

Reconciling Belief and Vulnerability in Information Flow

Sardaouna Hamadou, Vladimiro Sassone, Catuscia Palamidessi

► **To cite this version:**

Sardaouna Hamadou, Vladimiro Sassone, Catuscia Palamidessi. Reconciling Belief and Vulnerability in Information Flow. 31st IEEE Symposium on Security and Privacy, May 2010, Berkeley/Oakland, California, United States. IEEE Computer Society, pp.79-92, 2010, <10.1109/SP.2010.13>. <inria-00548007>

HAL Id: inria-00548007

<https://hal.inria.fr/inria-00548007>

Submitted on 19 Dec 2010

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Reconciling Belief and Vulnerability in Information Flow

Sardaouna Hamadou and Vladimiro Sassone
School of Electronics and Computer Science (ECS)
University of Southampton
Southampton, UK
{sh3, vs}@ecs.soton.ac.uk

Catuscia Palamidessi
INRIA and LIX
Ecole Polytechnique
Paris, France
catuscia@lix.polytechnique.fr

Abstract—Belief and vulnerability have been proposed recently to quantify information flow in security systems. Both concepts stand as alternatives to the traditional approaches founded on Shannon entropy and mutual information, which were shown to provide inadequate security guarantees. In this paper we unify the two concepts in one model so as to cope with (potentially inaccurate) attackers’ extra knowledge. To this end we propose a new metric based on vulnerability that takes into account the adversary’s beliefs.

Index Terms—Security; information hiding, information flow; quantitative and probabilistic models; uncertainty; accuracy;

I. I

Protecting *sensitive* and *confidential* data is becoming increasingly important in many fields of human activities, such as electronic communication, auction, payment and voting. Many protocols for protecting confidential information have been proposed in the literature. In recent years the frameworks for reasoning, designing, and verifying these protocols have considered probabilistic aspects and techniques for two reasons. First, the data to be protected often range in domains naturally subject to statistical considerations. Second and more important, the protocols often use randomised primitives to obfuscate the link between the information to be protected and the observable outcomes. This is the case, e.g., of the DCNets [8], Crowds [30], Onion Routing [37], and Freenet [13].

From the formal point of view, the *degree of protection* is the converse of the *leakage*, i.e. the amount of information about the secrets that can be deduced from the observables. Early approaches to information hiding in literature were the so-called *possibilistic approaches*, in which the probabilistic aspects were abstracted away and replaced by non-determinism. Some examples of these approaches are those based on *epistemic logic* [19], [36], on *function views* [21], and on *process calculi* [31], [32]. Recently, however, it has been recognised that the possibilistic view is too coarse, in that it tends to consider as equivalent systems which have very different degrees of protection.

The *probabilistic approaches* are therefore becoming increasingly more popular. At the beginning they were investigated mainly at their strongest form of protection, namely to express the property that the observables reveal no (quantitative) information about the secrets (*strong anonymity*, *no*

interference) [2], [8], [19]. More recently, weaker notions of protection have been considered, due to the fact that such strong properties are almost never achievable in practice. Still in the probabilistic framework, Rubin and Reiter have proposed the concepts of *possible innocence* and of *probable innocence* [30] as weak notions of anonymity protection (see also [4] for a generalisation of the latter). These are, however, still true-or-false properties. The need to express in a quantitative way the degree of protection has then lead naturally to explore suitable notions within the well-established fields of *Information Theory* and of *Statistics*.

Concepts from Information Theory [15] have indeed revealed quite useful in this domain. In particular, the notion of noisy channel has been used to model protocols for information-hiding, and the flow of information in programs. The idea is that the input $s \in \mathcal{S}$ of the channel represents the information to be kept secret, and the output $o \in \mathcal{O}$ represents the observable. The noise of the channel is generated by the efforts of the protocol to hide the link between the secrets and the observable, usually by means of randomised mechanisms. Consequently, an input s may generate several different outputs o , according to a conditional probability distribution $p(o|s)$. These probabilities constitute the *channel matrix* C . Similarly, for each output there may be several different corresponding inputs, according to the converse conditional probability $p(s|o)$ which is linked to the above by the Bayes law: $p(s|o) = p(o|s)p(s)/p(o)$. The probability $p(s)$ is the *a priori* probability of s , while $p(s|o)$ is the *a posteriori* probability of s , after we know that the output is o . These probability distributions determine the *entropy* and the *conditional entropy* of the input, respectively. They represent the uncertainty about the input, before and after observing the output. The difference between entropy and conditional entropy is called the *mutual information* and expresses how much information is carried by the channel, i.e. how much uncertainty about the input we lose by observing the output (i.e., equivalently, how much information about the input we gain by observing the output).

Even though several notions of entropy have been proposed in Information Theory, Shannon’s is by far the most famous of them, due to its relation with the *channel’s rate*, i.e., the speed by which information can be transmitted accurately on a channel. Consequently, there have been various attempts to

define the degree of protection by using concepts based on Shannon entropy, notably mutual information [10], [23], [24], [38] and the related notion of capacity, which is the supremum of the mutual information over all possible input distributions, and which therefore represents the worst case from the point of view of security [5], [27], [28].

A refinement of the above approaches came from the ideas of integrating the notions of extra knowledge and belief [14], [18]. The idea is that the gain obtained by looking at the output should be relative to the possible initial knowledge or belief that an attacker may have about the secret. For instance, assume that in a parliament composed by m Labourists and n Conservatives, m members voted against a proposal to remove minimum wages. Without any additional knowledge it is reasonable to believe that all Labourists voted against. If however we came to know that exactly one Conservative voted against, then it is more reasonable to believe that the most liberally-inclined Conservative voted against, and the least liberally-inclined Labourist voted in favour. In this case, the a posteriori belief is likely to be much more accurate than the a priori one, and the gain obtained using the knowledge about MPs' relative positioning on the left-to-right scale is much larger than the one computed as difference of entropies. Consequently, [14] proposes to define the protection of a system in terms of the difference (expressed in terms of Kullback-Leibler divergence) between the accuracy of the a posteriori belief and the accuracy of the a priori one.

In recent work, however, Smith has shown that the concepts based on Shannon entropy are not very suitable for modelling the information leakage in the typical scenario of protocol attacks, where the adversary attempts to guess the value of the secret in *one single* try [33]. He gave an example of two programs whose Shannon's mutual information is about the same, yet the probability of making the right guess after having observed the output is much higher in one program than in the other. In a subsequent paper [34], Smith has proposed to define the leakage in terms of a notion of mutual information based on Rényi *min-entropy*.

Recently in [20] the authors extended the vulnerability model of [34] in the context of the Crowds protocol for anonymous message posting to encompass the frequent situation where attackers have extra knowledge. They pointed out that in Crowds the adversary indeed has extra information (*viz.*, the target servers) and assumed that she knows the correlation between that and the secret (*viz.*, the users' preferences for servers). They proved that in such scenarios anonymity is more difficult to achieve.

In our opinion, a fundamental issue remains wide open: the need to measure and account for the *accuracy* of the adversary extra knowledge. Indeed, [20] assumes that the adversary's extra information is accurate, and this assumption is generally not warranted. Inaccuracy can indeed arise, e.g. from people giving deliberately wrong information, or simply from outdated data. As already noticed in [14] there is no reason in general to assume that the probability distributions

the attacker uses are correct, and therefore they must be treated as *beliefs*.

This paper tries to fill this gap by generalising the model on Rényi min-entropy to cope with the presence of the attacker's beliefs. To this end we propose a new metric based on the concept of vulnerability that takes into account the adversary's beliefs. The idea is that the attacker does not know the *actual* probability distributions (*i.e.*, the a priori distribution of the protocol's hidden input and its correlation with the extra information), and is assuming them. The *belief-vulnerability* is then the expected probability of guessing the value of the hidden input in *one try* given the adversary's belief. Informally, the adversary chooses the value of the secret input which has the maximum a posteriori probability according to her belief. Then the vulnerability of the secret input is expressed in terms of the actual a posteriori probabilities of the adversary's possible choices. We show the strength of our definitions both in terms of their theoretical properties and their utility by applying them to various threat scenarios and comparing the results to the previous approaches. Among its several advantages, our model allows to identify the levels of accuracy for the adversary's beliefs which are compatible with the security of a given program or protocol.

The rest of the paper is organised as follows: in §II we fix some basic notations and recall some fundamental notions of Information Theory; in §III we briefly revise previous approaches to quantitative information follow; §IV delivers our core technical contribution by extending the model on Rényi min entropy to the case of attacker's beliefs and investigating its theoretical properties; in §V we apply our approach to various threat scenarios and compare it to the previous approaches whilst §VI contains our concluding remarks.

II. P

In this section we briefly revise the elements of Information Theory which underpin the work in this paper, and illustrate our conceptual framework.

A. Some notions of information theory

Being in a purely probabilistic setting gives us the ability to use tools from information theory to reason about the uncertainty of a random variable and the inaccuracy of assuming a distribution for a random variable. In particular we are interested to the following notions: *entropy*, *mutual information*, *relative entropy* and *min-entropy*. We refer the reader to [16], [26] for more details.

