

Guess what?!

On the impossibility of unconditionally secure public-key encryption

Lorenz Panny

Department of Mathematics and Computer Science,
Technische Universiteit Eindhoven, The Netherlands
lorenz@yx7.cc

Abstract. We (once again) refute recurring claims about a public-key encryption scheme that allegedly provides unconditional security. This is approached from two angles: We give an information-theoretic proof of impossibility, as well as a concrete attack breaking the proposed scheme in essentially no time.

Keywords: public-key cryptography · perfect secrecy · information theory · impossibility · cryptanalysis

1 Introduction

In 2017, *Guess Again*, a public-key encryption scheme claiming *unconditional* security against passive eavesdroppers, was submitted to NIST's call for post-quantum cryptography [1]. Although we publicly broke that scheme with a fast attack script about three hours after the proposals were published by NIST [6], the authors still have not acknowledged the attack nor withdrawn their proposal (though NIST deselected it from advancing to the second round). About seven months later, a paper by a subset of the *Guess Again* authors describing essentially the same system, with no mention of our earlier attack, appeared at *ICMS 2018* [2]. The abstract of that paper states:

We offer a public-key encryption protocol where decryption of a single bit by a legitimate party is correct with probability p that is greater than $1/2$ but less than 1. At the same time, a computationally unbounded (passive) adversary correctly recovers the transmitted bit with probability exactly $1/2$.

In this note, we show that this claim is false, and in fact impossible to achieve.

Acknowledgements. I thank Tanja Lange and Andreas Hülsing for their helpful comments on an earlier draft.

* This work was supported in part by the Commission of the European Communities through the Horizon 2020 program under project number 643161 (ECRYPT-NET).
Date of this document: 2019.10.19.

2 Theory

A typical course in cryptography analyzes the one-time pad early on, deemed a helpful thing to study for multiple reasons: its tremendous historical significance, as a first basic (but illuminating) example of formally provable security, and to provide motivation and intuition for related concepts such as stream ciphers. Later in the cryptographic curriculum, public-key cryptography is introduced, and alas, it turns out that the absolute security guarantee of the one-time pad is unachievable in that setting—computational hardness assumptions must be made to separate honest users from attackers. This observation was pointed out as early as 1976, in the very paper in which Diffie and Hellman invented public-key cryptography [3]:

We note that neither public key cryptosystems nor one-way authentication systems can be unconditionally secure because the public information always determines the secret information uniquely among the members of a finite set. With unlimited computation, the problem could therefore be solved by a straightforward search.

Notice that this well-known impossibility argument assumes a scheme in which each public key comes from at most one private key. This is the apparent “gap” in the impossibility proof that the construction of [1, 2] tries to exploit, by making the scheme probabilistic and incorporating decryption errors (i.e., occasionally decrypting correctly generated ciphertexts into messages different from the original plaintext). Thus, indeed, the public information does not always uniquely identify the secret information anymore and the reasoning above does not apply immediately. However, a similar brute-force argument still works in the presence of decryption errors: Instead of defining a single unique plaintext, a given ciphertext now determines a *likelihood distribution* on the set of potential messages, and the attacker can (using massive but finite computation) simply evaluate this likelihood function and output the *most likely* messages as their guess for the plaintext. Lemma 2 below shows that this attacker “never loses”.

Remark 1. Another viewpoint to approach the question of unconditional public-key cryptography was pursued by Maurer in the nineties [4]: Using information-theoretical bounds on conditional entropy and mutual information, he concludes that two parties cannot possibly negotiate a shared secret over a public channel in an unconditionally secure manner:

[...] if Alice and Bob do not share at least some partially secret information initially, they cannot generate an information-theoretically secure secret key S [...] if they can only communicate over a public channel accessible to Eve, even if this channel is authenticated.

This fact can be rephrased as follows: There exists no unconditionally-secure public-key cryptosystem or public-key distribution protocol. [Footnote]

While, indeed, the machinery from [4] is certainly much stronger than necessary to imply Lemma 2, we will for simplicity present a more concrete and computational proof similar to the brute-force argument from [3] quoted above.

