

$$c \oplus k = m \oplus k \oplus k = m \oplus 0 = m.$$

Theorem 1.1. The Vernam Cipher has Prefect Secrecy if and only if $\beta = 1/2$.

Proof. Because $\beta = 1/2$, half of the 0 messages end up transmitting a 0, and half of the 1 messages end up transmitting a 0; so half of the transmissions are 0. Leaving that the other half of the transmissions are a 1.

$$\begin{aligned} P(C = 0) &= P(C = 0 \,|\, M = 0)\, P(M = 0) + P(C = 0 \,|\, M = 1)\, P(M = 1) \\ &= 1/2\,(P(M = 0) + P(M = 1)) = 1/2 \end{aligned}$$

And $P(C = 1) = 1/2$ likewise. Ciphertext c obtained from message m exactly when $k = c \oplus m$,

$$P(C = c \mid M = m) = P(k = c \oplus m) = 1/2.$$

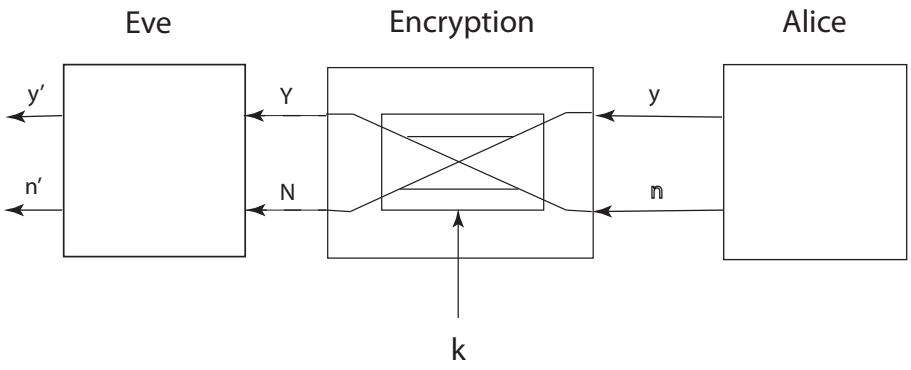

FIGURE 1. Vernam Cipher.

Using Bayes Theorem to put these facts together,

$$\begin{aligned} P(M = m \mid C = c) &= P(C = c \mid M = m)P(M = m)/P(C = c) \\ &= (1/2)P(M = m)/(1/2) \\ &= P(M = m) \end{aligned}$$

satisfying the definition of Perfect Secrecy.

□

Hence, to Eve, without access to k, the likelihood of a message after observation of C is the same as without any observation. The channel reveals nothing because Eve's beliefs can ignore it entirely without loss. However, Bob can update his probability to achieve certainty about which message was chosen by Alice.

1.2. **Odds Ratio and Bias.** Assume, through some fault of key generation that the key coin is not fair, $\beta \neq 1/2$. Then the system does not have perfect secrecy. In the case of a probability distribution that places weight $1/2$ on each of two elements in M, the ciphertext is a clear hint about the message chosen. Into the encryption box, both a 0 or a 1 are equally likely, but out of the encryption box, it is more likely a 0 if the input were a 0, and more likely a 1 if the input were a 1.

We calculate that $P(m = c) = \beta$, in this case. We can do this directly from the equation $k = m \oplus c$.

However, what of perfect secrecy in the case were the message distribution is not uniform. We can assume w.l.o.g. $\beta \geq 1/2$ and $\delta = P(M = 0) \geq P(M = 1)$. What

should Eve do? Let's assume that Eve chooses the most likely input to the encryption box given the observed output. We express this as an odds ratio:

$$P(M = 0 \mid C = c)/P(M = 1 \mid C = c)$$

Intuitively, if $c = 0$, then it is more likely that $m = 0$. $M = 0$ is already the a priori guess, and a biased coin only adds likelihood to this guess if $c = 0$.

So:

$$\begin{aligned}\frac{P(M = 1 \mid C = 1)}{P(M = 0 \mid C = 1)} &= \frac{P(C = 1 \mid M = 1)P(M = 1)/P(C = 1)}{P(C = 1 \mid M = 0)P(M = 0)/P(C = 1)} \\ &= \frac{P(C = 1 \mid M = 1)}{P(C = 1 \mid M = 0)} \frac{P(M = 1)}{P(M = 0)} \\ &= \frac{\beta}{1-\beta} \frac{1-\delta}{\delta}\end{aligned}$$

Eve will guess 1 if,

$$\frac{\beta}{1-\beta}\frac{1-\delta}{\delta} > 1,$$

which reduces to $\beta > \delta$.

Hence Eve continues to guess the more a priori outcome and achieves success probability $\delta = P(M = 0)$ unless the coin bias β rises about δ, in which case Eve guesses that m is c and achieves success probability β.