We use capital letters X, Y to denote discrete random variables and the corresponding small letters x, y and calligraphic letters \mathcal{X}, \mathcal{Y} for their values and set of values respectively. We denote by $p(x), p(y)$ the probability of x and y respectively and by $p(x, y)$ their joint probability.

Let X, Y be random variables. The (*Shannon*) *entropy* $H(X)$ of X is defined as

$$H(X) = - \sum_{x \in \mathcal{X}} p(x) \log p(x). \quad (1)$$

The entropy measures the *uncertainty* of a random variable. It takes its maximum value $\log|\mathcal{X}|$ when X is uniformly distributed and its minimum value 0 when X is a constant. We take the logarithm with a base 2 and thus measure entropy in *bits*. The *conditional entropy*

$$H(X|Y) = - \sum_{y \in \mathcal{Y}} p(y) \sum_{x \in \mathcal{X}} p(x|y) \log p(x|y) \quad (2)$$

measures the amount of uncertainty of X when Y is known. It can be shown that $0 \leq H(X|Y) \leq H(X)$ with the leftmost equality holding when Y completely determines the value of X and the rightmost one when Y reveals no information about X , i.e., X and Y are independent random variables.

Comparing $H(X)$ and $H(X|Y)$ give us the notion of *mutual information*, denoted $I(X; Y)$ and defined by

$$I(X; Y) = H(X) - H(X|Y). \quad (3)$$

It is non-negative, symmetric and bounded by $H(X)$. In other words $0 \leq I(X; Y) = I(Y; X) \leq H(X)$.

The *relative entropy* or *Kullback-Leibler distance* between two probability distribution p and q on the same set \mathcal{X} , denoted $D(p \parallel q)$, is defined as

$$D(p \parallel q) = \sum_{x \in \mathcal{X}} p(x) \log \frac{p(x)}{q(x)}. \quad (4)$$

It is non-negative (but not symmetric) and it is 0 if and only if $p = q$. The relative entropy measures the *inaccuracy* or *information divergence* of assuming that the distribution is q when the true distribution is p .

The *guessing entropy* $G(X)$ is the expected number of tries required to guess the value of X optimally. The optimal strategy is to guess the values of X in decreasing order of probability. Thus if we assume that $\mathcal{X} = \{x_1, x_2, \dots, x_n\}$ and x_i 's are arranged in decreasing order of probabilities, i.e., $p(x_1) \geq p(x_2) \geq \dots \geq p(x_n)$, then

$$G(X) = \sum_{1 \leq i \leq n} ip(x_i). \quad (5)$$

The min-entropy $H_\infty(X)$ of a random variable is given by

$$H_\infty(X) = - \log \max_{x \in \mathcal{X}} p(x) \quad (6)$$

and measures the difficulty for an attacker to correctly guess the value of X in *one* try (obviously using the optimal strategy above). It can be shown that $H_\infty(X) \leq H(X)$ with equality when X is uniformly distributed. In general, $H(X)$ can be arbitrary higher than $H_\infty(X)$, since it can be arbitrary high even if X assumes a given value with probability close to 1.

B. Framework

In this paper we consider a framework similar to the probabilistic approaches to anonymity and information flow used e.g. in [6], [19], [25], and [34]. We restrict ourselves to *total* protocols and programs with one *high level* input A , a random variable over a finite set \mathcal{A} , and one *low level* output (observable) O , a random variable over a finite set \mathcal{O} . We represent a protocol/program by the matrix of the conditional

probabilities $p(o_j|a_i)$, where $p(o_j|a_i)$ is the probability that the low output is o_j given that the high input is a_i . An adversary or eavesdropper can see the output of a protocol, but not its input, and she is interested in deriving the value of the input from the observed output in *one single* try.

In this paper we shall assume that the high input is generated according to an *a priori* probabilistic distribution $p_p(a_i)$ *unknown to the adversary*, as explained in the introduction. We also denote by $p_\beta(a_i)$ the adversary's assumed *a priori* distribution of A . Furthermore, we assume that the attacker has access to the value of a random variable B distributed over \mathcal{B} which summarises the additional knowledge (information) about A she has independently of the behaviour of the protocol. The matrix of the conditional probabilities $p_p(b_k|a_i)$ (resp. $p_\beta(b_k|a_i)$) expresses the *real* (resp. *the adversary's assumed*) correlation between the hidden input and the additional observables \mathcal{B} . An adversary's *belief* then consists of the pair $p_\beta(a_i), p_\beta(b_k|a_i)$ of her assumed probabilities.

When $|\mathcal{B}| = 1$ and the *a priori* distribution is publicly-known, i.e., $p_p(a_i) = p_\beta(a_i)$, the adversary's additional information about A is trivial and cannot help to determine the value of A . For example, knowing the length of a password in a fixed-length password system gives no advantage, as all passwords have the same length. Trivial information allow us to model the absence of additional information, and to see the standard framework in the literature as an instance of ours.

Example 1: Let A be a random variable with a publicly-known uniform a priori distribution over $\mathcal{A} = \{0, 1, 2, 3\}$. Assume that the adversary's additional observable is the parity of A , i.e. $\mathcal{B} = \{0, 1\}$, with the following deterministic belief's correlation $p_\beta(b_k|a_i) = p(a_i \bmod 2 = b_k)$. In other words, the adversary believes that her additional information accurately reflects that the value of A is an even number if $B = 0$ and odd otherwise.

Now suppose that A is the high input of the deterministic program C1 below, whose low output is

$$O = \begin{cases} 1 & \text{if } a \in \{0, 1\} \\ 2 & \text{otherwise.} \end{cases}$$

```

PROG C1:
  BEGIN
    O := [ log(A + 2) ]
  END

```

In the case of wrong belief, i.e., when the attacker believes that the value of A is even (resp. odd) when it actually is odd (resp. even), her low observation of PROG C1 does not allow her to correct her belief. Indeed, both observations can be induced by any number.

Now suppose that A is the high input of the probabilistic program C2 below, with low output $O \in \{-1, 0, 2\}$ and conditional probabilistic matrix as in Table I.

```

PROG C2:
  BEGIN
    R 'sampled from {0,2} with p(0) = λ and p(2) = 1 - λ';
    If A = R
      Then O := A

```

$p(o a)$	o_0	o_1	o_2
a_0	$1 - \lambda$	λ	0
a_1	1	0	0
a_2	λ	0	$1 - \lambda$
a_3	1	0	0

TABLE I

C

PROG C2

Else $O := -1$ END

Contrary to the PROG C1, the low output of PROG C2 may allow the adversary to correct her wrong belief. In particular if $B = 1$ and O is either 0 or 2 then the adversary knows that her belief is wrong. But the observation $O = -1$ cannot help her correct her wrong belief, as it is compatible with both beliefs.

III. U

This section reviews the existing definitions for quantifying information leakage. We begin by quantitative approaches to information flow based on Shannon entropy and mutual information, and recall why they fail to give good security guarantees. We then present an alternative approach based on the adversary's beliefs proposed by Clarkson, Myers and Schneider [14]. We conclude the section by presenting a more recent alternative approach based on the concept of vulnerability introduced by Smith [34].

A. Shannon entropy approach

There seems to be a general consensus in the literature for using Shannon entropy to measure uncertainty and mutual information to quantify information leakage [6], [9], [11], [12], [22]. We remind the reader that these approaches aim at quantifying information flow as a reduction of the adversary uncertainty about the high input and take no account of the adversary's initial knowledge. Shannon entropy $H(A)$ as a measure of the uncertainty of A seems adequate to express the adversary's initial uncertainty about A . Similarly, as the conditional entropy $H(A|O)$ measures the amount of uncertainty of A when O is known, it seems appropriate to express the adversary's remaining uncertainty. We thus have the following definitions.

- *initial uncertainty (IU)*: $H(A)$
- *remaining uncertainty (RU)*: $H(A|O)$
- *information leakage (IL)*: $IU - RU = H(A) - H(A|O) = I(A; O)$

Nevertheless, recent work by Smith [34] suggests that these notions do not support security guarantees satisfactorily. In particular the remaining uncertainty is generally of little value in characterising the real threat that the adversary could guess the value of A given her low observations. Smith uses the following example to prove that.

Example 2: Consider the following programs C3 and C4, where A is a uniformly distributed $8k$ -bit integer, $k \geq 2$, & denotes bitwise 'AND', and $0^{7k-1}1^{k+1}$ a binary constant.

```

PROG C3:
  BEGIN
    If  $A \bmod 8 = 0$ 
      Then  $O := A$ 
      Else  $O := 1$ 
    END

```

```

PROG C4:
  BEGIN
     $O := A \& 0^{7k-1}1^{k+1}$ 
  END

```

PROG C3 reveals completely the high input when A is a multiple of 8 while it reveals nothing about A otherwise (except of course the very fact that it is not a multiple of 8). On the contrary, PROG C4 reveals *always and only* the last $k + 1$ bits of A .