2.1 Proof of impossibility

Recall the protocol flow of a public-key encryption scheme:

- Bob generates a key pair $(\mathbf{pk}, \mathbf{sk})$ and publishes the public key \mathbf{pk} .
- Alice encrypts a message m and sends the ciphertext $c = \text{Enc}_{\mathbf{pk}}(m)$ to Bob.
- Bob decrypts c to obtain a message $m' = \text{Dec}_{\mathbf{sk}}(c)$.

Here, $\text{Enc}_{\mathbf{pk}}()$ and $\text{Dec}_{\mathbf{sk}}()$ are both probabilistic algorithms. Note that it is not required that any of the involved procedures be efficient; in fact, in the following, we only have to assume that $\text{Enc}_{\mathbf{pk}}$ will *eventually* halt almost surely, i.e., with probability 1, for all messages.

Lemma 1. *Consider a public-key encryption scheme with the interface above and message space $\{0, 1\}$, and fix a key pair $(\mathbf{pk}, \mathbf{sk})$.*

Let \mathcal{O} be a (not necessarily deterministic or efficient) algorithm that takes as input a ciphertext and outputs a single bit. Define the distinguishing advantage

$$\text{Adv}(\mathcal{O}) := \Pr[\mathcal{O}(\text{Enc}_{\mathbf{pk}}(1)) = 1] - \Pr[\mathcal{O}(\text{Enc}_{\mathbf{pk}}(0)) = 1],^1$$

where the probabilities are taken over the randomness consumed by \mathcal{O} and $\text{Enc}_{\mathbf{pk}}$. For shorthand, write $p_m(c) := \Pr[\text{Enc}_{\mathbf{pk}}(m) = c]$. Then it holds for all \mathcal{O} that

$$\text{Adv}(\mathcal{O}) \leq \sum_c \max\{0, p_1(c) - p_0(c)\}.$$

Proof. We first unroll the definition of $\text{Adv}(\mathcal{O})$, using linearity of expectations and the fact that $\Pr[X = 1] = E[X]$ when X is a random bit:

$$\begin{aligned} \text{Adv}(\mathcal{O}) &= \Pr[\mathcal{O}(\text{Enc}_{\mathbf{pk}}(1)) = 1] - \Pr[\mathcal{O}(\text{Enc}_{\mathbf{pk}}(0)) = 1] \\ &= E[\mathcal{O}(\text{Enc}_{\mathbf{pk}}(1))] - E[\mathcal{O}(\text{Enc}_{\mathbf{pk}}(0))] \\ &= E[\mathcal{O}(\text{Enc}_{\mathbf{pk}}(1)) - \mathcal{O}(\text{Enc}_{\mathbf{pk}}(0))] \\ &= \sum_c (\Pr[\text{Enc}_{\mathbf{pk}}(1) = c] - \Pr[\text{Enc}_{\mathbf{pk}}(0) = c]) \cdot E[\mathcal{O}(c)] \\ &= \sum_c (p_1(c) - p_0(c)) \cdot E[\mathcal{O}(c)]. \end{aligned} \quad (*)$$

There are positive and negative terms in the sum $(*)$, determined by the sign of $p_1(c) - p_0(c)$. To maximize the advantage, it is clearly optimal to maximize

¹ Note that typical definitions of the advantage take an absolute value to symmetrize the definition with respect to distinguishers which are *worse* than pure guessing—and are thus in fact *good* at distinguishing. In our scenario, this would clutter the notation and introduce unnecessary case distinctions, so we stick to the present definition and point out that a negative advantage can easily be turned into a positive one of equal magnitude by inverting the distinguisher’s output.

Also note that this definition of the advantage is equal to the “other” choice $\Pr[\mathcal{O}(\text{Enc}_{\mathbf{pk}}(0)) = 0] - \Pr[\mathcal{O}(\text{Enc}_{\mathbf{pk}}(1)) = 0]$.