1.3. **Conditions for Perfect Secrecy.** A few necessary conditions for perfect secrecy are immediate. It must be that the key space is at least the size the space of ciphertext messages. If not, then decrypting a given ciphertext by each key will give the space of possible messages, which must omit some possible messages. Any message missing from the collection of possible decipherings will have a posteri probability 0, and for the purposes of the proof, a distribution on messages can be assumed with a priori probability non-zero for this message.

The space of ciphertext messages must be at least as large as the space of messages, else there will be some two messages encrypting to the same ciphertext, and the requirement that decryption be the inverse of encrypting cannot be achieve.

Shannon's theorem for perfect secrecy assumes equal sizes for the key space, message space, and ciphertext space and gives two conditions necessary and sufficient for perfect secrecy.

(1) The choice if k from K is made uniformly at random; and
(2) For each m in M and c in C there is a unique k in K such that $c = E_k(m)$.

One can check that these theorem holds for the Vernam Cipher described above.

1.4. **Multibit Ciphers.** The exclusive-or over a 0-1 space can be replaced with a randomly chosen shift in the space of modular integers. A possibility is to use the letters A through Z as encodings of 0 through 25 mod 26. Then a rotational shift gives a perfectly secret cipher. However, if the same shift is used for multiple messages, then the cipher is not perfect. A constructive proof of non-secrecy is the use of frequency analysis to recover the key, and break the cipher. The a priori distribution of messages refines the guess of the key with each sample encryption.

If each shift were chosen independently, uniformly at random for each character in a text, then the requirements of perfect secrecy are satisfied. The reuse of a key breaks perfect secrecy.

However, shift ciphers do not have completely random choices for keys for each column. The Vigenere cipher improves a simple shift cipher by having several groups of columns which independently chosen shifts. However, this is not sufficient, let alone that the key is generally not drawn at random, but for convenience is drawn from a more limited space of words in, say, a dictionary. There are 170,000 words in the dictionary, and $26^4 = 456,976$. Hence the key space is notably sparse even for a 4 letter message.

2. Adversarial Indistinguishability

Shannon introduced encryption in the context of entropy and information. The model gives absolute answers. It does not depend on any assumptions about the attacker. The encryption keeps the message secret because the channel symbols support equally any hypothesis about which message is more likely, in the presence of that symbol. However, the price is that the number of bits in a key must be a large as the number of bits in the message. This is not practical.

Computational complexity allows for the possibility of a practical scheme that is effectively secure. A claim can be made that, although possible, it is not within the capacity of a practical computation to extract information from the ciphertext.

To help introduce these notions we rephrase perfect security with an adversary, that first will have unbounded power but then we will ask that the adversary fit within computational bounds.

2.1. **The Adversarial Indistinguishability game.** The Adversarial Indistinguishability Games has two players. An encryption system Π,

$$\Pi = (G, E_k, D_k),$$

and an Adversary $\mathcal{A}$,

$$\mathcal{A} = (\mathcal{A}_m, \mathcal{A}_c).$$

The Adversary is any Probabilistic Polynomial Time (PPT) algorithm. The game chooses one of two messages at random, chooses a random key, and give the encryption of the chosen message to the Adversary. The Adversary tries to guess which of the two messages was the message encrypted.

(1) The security parameter n is announced.
(2) The Adversary $\mathcal{A}_m(n)$ chooses two message, m_0 and m_1 from the message space. The messages must be of equal length, $|m_0| = |m_1|$, and of length bounded by a polynomial in n.
(3) An n bit key $G(n)$ is at random by and a fair coin $b \in \{0,1\}$ is tossed.
(4) The challenge cipher $c = E_k(m_b)$ is calculated.
(5) The Adversary attempts to predict b with a bit $\hat{b} = \mathcal{A}_c(n, c)$.
(6) The game is won by the Adversary if $b == \hat{b}$.

The game is illustrated in Figure 2. This figure shows the data flow between Π and $\mathcal{A}$ and in a top to bottom reading of the flow is the time sequence of events.

Denote by $\langle \Pi, \mathcal{A} \rangle(n)$ ($PrivK^{eav}_{\mathcal{A},\Pi}$ in the textbook's notation) the outcome of the interaction. This is a random variable,

$$\langle \Pi, \mathcal{A} \rangle(n) \quad : \quad \Omega_G \times \Omega_{\mathcal{A}} \times \{ 0, 1 \} \quad \rightarrow \quad \{ T, F \}$$

where Ω_G is the randomness consumed by G to generate the key; $\Omega_{\mathcal{A}}$ is the randomness consumed by the Adversary, and $\{ 0, 1, \}$ is the randomness of the bit b. Without loss of generality, the encryption E_k can be deterministic, as *the key is used only once*. Hence, any random behavior desired of E_k can be included in the key bits.