According to the consensus definitions, we have $IU = 8k$, $RU = 7k - 0.169$ and $IL = k + 0.169$ for PROG C3, and $IU = 8k$, $RU = 7k - 1$ and $IL = k + 1$ for PROG C4 (the reader is referred to [34] for the detailed calculations). So, under such definitions, PROG C4 appears actually *worse* than PROG C3, as $7k - 1 < 7k - 0.169$, even though intuitively C3 leaves A highly vulnerable to being guessed (e.g., when it is a multiple of 8) while C4 does not, at least for large k .

B. Belief approach

Recently Clarkson *et al.* [14] showed that the Shannon entropy approach is inadequate for measuring information flow when the adversary makes assumptions about the high-level secret and such assumptions might be incorrect. Based on the conviction that it is unavoidable that the attacker makes such (potentially inaccurate) assumptions, they proposed a new metric. They formalised the idea of an adversary's belief simply as a distribution of A assumed by the adversary: information flow is then expressed as an increase of the *accuracy* of such belief. The *initial accuracy* is the Kullback-Leibler distance between the adversary's initial *belief* and the *actual* distribution of A ; similarly the *remaining accuracy* is the Kullback-Leibler distance between the Bayesian-updated belief of the adversary after her low observation, and the actual distribution of A .

However, as noticed by Smith [34], when the adversary's belief coincides with the a priori distribution of A , then the belief approach reduces again to the inadequate standard approach illustrated above.

C. Vulnerability approach

Having observed that both the consensus and the belief approaches fail in general to give good security guarantees, Smith [34] proposes a new metric for quantitative information flow based on the notions of *vulnerability* and *min-entropy*. We briefly revise these concepts here.

The vulnerability of a random variable A is the worse-case probability that an adversary could guess the value of A correctly in *one try*. The vulnerability of A , denoted $V(A)$, is thus formally defined as follows.

Definition 1: $V(A) = \max_{a \in \mathcal{A}} p(a)$.

The *conditional vulnerability* of a A given O measures the expected probability of guessing A in one try given O . It is denoted $V(A|O)$ and defined as follows.

Definition 2: $V(A|O) = \sum_{o \in \mathcal{O}} p(o)V(A|o)$, where $V(A|o)$ is $\max_{a \in \mathcal{A}} p(a|o)$.

The initial uncertainty about A is then defined as the negative logarithm of $V(A)$, which turns out to be the min-entropy of the random variable A – cf. (6) above. And the remaining uncertainty about A after observing O is defined as the min-entropy of A given O . Thus we have the following vulnerability-based definitions:

- IU : $H_\infty(A) = -\log V(A)$
- RU : $H_\infty(A|O) = -\log V(A|O)$
- IL : $IU - RU = H_\infty(A) - H_\infty(A|O)$

Now on the security guarantees of the vulnerability-based approach. By applying these definitions to the programs of Example 2, we have $IU = 8k$, $RU = 8k - 3$ and $IL = 3$ for PROG C3, and $IU = 8k$, $RU = 7k - 1$ and $IL = k + 1$ for PROG C4. While these quantities remain the same as in the consensus approach for PROG C4, the new metric hugely increases the leakage ascribed to PROG C3 reflecting the fact that the low observations of PROG C3 leave the high input very vulnerable to being guessed.

A related line of research has explored methods of statistical inference, in particular those from the *hypothesis testing* framework. The idea is that the adversary’s best guess is that the true input is the one which has the maximum a posteriori probability (MAP rule) and that, therefore, the *a posteriori vulnerability* of the system is the complement of the *Bayes Risk*, which is the average probability of making the wrong guess when using the MAP rule [7]. This is always at least as high as the *a priori vulnerability*, which is the probability of making the right guess just based on the knowledge of the input distribution. It turns out that Smith’s notion of leakage actually corresponds to the ratio between the a posteriori and the a priori vulnerabilities [3], [34].

Concerning the efficient computation of the channel matrix and the leakage, the only work we are aware of is [1], in which the authors propose various automatic techniques. One of these is able to generate counterexamples, namely points on the execution where the channel exhibits an excessive amount of leakage. This method is therefore also useful to fix unsound protocols.

IV. U B V

We now propose an alternative approach based on the vulnerability concept that takes into account the adversary’s belief.

A. Belief-vulnerability

Let B be the adversary’s additional information about a random high level variable A . Then the *belief-vulnerability* of A is the expected probability of guessing A in one try given the adversary’s belief. Given an additional information $B = b$, the adversary will choose a value having the maximal a posteriori

probability according to her belief, that is a value $a' \in \Gamma_b$, where $\Gamma_b = \operatorname{argmax}_{a \in \mathcal{A}} p_\rho(a|b)$. The vulnerability of A given b is then the *real* probability that the adversary’s choice is correct given b , that is the a posteriori probability $p_\rho(a'|b)$. As there might be many values of A with the maximal belief a posteriori probability, the attacker will pick uniformly at random one element in Γ_b . Hence we have the following definition.

Definition 3: Let A be a random variable and B the adversary’s extra knowledge about A . Then the belief-vulnerability of A , denoted $V(A : B)$, is defined as

$$V(A : B) = \sum_{b \in \mathcal{B}} p_\rho(b)V(A : b) \quad (7)$$

where $V(A : b) = \frac{1}{|\Gamma_b|} \sum_{a' \in \Gamma_b} p_\rho(a'|b)$.

Next, we show how to compute $V(A : B)$ from the given probabilities.

$$\begin{aligned} V(A : B) &= \sum_{b \in \mathcal{B}} p_\rho(b)V(A : b) \\ &= \sum_{b \in \mathcal{B}} p_\rho(b) \left(\frac{1}{|\Gamma_b|} \sum_{a \in \Gamma_b} p_\rho(a|b) \right) \\ &= \sum_{b \in \mathcal{B}} \frac{1}{|\Gamma_b|} \sum_{a \in \Gamma_b} p_\rho(a|b)p_\rho(b) \\ &\quad \text{(by Bayes theorem)} \\ &= \sum_{b \in \mathcal{B}} \frac{1}{|\Gamma_b|} \sum_{a \in \Gamma_b} p_\rho(b|a)p_\rho(a). \end{aligned}$$

Thus the belief-vulnerability can be easily computed as follows.

Proposition 1: Let A be a random variable and B the adversary’s extra knowledge about A . Then

$$V(A : B) = \sum_{b \in \mathcal{B}} \frac{1}{|\Gamma_b|} \sum_{a \in \Gamma_b} p_\rho(b|a)p_\rho(a). \quad (8)$$

We then define the *initial uncertainty* as the min-entropy of $A : B$. Thus we have the following definition.

Definition 4: Let A be a random variable and B the adversary’s additional information about A . Then the initial threat to A given B , denoted $H_\infty(A : B)$, is defined as

$$H_\infty(A : B) = \log \left(\frac{1}{V(A : B)} \right). \quad (9)$$

Example 3: Suppose that A is uniformly distributed over $\{0, 1, 2, 3\}$ and the adversary’s extra information is about the parity of A . Assume that the a priori distribution of A is publicly-known, i.e. $\forall a \in \mathcal{A}, p_\beta(a) = p_\rho(a)$. Assume also that the adversary believes that her extra info is accurate, that is she assumes the following correlation:

$p_\beta(b a)$	b_0	b_1
a_0	1	0
a_1	0	1
a_2	1	0
a_3	0	1

$p_\rho(b a)$	$p_{\rho 1}$	$p_{\rho 2}$	$p_{\rho 3}$	$p_{\rho 4}$
$V(A : B)$	0.49	0.02	0.50	0
$H_\infty(A : B)$	1.03	5.56	1	$+\infty$

$p_{\rho 1}(b a)$	b_0	b_1	$p_{\rho 2}(b a)$	b_0	b_1
a_0	0.98	0.02	a_0	0.03	0.97
a_1	0.02	0.98	a_1	0.98	0.02
a_2	0.98	0.02	a_2	0.04	0.96
a_3	0.02	0.98	a_3	0.94	0.08

$p_{\rho 3}(b a)$	b_0	b_1	$p_{\rho 4}(b a)$	b_0	b_1
a_0	1	0	a_0	0	1
a_1	0	1	a_1	1	0
a_2	1	0	a_2	0	1
a_3	0	1	a_3	1	0

TABLE II

I

Then $\Gamma_0 = \{0, 2\}$ and $\Gamma_1 = \{1, 3\}$. Thus

$$\begin{aligned}
V(A : B) &= \sum_{b \in \mathcal{B}} \frac{1}{|\Gamma_b|} \sum_{a \in \Gamma_b} p_\rho(b|a) p_\rho(a) \\
&= \sum_{b \in \{0,1\}} \frac{1}{|\Gamma_b|} \sum_{a \in \Gamma_b} \frac{1}{4} p_\rho(b|a) \\
&= \frac{1}{4} \left[\frac{1}{2} \sum_{a \in \Gamma_0} p_\rho(b_0|a) \right. \\
&\quad \left. + \frac{1}{2} \sum_{a \in \Gamma_1} p_\rho(b_1|a) \right] \\
&= \frac{1}{8} \left[p_\rho(b_0|a_0) + p_\rho(b_0|a_2) \right. \\
&\quad \left. + p_\rho(b_1|a_1) + p_\rho(b_1|a_3) \right].
\end{aligned}$$