$E[\mathcal{O}(c)]$ when $p_1(c) > p_0(c)$ and minimize it when $p_0(c) > p_1(c)$. Since $\mathcal{O}(c)$ is a random variable on $\{0, 1\}$, its expectation $E[\mathcal{O}(c)]$ must lie in the interval $[0; 1]$; therefore, the value of $\text{Adv}(\mathcal{O})$ is upper bounded by the quantity

$$\sum_c (p_1(c) - p_0(c)) \cdot \begin{cases} 0 & \text{if } p_0(c) \geq p_1(c); \\ 1 & \text{if } p_1(c) > p_0(c). \end{cases}$$

This is the same thing as

$$\sum_c \max\{0, p_1(c) - p_0(c)\}. \quad \square$$

Example. For a public-key encryption scheme on $\{0, 1\}$ with guaranteed correct decryption, Dec_{sk} has advantage one: Since $\text{Dec}_{\text{sk}}(\text{Enc}_{\text{pk}}(m)) = m$ for all m , the definition immediately simplifies to $\text{Adv}(\text{Dec}_{\text{sk}}) = \Pr[1 = 1] - \Pr[0 = 1] = 1$.

On the other hand, a distinguisher \mathcal{O} that ignores its input and simply outputs a uniformly random bit has advantage zero, since both probabilities in the definition of $\text{Adv}(\mathcal{O})$ equal $1/2$ in that case.

Lemma 2. *As above, consider a public-key encryption scheme with message space $\{0, 1\}$,² and fix a public key (pk, sk) .*

Using Enc_{pk} , one can construct a deterministic (but not necessarily efficient) algorithm \mathcal{A}_{pk} that achieves the upper bound proved in Lemma 1. In particular, it is at least as good as Dec_{sk} at distinguishing the two random variables $\text{Enc}_{\text{pk}}(0)$ and $\text{Enc}_{\text{pk}}(1)$, despite not knowing the secret key.

Proof. The argument consists of two parts: Constructing a maximum-likelihood estimator \mathcal{A}_{pk} for the message m from (a single sample of) the random variable $\text{Enc}_{\text{pk}}(m)$, and then showing that \mathcal{A}_{pk} achieves the upper bound from Lemma 1 to conclude that Dec_{sk} cannot possibly do better than \mathcal{A}_{pk} .

The adversary \mathcal{A}_{pk} is constructed as follows: Given a ciphertext c , compute the two probabilities $p_0(c) = \Pr[\text{Enc}_{\text{pk}}(0) = c]$ and $p_1(c) = \Pr[\text{Enc}_{\text{pk}}(1) = c]$ by iterating through all possible values of the randomness consumed by Enc_{pk} ³ and counting how often the ciphertext c is observed for either message. Then simply output the plaintext guess 0 if $p_0(c) \geq p_1(c)$ and the guess 1 if $p_1(c) > p_0(c)$. Notice that the adversary \mathcal{A}_{pk} indeed implements a maximum-likelihood estimator for the message $m \in \{0, 1\}$ from a single encryption $\text{Enc}_{\text{pk}}(m)$ of m . Moreover, the algorithm \mathcal{A}_{pk} achieves the upper bound from Lemma 1 (in fact, it is crafted precisely to do so), hence no Dec_{sk} can ever do better than \mathcal{A}_{pk} . \square

The claims from [1, 2] amount to the assertion that knowledge of the secret key permits honest users to succeed in decrypting with probability

$$\Pr[\text{Dec}_{\text{sk}}(\text{Enc}_{\text{pk}}(m)) = m] > 1/2,$$

² The restriction to one-bit messages does not sacrifice any generality: Any message space of cardinality at least two can represent single bits, and conversely, bit strings suffice to encode any other message space.

³ By the assumption that Enc_{pk} halts almost surely, it consumes only a finite amount of randomness, hence \mathcal{A}_{pk} halts after finitely many steps as well.

while at the same time, for *all* adversaries \mathcal{A}_{pk} given only the public key,

$$\Pr[\mathcal{A}_{\text{pk}}(\text{Enc}_{\text{pk}}(m)) = m] = 1/2.$$

These statements imply $\text{Adv}(\text{Dec}_{\text{sk}}) > 0$ and $\text{Adv}(\mathcal{A}_{\text{pk}}) = 0$, which is impossible due to Lemma 2. Essentially, the only way to make sure \mathcal{A}_{pk} cannot do better than guessing is to transmit no information whatsoever about the plaintext in the ciphertext, but then Dec_{sk} must also resort to plainly guessing what the encrypted message was supposed to be, which is utterly useless; in fact, it entirely defeats the purpose of communication as defined by Shannon in his seminal paper [7]: “*The fundamental problem of communication is that of reproducing at one point either exactly or approximately a message selected at another point.*”

3 Practice

We now show how the concrete proposal *Guess Again* works, how it is broken in practice, and where exactly the flaw in the authors’ security argument lies.⁴

3.1 Description of *Guess Again*

The *Guess Again* scheme makes use of random walks. For each $x \in \mathbb{Z}$ and $k \geq 0$, we thus define the random variable $\text{walk}(x, k)$ as the end point of a random walk of k steps, each adding either $+1$ or -1 with probability $1/2$, starting from x . Write $s \leftarrow S$ for sampling s uniformly at random from a finite set S .