Definition 2.1. Given an encryption scheme Π, the *success probability* is for an Adversary $\mathcal{A}$ is the probability $Pr(\langle \Pi, \mathcal{A} \rangle(n))$ that the Adversary wins the game. The encryption scheme has *perfect adversarial Indistinguishability* if for all PPT Adversaries, the success probability is no better than chance,

$$\forall \mathcal{A}, \ Pr(\langle \Pi, \mathcal{A} \rangle(n)) = 1/2.$$

Theorem 2.1. An encryption scheme has Perfect Adversarial Indistinguishability if and only if it has Perfect Secrecy.

Lemma 2.1. Perfect Secrecy implies Perfect Adversarial Indistinguishability.

Proof. Assume the systems does not have perfect adversarial indistinguishability. Then by the definition of perfect adversarial indistiguishability, there are two messages m_0 and m_1 for which some adversary $\mathcal{A}$ has an advantage. What the Adversary sees is a sample of a ciphertext-valued random variable,

$$X_i \quad : \quad \Omega_G \times \Omega_{\mathcal{A}} \times \{ 0, 1 \} \quad \rightarrow \quad E_k(m_b),$$

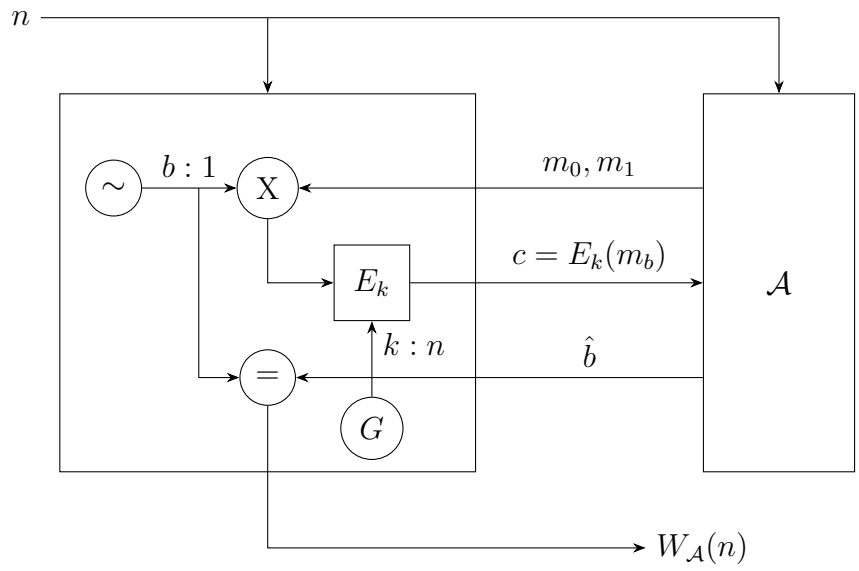

FIGURE 2. Adversarial Indistinguishability Game

Note that,

$$\mathcal{P}(X_i) = \mathcal{P}(C \,|\, M = m_i),$$

This is the measure in the space $\Omega_G \times \Omega_{\mathcal{A}}$ of those keys and randomness that map m_i to a particular $c \in C$. Since there is the possibility of an adversary with an advantage, it must be that these distributions differ.

Hence they must differ for a particular $c^* \in C$,

$$\mathcal{P}(C = c^* \,|\, M = m_0) \neq \mathcal{P}(C = c^* \,|\, M = m_1),$$

Perfect secrecy states that a certain equality holds for all probability distributions on M, so choose the distribution where m_0 and m_1 both have probability $1/2$, and no other message is possible. Applications of Baye's law then gives,

$$\begin{aligned}
\mathcal{P}(M = m_0 \,|\, C = c^*) &= \mathcal{P}(C = c^* \,|\, M = m_0)\mathcal{P}(M = m_0)/\mathcal{P}(C = c^*) \\
&= \mathcal{P}(C = c^* \,|\, M = m_0)\mathcal{P}(M = m_1)/\mathcal{P}(C = c^*) \\
&\neq \mathcal{P}(C = c^* \,|\, M = m_1)\mathcal{P}(M = m_1)/\mathcal{P}(C = c^*) \\
&= \mathcal{P}(M = m_1 \,|\, C = c^*)
\end{aligned}$$

For perfect secrecy both these conditional probabilities must be equal to the a priori probability of the message,

$$\mathcal{P}(M = m_i \,|\, C = c^*) = \mathcal{P}(M = m_i) = 1/2.$$

Which is a contradiction.