Table II summarizes the initial uncertainty about A when the real correlation between the high level input and the extra observables is $p_{\rho 1}(b|a)$, $p_{\rho 2}(b|a)$, $p_{\rho 3}(b|a)$, and $p_{\rho 4}(b|a)$ respectively. The correlation $p_{\rho 1}(b|a)$ means that the adversary's extra information is slightly noisy: when the high input is an even (resp. odd) number, the extra observable is usually (with probability 0.98) even (resp. odd). But with a small probability the adversary is wrong as the parity is reversed. The contrary holds for the second correlation $p_{\rho 2}(b|a)$, that is the correlation is highly noisy. The third correlation is not a noisy one and coincide with the adversary's assumed one: the adversary's belief is therefore 100% accurate since we assumed that the actual a priori distribution of A is publicly-known. Finally, the last correlation $p_{\rho 4}(b|a)$ always fools the attacker by reverting the parity. Note that in this last case the adversary's initial uncertainty is infinite. This means that it is impossible for her to guess the value of the secret in one try when her initial belief is 100% inaccurate.

More generally, let $\text{Beliefs}(A, B)$ denote the set of adversary's beliefs about A according to the extra information B , and let $\text{Beliefs}_\perp(A, B)$ be the set of totally inaccurate beliefs,

$$\begin{aligned}
\text{Beliefs}_\perp(A, B) &= \left\{ (p_\beta(a), p_\beta(b|a)) \in \text{Beliefs}(A, B) \mid \right. \\
&\quad \left. \forall b \in \mathcal{B}, a \in \Gamma_b \text{ implies } p_\beta(a|b) = 0 \right\}
\end{aligned}$$

Then the following result holds.

Proposition 2: Let A be a random variable and B the adversary's extra information about A . Let $(p_\beta(a), p_\beta(b|a))$ be the adversary's belief, then

$$(p_\beta(a), p_\beta(b|a)) \in \text{Beliefs}_\perp(A, B) \text{ implies } H_\infty(A : B) = +\infty.$$

Proof: Direct consequence of Proposition 1 and Definition 4. \blacksquare

In order to avoid such infinite values when computing the reduction of uncertainty, we shall exclude 100% always inaccurate beliefs. Thus we define *admissible beliefs* up to a positive number ϵ so that we can approximate 100% always inaccurate beliefs by making ϵ as close to zero as possible.

Definition 5: An adversary's initial belief is ϵ -admissible ($0 < \epsilon \leq 1$) if the following holds.

$$\forall b \in \mathcal{B}, a \in \Gamma_b \text{ implies } p_\rho(a|b) \geq \epsilon.$$

$\text{Beliefs}_\epsilon(A, B)$ denotes the set of ϵ -admissible beliefs.

Note that in the above definition, ϵ is a lower bound on the probability that the adversary's guess is correct. Note also that if a belief is ϵ -admissible then it is also ϵ' -admissible for all $\epsilon' \leq \epsilon$. Thus $\text{Beliefs}_\epsilon(A, B) \subseteq \text{Beliefs}_{\epsilon'}(A, B)$.

Next we show that, contrary to information, belief may actually hurt. Indeed, in the above example since A is uniformly (and publicly) distributed over $\{0, 1, 2, 3\}$ then $H_\infty(A) = 2$. Hence $H_\infty(A) > H_\infty(A : B)$ when $\rho \in \{\rho_1, \rho_3\}$ and $H_\infty(A) < H_\infty(A : B)$ when $\rho \in \{\rho_2, \rho_4\}$. In particular, the following two results hold.

Lemma 1: If $\forall b \in \mathcal{B}, a \in \Gamma_b$ implies $p_\rho(b|a) \leq \frac{1}{|\mathcal{B}|}$ then

$$H_\infty(A) \leq H_\infty(A : B)$$

Proof:

$$\begin{aligned}
V(A : B) &= \sum_{b \in \mathcal{B}} \frac{1}{|\Gamma_b|} \sum_{a \in \Gamma_b} p_\rho(b|a) p_\rho(a) \\
&\leq \sum_{b \in \mathcal{B}} \frac{1}{|\Gamma_b|} \sum_{a \in \Gamma_b} \frac{1}{|\mathcal{B}|} p_\rho(a) \\
&\leq \sum_{b \in \mathcal{B}} \frac{1}{|\Gamma_b|} \sum_{a \in \Gamma_b} \frac{1}{|\mathcal{B}|} \max_{a \in \mathcal{A}} p_\rho(a) \\
&\leq \sum_{b \in \mathcal{B}} \frac{1}{|\Gamma_b|} \frac{|\Gamma_b|}{|\mathcal{B}|} \max_{a \in \mathcal{A}} p_\rho(a) \\
&\leq \frac{1}{|\mathcal{B}|} \sum_{b \in \mathcal{B}} V(A) \\
&\leq \frac{|\mathcal{B}| V(A)}{|\mathcal{B}|} = V(A).
\end{aligned}$$

Hence $H_\infty(A) \leq H_\infty(A : B)$. \blacksquare

The next result states that a 100% accurate belief is informative and hence may only reduce the uncertainty about A .

Lemma 2: If $\forall b \in \mathcal{B}, a \in \Gamma_b$ implies that $p_\rho(b|a) = \max_{a' \in \mathcal{A}} p_\rho(a'|b)$, then

$$H_\infty(A : B) = H_\infty(A|B) \leq H_\infty(A).$$

Proof:

$$\begin{aligned}
V(A : B) &= \sum_{b \in \mathcal{B}} p_\rho(b) V(A : b) \\
&= \sum_{b \in \mathcal{B}} p_\rho(b) \left(\frac{1}{|\Gamma_b|} \sum_{a \in \Gamma_b} p_\rho(a|b) \right) \\
&= \sum_{b \in \mathcal{B}} p_\rho(b) \left(\frac{1}{|\Gamma_b|} \sum_{a \in \Gamma_b} \max_{a' \in \mathcal{A}} p_\rho(a'|b) \right) \\
&= \sum_{b \in \mathcal{B}} p_\rho(b) \left(\frac{1}{|\Gamma_b|} \sum_{a \in \Gamma_b} V(A|b) \right) \\
&= \sum_{b \in \mathcal{B}} p_\rho(b) \left(\frac{1}{|\Gamma_b|} |\Gamma_b| V(A|b) \right) \\
&= \sum_{b \in \mathcal{B}} p_\rho(b) V(A|b) \\
&= V(A|B)
\end{aligned}$$

Hence $H_\infty(A : B) = H_\infty(A|B) \leq H_\infty(A)$. \blacksquare

We conclude this subsection by establishing both a lower and an upper bounds of our initial uncertainty in term of min-entropy. The following auxiliary definitions and results serve this purpose.

Definition 6: An adversary's initial belief is (at least) ω -accurate ($0 < \omega \leq 1$), denoted $(p_{\beta^\omega}(a), p_{\beta^\omega}(b|a))$ if the following holds.

$$\forall b \in \mathcal{B}, \quad \frac{1}{|\Gamma_b|} \sum_{a \in \Gamma_b} p_\rho(a|b) \geq \omega \cdot V(A|b).$$

In other words, an adversary's belief is ω -accurate if the belief-vulnerability of A is never off by more than a factor ω from the real vulnerability of A given the additional information. Similarly, we say that an adversary's belief is (ω) -accurate if it is *exactly* ω -accurate, i.e.,

$$\forall b \in \mathcal{B}, \quad \frac{1}{|\Gamma_b|} \sum_{a \in \Gamma_b} p_\rho(a|b) = \omega \cdot V(A|b).$$

Fixed the actual probabilities $p_\rho(a)$ and $p_\rho(b|a)$, let us consider the partial order \ll on $\text{Beliefs}(A, B)$, such that $(p_\beta(a), p_\beta(b|a)) \ll (p_{\beta'}(a), p_{\beta'}(b|a))$ if and only if

$$\forall b \in \mathcal{B}, \quad \frac{1}{|\Gamma_b|} \sum_{a \in \Gamma_b} p_\rho(a|b) \leq \frac{1}{|\Gamma'_b|} \sum_{a \in \Gamma'_b} p_\rho(a|b).$$

Then the following lemma states that $(p_{\beta^1}(a), p_{\beta^1}(b|a))$ is an upper bound on $\text{Beliefs}(A, B)$.

Lemma 3: For all $(p_\beta(a), p_\beta(b|a))$ in $\text{Beliefs}(A, B)$ we have $(p_\beta(a), p_\beta(b|a)) \ll (p_{\beta^1}(a), p_{\beta^1}(b|a))$.

Proof: Follows easily from the definitions of \ll and $(p_{\beta^1}(a), p_{\beta^1}(b|a))$. \blacksquare

We now show that the uncertainty based on belief-vulnerability decreases when the accuracy of the adversary's belief increase.