- **Parameters:** Positive integer constants n, f, g, h . [2] recommends $n = 256$, $g = h = 2000$, and $f = 100000$. Note [1] suggested $f = 120000$ instead.
- **Key generation:** Bob picks a private key $b \leftarrow \{0, \dots, n-1\}$ and computes the public key $B := \text{walk}(b, h)$. If $B \geq n - 1$, Bob starts over with a new b .
- **Encryption:** Alice samples values $(a, t, \delta) \leftarrow \{B, \dots, n-1\} \times \{f, g\} \times \{\pm \frac{1}{2}\}$ and computes $A := \text{walk}(a, t) + \delta$. She repeats this step “sufficiently often” and groups the resulting pairs (a, A) into four lists $L_f^<$, $L_g^<$, $L_f^>$, and $L_g^>$, according to the condition $A \leq B$ and the value of t . She then selects $r \leftarrow \{<, >\}$ and $s \leftarrow \{f, g\}$, and a pair $(a_0, A_0) \leftarrow L_s^r$. Her ciphertext equals (m, a_0) if $(r, s) \in \{(<, f), (>, g)\}$ and $(1 \oplus m, a_0)$ else.⁵ Upon receiving such a ciphertext (c, a_0) , Bob outputs either the bit $m' = c$ when $b < a_0$, or $m' = 1 \oplus c$ when $b \geq a_0$.

It is stated in [2, Section 6] that (for the example parameters above) this protocol has a ≈ 0.55 chance of transmitting a bit correctly, i.e., achieving $m' = m$.

⁴ The newer paper [2] does not use the name “Guess Again”, nor does it mention the NIST process [5]. Since the proposals in [1] and [2] are essentially identical, we shall use the name *Guess Again* to refer to both, and point out differences where relevant.

⁵ The bit-flipping description deviates from [1, 2], which uses the equivalent viewpoint of labelling intervals instead.

3.2 Biases in the ciphertexts

The security argument in [2, Section 4.2] states that the bit m is flipped with probability $1/2$ during the encryption procedure, hence c contains no information whatsoever about the plaintext bit m . Moreover, the additional component a_0 of the ciphertext is uniform in a public range $\{B, \dots, n-1\}$, hence also does not leak information about m . Probabilistically speaking, these observations can be phrased as the (indeed true) fact that the distributions of the values c and a_0 are both independent of the plaintext m .

However, the crucial flaw in the security argument of *Guess Again* is that **the joint distribution of (c, a_0) is not independent of m** . To see this, notice that those values of $a \leftarrow \{B, \dots, n-1\}$ that lead to $A < B$ must, on average, be closer to B than those that lead to $A > B$. Moreover, this tendency clearly depends on the chosen number of steps of the random walk: Shorter walks lead to closer nodes; in fact, the expected (absolute) distance travelled by a simple random walk of k steps is $\Theta(\sqrt{k})$. Therefore, there is a probabilistic dependency between the choice of $t \in \{f, g\}$ and the condition $A \leq B$. Working out the distribution of a conditioned on the choice of $A \leq B$ and $t \leftarrow \{f, g\}$ is not particularly difficult, but rather tedious, hence we only give qualitative simulation results showing the distribution of a relative to each of the four choices. See Figure 1 and Table 1. The bottom line is that *“flipped” ciphertexts are smaller on average*, which is precisely the leakage exploited by the (simple) attack script in Section 3.3.

Table 1. Biases in a depending on the secret random choices during encryption. Notice in particular that the decision whether the message bit is flipped or not is correlated with the bias in a .