□

Lemma 2.2. Perfect Adversarial Indistinguishability implies Perfect Secrecy

Proof. Assume that the scheme does not have Perfect Secrecy. An equivalent definition for Perfect Secrecy is that $\forall m_0, m_1 \in M$ and any $c \in C$ which has non-zero probability under the message distribution $\mathcal{P}(M)$,

$$\mathcal{P}(C = c \mid M = m_0) = \mathcal{P}(C = c \mid M = m_1)$$

because for any m and m' of non-zero probability, the perfect secrecy condition,

$$\mathcal{P}(M = m \mid C = c) = \mathcal{P}(M = m)$$

entails,

$$\mathcal{P}(c \mid m) = \mathcal{P}(m \mid c)\,\mathcal{P}(c)\,/\,\mathcal{P}(m) = \mathcal{P}(c) = \mathcal{P}(c \mid m').$$

Therefore since the scheme is assumed not perfectly secret, there exists m_0, m_1 and c^* such that

$$\mathcal{P}(C = c^* \mid M = m_0) \neq \mathcal{P}(C = c^* \mid M = m_1).$$

Because we are in a uniform model of computation, we assume not only does such a pair of messages exist, but they are discoverable for any n by the Adversary within its computational bounds.

Construct $\mathcal{A}$ as follows. $\mathcal{A}$ offers messages m_0 and m_1. On receiving c, if $c \neq c^*$, $\mathcal{A}$ answers a random bit b'. If $c = c^*$, $\mathcal{A}$ answers 0. Note,

$$\begin{aligned}
\mathcal{P}(C = c^*) &= \mathcal{P}(C = c^* \mid b = 0)\,\mathcal{P}(b = 0) + \mathcal{P}(C = c^* \mid b = 1)\,\mathcal{P}(b = 1) \\
&= 1/2\,(\mathcal{P}(C = c^* \mid b = 0) + \mathcal{P}(C = c^* \mid b = 1)) \\
&= 1/2\,(\mathcal{P}(C = c^* \mid M = m_0) + \mathcal{P}(C = c^* \mid M = m_1)) \\
&\neq \mathcal{P}(C = c^* \mid M = m_0) \\
&= \mathcal{P}(C = c^* \mid b = 0).
\end{aligned}$$

As a consequence,

$$\begin{aligned}
\mathcal{P}(PrivK^{eav}_{\mathcal{A},\Pi} = 1 \mid C = c^*) &= \mathcal{P}(b = 0 \mid C = c^*) \\
&= \mathcal{P}(C = c^* \mid b = 0)\,\mathcal{P}(b = 0)\,/\,\mathcal{P}(C = c^*) \\
&\neq \mathcal{P}(b = 0) \\
&= 1/2
\end{aligned}$$

The advantage when c is not c^* is $1/2$, but not $1/2$ when c is c^*. Averaging over the cases, the average is not $1/2$. Hence the scheme is not perfect adversary indistinguishable.

□

3. Practical Indistinguishability

Having rephrased perfect secrecy in the model of adversarial indistinguishability, we consider now only practical adversaries, and allow them a negligible possibility of success. A practical adversary is modeled as a probabilistic polynomial-time bounded algorithm. Brute force attacks are exponential in key size, and if we have fewer keys than plaintexts, are bound to ultimately succeed. But a sufficient large key will keep these attacks only theoretically successful. There will not be enough computing power to carry out the attack in time to act upon the information.

As an aside, in physical security, a similar approach is used. Safes are not considered unbreakable, but are gauged by the time required to break them. For instance, thicker steel will take more time to cut. Note however, this is an absolute time; whereas we are invoking the models of algorithm complexity, so our time is a function, and by selection of key size the estimated time to break can be increased at will.

The possibility of of key compromise, even though the probability of that is exponentially remote, means that there is intact a small advantage to the attacker. It cannot be denied. However, the practical definition of indistinguishability looses as little ground as possible to this reality by allowing only a negligible advantage.

Definition 3.1. A functionis *negligible* if it $O(1/f(n))$ for any polynomial function $f(n)$.

Although this is not decaying as fast as an exponential function, is is fast enough to allow polynomial repetitions of any attack to still have negligible success rate.

Definition 3.2. An encryption scheme Π has *Adversarial Indistinguishability* if, for any probabilistic polynomial-time adversary $\mathcal{A}$ there is a negligible function $g(n)$ such that for all n,

$$P(\langle \Pi, \mathcal{A} \rangle(n)) \leq 1/2 + g(n).$$

A possible implementation of such an encryption is by analogy to the Vernum cipher, except that instead of using a truly random pad, a pseudorandom pad is used. A pseudorandom pad is a stream, a function of a random seed, that appears random. For instance, the next bit is unpredictable even when the adversary examines the stream up to that point, except for the seed.

The existence of good ciphers by this definition, then, becomes the question of the existence of such pseudorandom generators.