Lemma 4: $(p_\beta(a), p_\beta(b|a)) \ll (p_{\beta'}(a), p_{\beta'}(b|a))$ implies $H_\infty(A : B) \geq H_\infty(A : B')$.

Proof:

$$\begin{aligned}
(p_\beta(a), p_\beta(b|a)) &\ll (p_{\beta'}(a), p_{\beta'}(b|a)) \\
&\Downarrow \\
\forall b \in \mathcal{B}, \quad \frac{1}{|\Gamma_b|} \sum_{a \in \Gamma_b} p_\rho(a|b) &\leq \frac{1}{|\Gamma'_b|} \sum_{a \in \Gamma'_b} p_\rho(a|b) \\
&\Downarrow \\
\forall b \in \mathcal{B}, \quad p_\rho(b) \frac{1}{|\Gamma_b|} \sum_{a \in \Gamma_b} p_\rho(a|b) &\leq p_\rho(b) \frac{1}{|\Gamma'_b|} \sum_{a \in \Gamma'_b} p_\rho(a|b) \\
&\Downarrow \\
\sum_{b \in \mathcal{B}} p_\rho(b) \frac{1}{|\Gamma_b|} \sum_{a \in \Gamma_b} p_\rho(a|b) &\leq \sum_{b \in \mathcal{B}} p_\rho(b) \frac{1}{|\Gamma'_b|} \sum_{a \in \Gamma'_b} p_\rho(a|b) \\
&\Downarrow \\
V(A : B) &\leq V(A : B').
\end{aligned}$$

Hence $H_\infty(A : B) \geq H_\infty(A : B')$. \blacksquare

Now we show that a 1-accurate belief is also 100% accurate.

Lemma 5: An adversary's belief is 1-accurate if and only if

$$\forall b \in \mathcal{B}, \quad \forall a \in \Gamma_b, \quad p_\rho(a|b) = \max_{a' \in \mathcal{A}} p_\rho(a'|b).$$

Proof: An adversary's belief is 1-accurate if and only if

$$\forall b \in \mathcal{B}, \quad \frac{1}{|\Gamma_b|} \sum_{a \in \Gamma_b} p_\rho(a|b) \geq 1 \cdot \max_{a' \in \mathcal{A}} p_\rho(a'|b). \quad (10)$$

But

$$\frac{1}{|\Gamma_b|} \sum_{a \in \Gamma_b} p_\rho(a|b) \leq 1 \cdot \max_{a' \in \mathcal{A}} p_\rho(a'|b)$$

since

$$\forall b \in \mathcal{B}, \quad \forall a \in \Gamma_b, \quad p_\rho(a|b) \leq \max_{a' \in \mathcal{A}} p_\rho(a'|b).$$

Therefore (10) is equivalent to

$$\forall b \in \mathcal{B}, \quad \frac{1}{|\Gamma_b|} \sum_{a \in \Gamma_b} p_\rho(a|b) = 1 \cdot \max_{a' \in \mathcal{A}} p_\rho(a'|b) \quad (11)$$

which itself is equivalent to

$$\forall b \in \mathcal{B}, \quad \forall a \in \Gamma_b, \quad p_\rho(a|b) = \max_{a' \in \mathcal{A}} p_\rho(a'|b).$$

Indeed, if there exists a_i in Γ_b such that $p_\rho(a_i|b) < \max_{a' \in \mathcal{A}} p_\rho(a'|b)$ there is a_j in Γ_b such that $p_\rho(a_j|b) > \max_{a' \in \mathcal{A}} p_\rho(a'|b)$ in order for (11) to holds; but this is impossible. \blacksquare

The next result shows that an ω -accurate belief impacts the vulnerability of A in presence of extra information by a factor at least ω .

Lemma 6: If the adversary's initial belief is ω -accurate then

$$H_\infty(A : B) \leq \omega \cdot H_\infty(A|B)$$

Proof: By definition we have:

$$\begin{aligned}
V(A : B) &= \sum_{b \in \mathcal{B}} p_\rho(b) V(A : b) \\
&= \sum_{b \in \mathcal{B}} p_\rho(b) \left(\frac{1}{|\Gamma_b|} \sum_{a \in \Gamma_b} p_\rho(a | b) \right) \\
&\quad (\text{by hypotheses}) \\
&\geq \sum_{b \in \mathcal{B}} p_\rho(b) (\omega \cdot V(A | b)) \\
&\geq \omega \cdot \sum_{b \in \mathcal{B}} p_\rho(b) V(A | b) \\
&\geq \omega \cdot V(A | b).
\end{aligned}$$

We can then establish our bounds on the initial uncertainty as follows.

Theorem 1: Let A be a random variable and B be the additional information about A . Then

$$H_\infty(A | B) \leq H_\infty(A : B) \leq H_\infty(A | B) + \log\left(\frac{1}{\zeta}\right),$$

where $\zeta = \min_{b \in \mathcal{B}} \left(\frac{1}{|\Gamma_b|} \sum_{a \in \Gamma_b} p_\rho(a | b) \right) / V(A | b)$.

Note that ζ in the above theorem is strictly greater than zero since we consider admissible beliefs for some positive ϵ . Hence the upper bound is well defined.

Proof: From Lemma 10, Lemma 2, Lemma 3 and Lemma 4 we have

$$H_\infty(A | B) \leq H_\infty(A : B).$$

The second part of the inequality follows from Lemma 6 and from the fact that the adversary's belief is ζ -accurate. ■

Finally, we show that when A is uniformly distributed, we can obtain a better upper bound. We begin by recalling a result proven in [35].

Lemma 7: If A is uniformly distributed and the program is deterministic then $H_\infty(A | O) = \log(|\mathcal{A}| / |O|)$.

Thus we have the following corollary of Theorem 1.

Corollary 1: If A is uniformly distributed and the actual correlation $p_\rho(b | a)$ is deterministic then

$$\log\left(\frac{|\mathcal{A}|}{|\mathcal{B}|}\right) \leq H_\infty(A : B) \leq \log\left(\frac{|\mathcal{A}|}{|\mathcal{B}|}\right) + \log\left(\frac{1}{\zeta}\right),$$

where ζ is defined as in the above theorem.

B. A posteriori belief-vulnerability

We now define our belief-vulnerability conditioned to the low observations of the adversary. Note that in this case, the adversary's low observations could help her sort out inaccurate beliefs. Indeed, if an observation o contradicts her initial belief about extra information b , that is there is no high input a in Γ_b such that $p_\rho(o | a) > 0$, then to try values in Γ_b is pointless. A belief b is *compatible* (from the adversary's point of view) to an observation o , denoted $b \diamond o$, if there exists a in Γ_b such that $p_\rho(o | a) > 0$. For instance, if the adversary initially believes that A is an odd number while observing a low output o of the program which is only possible for even numbers high inputs,

then her belief and her observation are incompatible. Let o and b be the adversary's observation and initial belief respectively. She will then only try values a in Γ_b for which $p_\rho(o | a) > 0$ if her belief and observation are compatible. Otherwise, as the evidence contradicts her belief, she will throw it away and only use the observation.

Now let $\Gamma_{b,o}$ denote the set possible adversary's choices according to both her belief and her low observation. Then

$$\Gamma_{b,o} = \begin{cases} \operatorname{argmax}_{a \in \mathcal{A}} p_\beta(a | b, o) & \text{if } b \diamond o, \\ \operatorname{argmax}_{a \in \mathcal{A}} p_\rho(a | o) & \text{otherwise.} \end{cases}$$

Then we define the a posteriori belief-vulnerability as follows.

Definition 7: Let A be the high input of a program, O its low output and B the adversary's initial belief about A . Then the belief-vulnerability of A given O , denoted $V(A | O : B)$, is defined as

$$V(A | O : B) = \sum_{o \in O} \sum_{b \in \mathcal{B}} p_\rho(b, o) V(A | o : b), \quad (12)$$

where $V(A | o : b) = \frac{1}{|\Gamma_{b,o}|} \sum_{a \in \Gamma_{b,o}} p_\rho(a | o, b)$.

We now show how to compute $V(A | O : B)$ under the assumption that the extra information B and the low observable O are actually *independent*.

$$\begin{aligned}
V(A | O : B) &= \sum_{o \in O} p_\rho(b, o) V(A | o : b) \\
&= \sum_{o \in O} \sum_{b \in \mathcal{B}} p_\rho(b, o) \left(\frac{1}{|\Gamma_{b,o}|} \sum_{a \in \Gamma_{b,o}} p_\rho(a | o, b) \right) \\
&= \sum_{o \in O} \sum_{b \in \mathcal{B}} \frac{1}{|\Gamma_{b,o}|} \sum_{a \in \Gamma_{b,o}} p_\rho(a | o, b) p_\rho(b, o) \\
&= \sum_{o \in O} \sum_{b \in \mathcal{B}} \frac{1}{|\Gamma_{b,o}|} \sum_{a \in \Gamma_{b,o}} p_\rho(o, b | a) p_\rho(a) \\
&= \sum_{o \in O} \sum_{b \in \mathcal{B}} \frac{1}{|\Gamma_{b,o}|} \sum_{a \in \Gamma_{b,o}} p_\rho(b | a) p_\rho(o | a) p_\rho(a)
\end{aligned}$$

Thus we have the following proposition.