$A \leq B$	$t \leftarrow \{f, g\}$	bias in a	c
<	f	slightly biased towards B	m
<	g	strongly biased towards B	$1 \oplus m$
>	f	no significant bias	$1 \oplus m$
>	g	slightly biased away from B	m

3.3 The attack

Recall from Section 3.1. that a *Guess Again* ciphertext consists of a tuple (c, a_0) with c a single bit and a_0 a value between B and $n-1$. The bit c is either equal to the message bit m , or it is negated. In Section 3.2, we showed that the choice between $c = m$ and $c = 1 \oplus m$ influences the distribution of the corresponding value a_0 ; concretely, a_0 is biased towards B when $c = 1 \oplus m$. This means we can mount a distinguishing attack by comparing a_0 to its expected values relative to the assumptions $c = m$ and $c = 1 \oplus m$, and output a guess according to which distribution a_0 is more likely to come from. This distinguisher has a better than $1/2$ chance of identifying the message given a ciphertext.

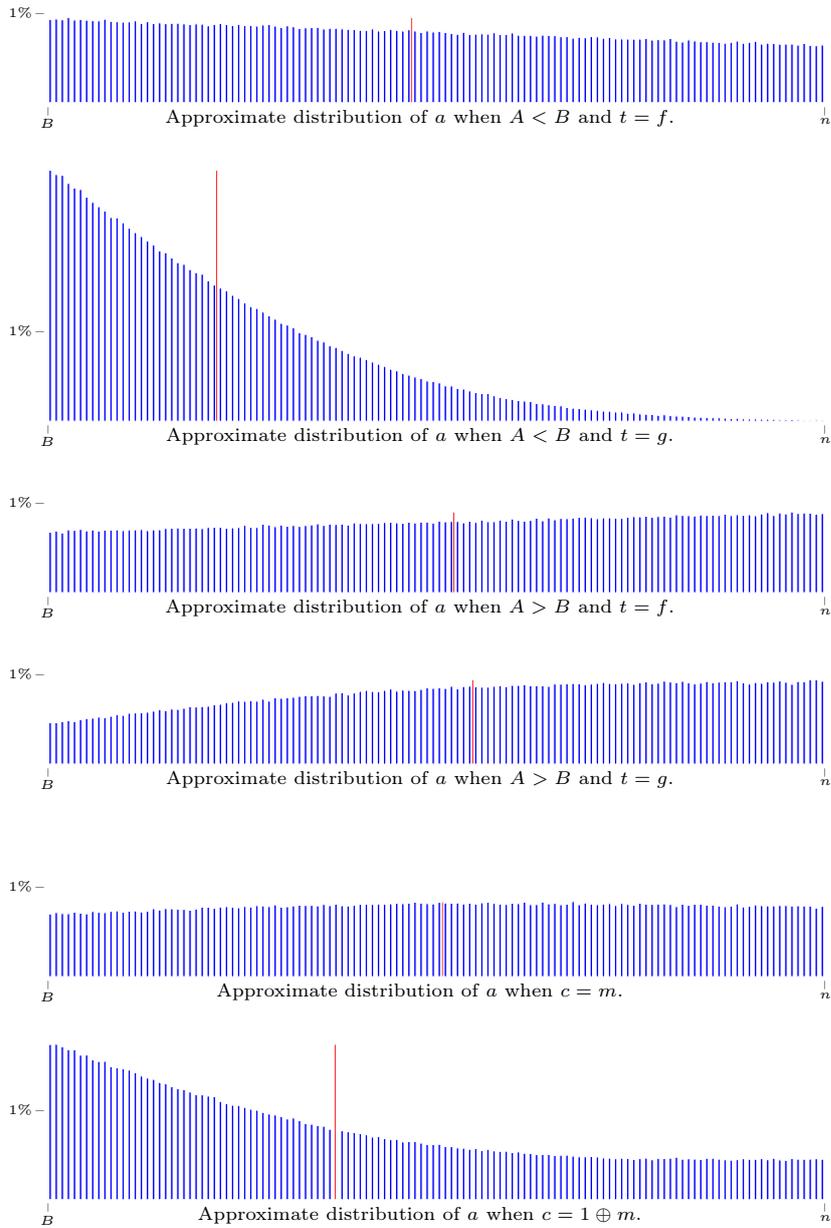


Figure 1. Approximate distribution of a conditioned on the property $A \leq B$ and the number of steps t , where $(n, f, g) = (256, 100\,000, 2000)$ and B has been fixed as $n/2$. (Similar graphs result from other public keys.) The mean of each distribution is marked by a red vertical line. For efficiency reasons, we approximated the distribution of the random walks $walk(x, k)$ by the normal distribution $\mathcal{N}(x, k)$. Each of these histograms contains one million data points.

When more than one ciphertext encrypting the same message is sent at a time to boost the success rate, as was done in [1], accumulating this leakage allows for a very reliable distinguisher: Group the ciphertexts by their bit c , compute the average value of a_0 in each group, and compare. Output the c with the bigger average a_0 as a guess for m . This approach is implemented in the script below and works on 100% of the example ciphertexts included in [1].

In [2, Section 5], the authors—perhaps in response to, but with no mention of, our earlier attack—warn against naïvely encrypting the same bit multiple times, foreshadowing that “statistical attacks [...] may be possible”. However, if the construction truly were information-theoretically secure, any number of repetitions would never leak information (assuming uncorrelated randomness), hence acknowledging the reality of “statistical attacks” is clearly incompatible with their claim of unconditional security upheld elsewhere in the paper.

Figure 2 shows the essential part of our Python script [6] sent to NIST’s pqc-forum mailing list on December 21st, 2017, not including boilerplate code to parse and break the example ciphertext files contained in *Guess Again*’s NIST submission package.

```
l = 4000

def recover_bit(ct, bit):
    assert bit < len(ct) // 1
    ts = [struct.unpack('BB', ct[i:i+2]) for i in range(1*bit, 1*(bit+1), 2)]
    xs, ys = [a for a, b in ts if b == 1], [a for a, b in ts if b == 2]
    return sum(xs) / len(xs) >= sum(ys) / len(ys)

def decrypt(ct):
    res = sum(recover_bit(ct, b) << b for b in range(len(ct) // 1))
    return int.to_bytes(res, len(ct) // 1 // 8, 'little')
```

Figure 2. Excerpt of the attack script sent to NIST’s mailing list [6].

Note that this script does *not* implement the adversary from Lemma 2; it does not brute-force anything and in fact runs in time linear in the input size. We remark that it is certainly possible to tweak *Guess Again* in such a way that any concrete attack would be (much) less efficient—for an extreme example, simply replace the entire construction by a computationally secure public-key encryption scheme—but the fundamental impossibility of *information-theoretic* security remains. Therefore, we conclude that the security claims of *Guess Again* are both theoretically impossible as well as demonstrably broken in practice.

References

- [1] Mariya Bessonov, Dima Grigoriev, Alexey Gribov, and Vladimir Shpilrain. Guess Again, 2017. Submission to [5].
- [2] Mariya Bessonov, Dima Grigoriev, and Vladimir Shpilrain. A framework for unconditionally secure public-key encryption (with possible decryption errors). In James H. Davenport, Manuel Kauers, George Labahn, and Josef Urban, editors, *Mathematical Software – ICMS 2018*, pages 45–54. Springer, July 2018.
- [3] Whitfield Diffie and Martin E. Hellman. New directions in cryptography. *IEEE Trans. Information Theory*, 22(6):644–654, November 1976.
- [4] Ueli Maurer. Information-theoretic cryptography. In Michael J. Wiener, editor, *Advances in Cryptology – CRYPTO ’99*, volume 1666 of *LNCS*, pages 47–64. Springer, August 1999.
<ftp://ftp.inf.ethz.ch/pub/crypto/publications/Maurer99.pdf>.
- [5] National Institute of Standards and Technology. Post-quantum cryptography standardization: Call for proposals announcement, December 2016.
<https://csrc.nist.gov/Projects/Post-Quantum-Cryptography/Post-Quantum-Cryptography-Standardization>.
- [6] Lorenz Panny. OFFICIAL COMMENT: Guess Again, 21 December 2017. Email message to the `pqc-forum@list.nist.gov` mailing list.
<https://csrc.nist.gov/CSRC/media/Projects/Post-Quantum-Cryptography/documents/round-1/official-comments/guess-again-official-comment.pdf>.
- [7] Claude E. Shannon. A mathematical theory of communication. *Bell System Technical Journal*, 27(3):379–423, July 1948.