Proposition 3: Let A be the high input of a program, O its low output and B the adversary's extra information about A . If O and B are independent, then

$$V(A | O : B) = \sum_{o \in O} \sum_{b \in \mathcal{B}} \frac{1}{|\Gamma_{b,o}|} \sum_{a \in \Gamma_{b,o}} p_\rho(b | a) p_\rho(o | a) p_\rho(a).$$

We then define our remaining uncertainty as follows.

Definition 8: Let A be the high input of a program, O its low output and B the adversary's initial belief about A . The remaining uncertainty about A after observing O , denoted $H_\infty(A | O : B)$, is defined as

$$H_\infty(A | O : B) = \log\left(\frac{1}{V(A | O : B)}\right)$$

Example 4: Consider the following program, with A and B as in Example 3, then $O = \{0, 1, 2\}$.

BEGIN

$O := \lfloor \log(A + 1) \rfloor$;

END

Both the program and the adversary's assumed correlation are deterministic, it is therefore easy to compute the adversary's belief conditional matrix $p_\beta(a|o, b)$ and the associated possible choices $\Gamma_{o,b}$. The result is shown in Table III. Thus under the assumption that B and O are actually independent we have

$$\begin{aligned}
V(A|O : B) &= \sum_{b,o} \frac{1}{|\Gamma_{b,o}|} \sum_{a \in \Gamma_{b,o}} p_\rho(b|a) p_\rho(o|a) p_\rho(a) \\
&= \frac{1}{4} \sum_{b,o} \frac{1}{|\Gamma_{b,o}|} \sum_{a \in \Gamma_{b,o}} p_\rho(b|a) p_\rho(o|a) \\
&= \frac{1}{4} [p_\rho(b_0|a_0) p_\rho(o_0|a_0) \\
&\quad + p_\rho(b_0|a_2) p_\rho(o_1|a_2) \\
&\quad + p_\rho(b_0|a_3) p_\rho(o_2|a_3) \\
&\quad + p_\rho(b_1|a_0) p_\rho(o_0|a_0) \\
&\quad + p_\rho(b_1|a_1) p_\rho(o_1|a_1) \\
&\quad + p_\rho(b_1|a_3) p_\rho(o_2|a_3)] \\
&= \frac{1}{4} [p_\rho(b_0|a_0) + p_\rho(b_0|a_2) \\
&\quad + p_\rho(b_0|a_3) + p_\rho(b_1|a_0) \\
&\quad + p_\rho(b_1|a_1) + p_\rho(b_1|a_3)] \\
&= \frac{1}{4} [(p_\rho(b_0|a_0) + p_\rho(b_1|a_0)) \\
&\quad + (p_\rho(b_0|a_3) + p_\rho(b_1|a_3)) \\
&\quad + p_\rho(b_0|a_2) + p_\rho(b_1|a_1)] \\
&= \frac{1}{4} [2 + p_\rho(b_0|a_2) + p_\rho(b_1|a_1)].
\end{aligned}$$

Hence $V(A|O : B) \geq \frac{1}{2}$, meaning that the remaining uncertainty $H_\infty(A|O : B)$ is always less than or equal to 1 regardless the actual correlation between A and B . Thus PROG C5 leaves the high value very vulnerable to be guessed. Recall that the initial uncertainty $H_\infty(A : B)$ (see Example 3) can be arbitrary high when the accuracy of the adversary's belief is very low. This implies that a deliberate reverting of the parity of the high value in order to confuse the adversary is of very little use when the adversary can see the output of PROG C5. Indeed, even if her initial belief is wrong, the observation allows her to correct it. Table IV summarises the remaining uncertainty for PROG C5 when the actual correlation is ρ_1, ρ_2 and ρ_3 .

We now establish both a lower bound and upper bound to our remaining uncertainty. To this end, we establish some auxiliary results.

We first extend the notion of initial belief's accuracy and the partial order \ll to the adversary's post-beliefs as follows. An adversary's post-belief is (at least) ω -accurate if

$$\forall o \in O, \forall b \in \mathcal{B}, \frac{1}{|\Gamma_{b,o}|} \sum_{a \in \Gamma_{b,o}} p_\rho(a|o, b) \geq \omega \cdot V(A|o, b).$$

Let \ll_O denote the partial order on Beliefs induced by the observations O such that $(p_\beta(a), p_\beta(b|a)) \ll_O (p'_\beta(a), p'_\beta(b|a))$

$p_\beta(a o, b)$	a_0	a_1	a_2	a_3	Γ_{b_k, o_j}
b_0, o_0	1	0	0	0	$\{a_0\}$
b_0, o_1	0	0	1	0	$\{a_2\}$
b_0, o_2	0	0	0	1	$\{a_3\}$
b_1, o_0	1	0	0	0	$\{a_0\}$
b_1, o_1	0	1	0	0	$\{a_1\}$
b_1, o_2	0	0	0	1	$\{a_3\}$

TABLE III

T

$p_\rho(b a)$	$p_{\rho 1}$	$p_{\rho 2}$	$p_{\rho 3}$
$V(A : B)$	0.99	0.515	1
$H_\infty(A : B)$	0.0145	0.957	0

TABLE IV

R

if and only if $\forall o \in O, \forall b \in \mathcal{B}$,

$$\frac{1}{|\Gamma_{b,o}|} \sum_{a \in \Gamma_{b,o}} p_\rho(a|o, b) \leq \frac{1}{|\Gamma'_{b,o}|} \sum_{a \in \Gamma'_{b,o}} p_\rho(a|o, b).$$

As in the previous subsection, we can show that a 1-accurate post-belief is an information and that the remaining uncertainty is a decreasing function of the accuracy of the adversary post-belief.

Lemma 8: Let A be the high input of a program, O its low output and B be an additional information about A . If the adversary's post-belief is 1-accurate then

$$H_\infty(A|O : B) = H_\infty(A|B, O).$$

Lemma 9: $(p_\beta(a), p_\beta(b|a)) \ll_O (p'_\beta(a), p'_\beta(b|a))$ implies $H_\infty(A|O : B) \geq H_\infty(A|O : B')$.

We then establish the following bounds for the belief-vulnerability based remaining uncertainty.

Theorem 2: Let A be a random variable, B be the additional information about A and O be the low output of the program. Then

$$H_\infty(A|O, B) \leq H_\infty(A|O : B) \leq H_\infty(A|O, B) + \log\left(\frac{1}{\eta}\right),$$

where $\eta = \min_{o \in O, b \in \mathcal{B}} \left(\frac{1}{|\Gamma_{o,b}|} \sum_{a \in \Gamma_{o,b}} p_\rho(a|o, b) \right) / V(A|o, b)$.

Finally we have the following corollary of Theorem 2.

Corollary 2: If A is uniformly distributed and both the protocol and the actual correlation between A and B are deterministic then

$$\log\left(\frac{|\mathcal{A}|}{|\mathcal{O}| \cdot |\mathcal{B}|}\right) \leq H_\infty(A|O : B) \leq \log\left(\frac{|\mathcal{A}|}{|\mathcal{O}| \cdot |\mathcal{B}|}\right) + \log\left(\frac{1}{\eta}\right).$$

We conclude this section by showing that in case of belief's absence (i.e., the initial knowledge of the adversary is reduce to trivial information) then our definitions are equivalent to the vulnerability-based definitions of Smith.

Theorem 3: The following statements are equivalent.

- 1) $|\mathcal{B}| = 1$ and the a priori probability of A is publicly-known.
- 2) For each adversary's initial belief $(p_\rho(a), p_\rho(b|a))$ in $\text{Beliefs}(A, B)$ and for each program $p_\rho(o|a)$ we have $V(A : B) = V(A)$ and $V(A|O : B) = V(A|O)$.
- 3) For each adversary's initial belief $(p_\rho(a), p_\rho(b|a))$ in $\text{Beliefs}(A, B)$ and for each program $p_\rho(o|a)$ we have $H_\infty(A|B) = H_\infty(A)$ and $H_\infty(A|O : B) = H_\infty(A|O)$.

Proof: (1) \Rightarrow (2): $|\mathcal{B}| = 1$ implies that B is a constant. Hence B is independent of both A and O . Furthermore, the only possible adversary's belief, which is the publicly-known a priori distribution of A , is 1-accurate. Thus $V(A : B) = V(A|B)$. But $V(A|B) = V(A)$ since A and B are independent. Similarly, $V(A|O : B) = V(A|O)$.

(2) \Rightarrow (1): (By contradiction). Assume that (2) holds and $|\mathcal{B}| > 1$ or the adversary does not know the a priori distribution of A . Let first consider the case $|\mathcal{B}| > 1$. Then we can create an adversary's belief which is *exactly* ω -accurate for any $0 < \omega < 1$. Thus for such belief we have $V(A : B) = \omega \cdot V(A|B)$. Therefore if we choose $\omega \neq V(A)/V(A|B)$, then $V(A : B) \neq V(A)$. Hence, it contradicts our initial hypotheses that (2) holds.

Now assume that $|\mathcal{B}| = 1$ but the adversary does not know the a priori distribution of A . Again B is a constant and thus irrelevant. If A is not uniformly distributed then it is easy to construct an adversary's assumed a priori distribution of A such that $V(A : B) \neq V(A|B) = V(A)$. If however A is uniformly distributed then we can still create a program such that $V(A|O : B) \neq V(A|O, B) = V(A|O)$. Again, this contradicts our initial hypotheses that (2) holds.

Finally, the equivalence (2) \Leftrightarrow (3) follows because functions $g(x) = -\log(x)$ and $g'(x) = 2^{-x}$ are strictly monotone. \blacksquare

V.	O	A	A	B	-
			A		

The previous section establishes the reasonableness of our definitions in terms of their theoretical properties. Now we show the utility of our approach by applying it to various threat scenarios and comparing the results to the previous approaches.

We proceed now to the analysis of the programs presented in this paper, and compare the results with previous approaches. Each of the programs is analysed under the following hypothesis.

- The high input A is uniform and publicly-known. Thus

$$\forall a \in \mathcal{A} \quad p_\rho(a) = p_\beta(a) = \frac{1}{|\mathcal{A}|}.$$

- The adversary believes that her extra info is accurate, that is she assumes the correlation shown in Table V. Thus $\Gamma_0 = \{0, 2\}$ and $\Gamma_1 = \{1, 3\}$.
- The real correlation between A and B is of the form of the matrix shown in Table V. It is easy to see that the adversary's initial belief is therefore ω -accurate.
- B and O are independent.

$p_\beta(b a)$	b_0	b_1	$p_\rho(b a)$	b_0	b_1
a_0	1	0	a_0	$\frac{\omega}{1+\omega}$	$\frac{1}{1+\omega}$
a_1	0	1	a_1	$\frac{1}{1+\omega}$	$\frac{\omega}{1+\omega}$
a_2	1	0	a_2	$\frac{\omega}{1+\omega}$	$\frac{1}{1+\omega}$
a_3	0	1	a_3	$\frac{1}{1+\omega}$	$\frac{\omega}{1+\omega}$

TABLE V
C

We denote by IU_x the initial uncertainty computed using approach $x \in \{c, v, bv\}$ where c, v and bv denote the consensus, vulnerability and belief-vulnerability approaches respectively. Ditto for RU_x and IL_x .

We begin by PROG C1 of Example 1. Since A is uniformly distributed then $IU_c = IU_v = \log|\mathcal{A}| = 2$. Furthermore, $RU_v = \log(|\mathcal{A}|/|O|) = \log\frac{4}{2} = 1 = RU_c$ since PROG C1 is deterministic. Thus, when we do not take into account the attacker's belief, then $IL_c = IL_v = 1$. Now let consider the uniformly ω -accurate attacker's belief. Then from the calculation in Example 3 we have $IU_{bv} = -\log(\frac{\omega}{2(1+\omega)})$. And from Proposition 3, we get $RU_{bv} = -\log(\frac{\omega}{1+\omega})$. Therefore for all ω , $IL_{bv} = 1$. Thus, the adversary's initial knowledge about the parity of A does not affect the quantity of information leaked by PROG C1. However, the real question is not how much information is leaked by this program, but what the remaining uncertainty represents in term of security threat to




Fig. 1. Information flow of PROG C1




Fig. 2. Information flow of PROG C5

the high input. Even though the adversary's belief does not affect the quantity of information leaked, it dramatically affects both the initial and remaining uncertainty. Indeed, as illustrated by Figure 1, both IU_{bv} and RU_{bv} tend toward infinity as ω tends toward zero. On the other hand, IU_{bv} and RU_{bv} tend toward two and one, respectively, as ω tends toward one. In other words inaccurate beliefs strengthen the security of the program (by confusing the adversary), whilst accurate beliefs may weaken it. Thus, a deliberate randomization of the parity of the high input in order to confuse the adversary is a good strategy to strengthen the security of this program.

We continue our analysis with PROG C5 of Example 4 which is a slight modification of PROG C1. Again $IU_c = IU_v = \log|\mathcal{A}| = 2$ and $IU_{bv} = -\log(\frac{\omega}{2(1+\omega)})$. For the remaining uncertainty we have $RU_c = 0.585$, $RU_v = 0.415$ and $RU_{bv} = -\log(\frac{2\omega+1}{2(1+\omega)})$. Therefore, $IL_c = 1.415$, $IL_v = 1.585$ and $IL_{bv} = \log(\frac{2\omega+1}{\omega})$. The information flow ascribed by our approach to this program is illustrated by Figure 2. Unlike PROG C1, the information leakage of this program can be arbitrary high when the inaccuracy of the adversary's belief is high whilst its remaining uncertainty RU_{bv} remains very low even for inaccurate beliefs. As already noticed in Example 4, this program leaves A highly vulnerable of being guessed and a deliberate padding of A in order to confuse the adversary is of little help. Note however that RU_{bv} tends toward one, which is higher than both RU_c and RU_v , as ω tends toward zero. It means that highly inaccurate beliefs slightly strengthen the security of PROG C5.

We proceed with the probabilistic program PROG C2 (see Example 1). Again $IU_c = IU_v = \log|\mathcal{A}| = 2$ and $IU_{bv} = -\log(\frac{\omega}{2(1+\omega)})$. For the remaining uncertainty, we have $RU_v = 1$, $RU_c = \frac{1}{4}(3 \log 3 - \log[(1-\lambda)^{(1-\lambda)}\lambda^\lambda])$ and $RU_{bv} = -\log(\frac{1}{4} + \frac{\omega}{\omega+1}[\frac{1}{4} + \frac{1}{4} \max(\lambda, 1-\lambda)])$. The information flow ascribed by the consensus definitions is illustrated in Figure 3 and those of the belief-vulnerability approach in Figures 4 and 5. We first note that in the case of belief's absence, our approach – which coincides with the vulnerability one– ascribes the same information flow quantities to both PROG C2 and PROG C1, even though they seem to present rather different threats.¹ The reason is simply that after her low observation, the adversary needs on average 2^1 tries to guess the value of A for both programs. We also note that the randomisation parameter λ of PROG C2 has no effect on the vulnerability approach, and has only a little one on ours when the accuracy of the adversary's beliefs tends to 1 and λ to $\frac{1}{2}$. This is due to the fact that these metrics focus on the single probability that brings greatest risk and values 0 and 2 of A play symmetric roles with respect to λ . Finally, comparing Figures 4 and 5 to Figure 1, our approach allows us to assert that the security performance of PROG C1 is better than those of PROG C2, except for highly accurate beliefs. Indeed, the remaining uncertainty of PROG C2 is always less than or equal to 2 whilst those of PROG C1 can be arbitrary high for highly inaccurate beliefs. In fact, we have the following result relating the security performance of these




Fig. 3. Shannon entropy-based information flow of PROG C2

programs, the randomisation parameter λ , and the accuracy of the adversary's beliefs.

Proposition 4: The security performance of PROG C1 is better than those of PROG C2 if and only if the randomisation parameter λ of PROG C2 and the accuracy of the adversary's beliefs ω satisfy the following relation.

$$\omega \leq \frac{1}{3 - \max(\lambda, 1 - \lambda)}$$

The few elementary examples above illustrate the applicability of our metric to various threats scenarios. In particular, when it is unavoidable for the attacker to initially have access to some (potentially inaccurate) information about the high input, our approach allows to establish the adversary's beliefs accuracy limit that is tolerable given a specific program. For instance, adversary's beliefs which are less than or equal to 50% accurate are tolerable for PROG C1, since they happen to confuse the adversary instead of helping her. Furthermore, given a collection of programs with the same security objective, we can design a more complex program that adapts dynamically to the context of the adversary, when the accuracy of her beliefs changes. For example, proposition 4 tells us that it is more secure to use PROG C1 than PROG C2 in a context where one can assume that the accuracy of the initial information of the adversary is less than two-fifth; on the other hand, the contrary holds for higher accurate beliefs.

VI. C

This paper presents a new approach to quantitative information flow that incorporates the attacker's beliefs in the model on Rényi min entropy. We investigate the impact of such adversary's extra knowledge on the security of the secret information. Our analyses reveals that inaccurate extra information tend to confuse the adversary by increasing her uncertainty about the hidden secret while accurate information may increase dramatically its vulnerability. We showed the strength of our definitions both in terms of their theoretical properties and their utility by applying them to various threat scenarios and comparing the results to the previous approaches. Our model allows to identify the levels of accuracy for the adversary's

¹See the discussion on the last paragraph of page 298 of [34].




Fig. 4. RU_{bv} of PROG C2




Fig. 5. IL_{bv} of PROG C2

beliefs which are compatible with the security of a given program or protocol.

R

- [1] Miguel Andrés, Catuscia Palamidessi, Peter van Rossum, and Geoffrey Smith. Computing the amount of leakage in information-hiding systems. Technical report, LIX, Ecole Polytechnique, 2009.
- [2] Mohit Bhargava and Catuscia Palamidessi. Probabilistic anonymity. In Martín Abadi and Luca de Alfaro, editors, *Proceedings of CONCUR*, volume 3653 of *Lecture Notes in Computer Science*, pages 171–185. Springer, 2005. <http://www.lix.polytechnique.fr/~catuscia/papers/Anonymity/concur.pdf>.
- [3] Christelle Braun, Konstantinos Chatzikokolakis, and Catuscia Palamidessi. Quantitative notions of leakage for one-try attacks. In *Proceedings of the 25th Conference on Mathematical Foundations of Programming Semantics (MFPS 2009)*, volume 249 of *Electronic Notes in Theoretical Computer Science*, pages 75–91. Elsevier B.V., 2009.
- [4] Konstantinos Chatzikokolakis and Catuscia Palamidessi. Probable innocence revisited. *Theoretical Computer Science*, 367(1-2):123–138, 2006. <http://www.lix.polytechnique.fr/~catuscia/papers/Anonymity/tcsPI.pdf>.
- [5] Konstantinos Chatzikokolakis, Catuscia Palamidessi, and Prakash Panagaden. Anonymity protocols as noisy channels. *Information and Computation*, 206(2-4):378–401, 2008. <http://www.lix.polytechnique.fr/~catuscia/papers/Anonymity/Channels/full.pdf>.
- [6] Konstantinos Chatzikokolakis, Catuscia Palamidessi, and Prakash Panagaden. Anonymity protocols as noisy channels. *Inf. Comput.*, 206(2-4):378–401, 2008.
- [7] Konstantinos Chatzikokolakis, Catuscia Palamidessi, and Prakash Panagaden. On the bayes risk in information-hiding protocols. *Journal of Computer Security*, 16(5):531–571, 2008.
- [8] David Chaum. The dining cryptographers problem: Unconditional sender and recipient untraceability. *Journal of Cryptology*, 1:65–75, 1988.
- [9] David Clark, Sebastian Hunt, and Pasquale Malacaria. Quantitative analysis of the leakage of confidential data. *Electr. Notes Theor. Comput. Sci.*, 59(3), 2001.
- [10] David Clark, Sebastian Hunt, and Pasquale Malacaria. Quantitative information flow, relations and polymorphic types. *Journal of Logic and Computation, Special Issue on Lambda-calculus, type theory and natural language*, 18(2):181–199, 2005.
- [11] David Clark, Sebastian Hunt, and Pasquale Malacaria. Quantitative information flow, relations and polymorphic types. *J. Log. Comput.*, 15(2):181–199, 2005.
- [12] David Clark, Sebastian Hunt, and Pasquale Malacaria. A static analysis for quantifying information flow in a simple imperative language. *Journal of Computer Security*, 15(3):321–371, 2007.
- [13] Ian Clarke, Oskar Sandberg, Brandon Wiley, and Theodore W. Hong. Freenet: A distributed anonymous information storage and retrieval system. In *Designing Privacy Enhancing Technologies, International Workshop on Design Issues in Anonymity and Unobservability*, volume 2009 of *Lecture Notes in Computer Science*, pages 44–66. Springer, 2000.
- [14] Michael R. Clarkson, Andrew C. Myers, and Fred B. Schneider. Belief in information flow. *Journal of Computer Security*, 2008. To appear. Available as Cornell Computer Science Department Technical Report TR 2007-207.
- [15] Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory*. John Wiley & Sons, Inc., 1991.
- [16] Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory (Wiley Series in Telecommunications and Signal Processing)*. Wiley-Interscience, 2006.
- [17] Úlfar Erlingsson and Marco Pistoia, editors. *Proceedings of the 2008 Workshop on Programming Languages and Analysis for Security, PLAS 2008, Tucson, AZ, USA, June 8, 2008*. ACM, 2008.
- [18] Matthias Franz, Bernd Meyer, and Andreas Pashalidis. Attacking unlinkability: The importance of context. In Nikita Borisov and Philippe Golle, editors, *Privacy Enhancing Technologies, 7th International Symposium, PET 2007 Ottawa, Canada, June 20-22, 2007, Revised Selected Papers*, volume 4776 of *Lecture Notes in Computer Science*, pages 1–16. Springer, 2007.
- [19] Joseph Y. Halpern and Kevin R. O’Neill. Anonymity and information hiding in multiagent systems. *Journal of Computer Security*, 13(3):483–512, 2005.
- [20] S. Hamadou, C. Palamidessi, V. Sassone, and E. ElSalamouny. Probable Innocence in the presence of independent knowledge. In *To appear in the Proc. of the sixth International Workshop on Formal Aspects in Security and Trust (FAST2009)*, LNCS. Spr.-Ver., 2009.
- [21] Dominic Hughes and Vitaly Shmatikov. Information hiding, anonymity and privacy: a modular approach. *Journal of Computer Security*, 12(1):3–36, 2004.
- [22] Boris Köpf and David A. Basin. An information-theoretic model for adaptive side-channel attacks. In Ning et al. [29], pages 286–296.
- [23] Pasquale Malacaria. Assessing security threats of looping constructs. In Martin Hofmann and Matthias Felleisen, editors, *Proceedings of the 34th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2007, Nice, France, January 17-19, 2007*, pages 225–235. ACM, 2007.
- [24] Pasquale Malacaria and Han Chen. Lagrange multipliers and maximum information leakage in different observational models. In Úlfar Erlingsson and Marco Pistoia, editor, *Proceedings of the 2008 Workshop on Programming Languages and Analysis for Security (PLAS 2008)*, pages 135–146. Tucson, AZ, USA, June 8, 2008 2008. ACM.
- [25] Pasquale Malacaria and Han Chen. Lagrange multipliers and maximum information leakage in different observational models. In Erlingsson and Pistoia [17], pages 135–146.
- [26] James L. Massey. Guessing and entropy. In *In Proceedings of the 1994 IEEE International Symposium on Information Theory*, page 204, 1994.

- [27] Ira S. Moskowitz, Richard E. Newman, Daniel P. Crepeau, and Allen R. Miller. Covert channels and anonymizing networks. In Sushil Jajodia, Pierangela Samarati, and Paul F. Syverson, editors, *WPES*, pages 79–88. ACM, 2003.
- [28] Ira S. Moskowitz, Richard E. Newman, and Paul F. Syverson. Quasi-anonymous channels. In *IASTED CNIS*, pages 126–131, 2003.
- [29] Peng Ning, Sabrina De Capitani di Vimercati, and Paul F. Syverson, editors. *Proceedings of the 2007 ACM Conference on Computer and Communications Security, CCS 2007, Alexandria, Virginia, USA, October 28-31, 2007*. ACM, 2007.
- [30] Michael K. Reiter and Aviel D. Rubin. Crowds: anonymity for Web transactions. *ACM Transactions on Information and System Security*, 1(1):66–92, 1998.
- [31] Peter Y. Ryan and Steve Schneider. *Modelling and Analysis of Security Protocols*. Addison-Wesley, 2001.
- [32] Steve Schneider and Abraham Sidiropoulos. CSP and anonymity. In *Proc. of the European Symposium on Research in Computer Security (ESORICS)*, volume 1146 of *Lecture Notes in Computer Science*, pages 198–218. Springer, 1996.
- [33] Geoffrey Smith. Adversaries and information leaks (tutorial). In Gilles Barthe and Cédric Fournet, editors, *Proceedings of the Third Symposium on Trustworthy Global Computing*, volume 4912 of *Lecture Notes in Computer Science*, pages 383–400. Springer, 2007.
- [34] Geoffrey Smith. On the foundations of quantitative information flow. In Luca De Alfaro, editor, *Proceedings of the Twelfth International Conference on Foundations of Software Science and Computation Structures (FOSSACS 2009)*, volume 5504 of *Lecture Notes in Computer Science*, pages 288–302, York, UK, March 2009 2009. Springer.
- [35] Geoffrey Smith. On the foundations of quantitative information flow. In Luca de Alfaro, editor, *FOSSACS*, volume 5504 of *Lecture Notes in Computer Science*, pages 288–302. Springer, 2009.
- [36] Paul F. Syverson and Stuart G. Stubblebine. Group principals and the formalization of anonymity. In *World Congress on Formal Methods (1)*, pages 814–833, 1999.
- [37] P.F. Syverson, D.M. Goldschlag, and M.G. Reed. Anonymous connections and onion routing. In *IEEE Symposium on Security and Privacy*, pages 44–54, Oakland, California, 1997.
- [38] Ye Zhu and Riccardo Bettati. Anonymity vs. information leakage in anonymity systems. In *Proc. of ICDCS*, pages 514–524. IEEE Computer Society, 2005.