# Bit Security of Quantum Key Search

Marc Fischlin and Evangelos Gkoumas

Cryptoplexity, Technische Universität Darmstadt, Germany,
first.last@tu-darmstadt.de

**Abstract.** A common presumption in cryptography is that quantum key search effectively halves the level of bit security compared to the classical setting, leading for example to the recommendation to use AES with 256-bit keys instead of 128-bit keys. From a very coarse point of view, this is perspicuous by the speed-up obtained via Grover's search algorithm. On closer inspection, however, it lacks a formal justification, especially if the AES key is not perfectly random but only statistically close to uniform if generated by a quantum key distribution scheme. In other words, the question is how the statistical distance influences the quantum key search, or viewed from an implementation point of view, how one should choose the statistical distance to achieve the best security bounds.

Our starting points are the recent works about bit security of Micciancio and Walter (Eurocrypt 2018), Watanabe and Yasunaga (Asiacrypt 2021), and Lee (Communication in Cryptology 2024) in the classical setting. We transfer them to the quantum setting and discuss the security against quantum key search if the keys are close to uniform. We then argue that to achieve an optimal bit security level of $\lambda/2$ bits for $\lambda$-bit keys against quantum search, it is advisable to set the statistical distance of the keys from uniform to a value in the range from $2^{-\lambda}$ to $2^{-\lambda/2}$. Going below these bounds does not yield any advantage for the bit security, and going above this range gives worse bit security guarantees.

**Keywords:** Bit Security · Quantum Key Search · Hellinger distance

## 1 Introduction

An easy way to compare the strength of cryptographic primitives like block ciphers, hash functions, or even public-key schemes is via the *bit security* of the primitives. For example, a hash function like SHA-256 is attributed a security level of $\lambda' = 128$, and thus in the same realm as 256-bit elliptic curves. This is based on the observation that the currently best attack on SHA-256 via the general birthday collision search requires approximately $O(2^{\lambda/2})$ steps for $\lambda = 256$. Similarly, for $\lambda$-bit elliptic curves, Pollard's rho method also needs $O(2^{\lambda/2})$ steps. Hence, for a hash-and-sign signature scheme of security level $\lambda' = 128$, one could deploy SHA-256 with a 256-bit curve (albeit one would need to take the security loss due to the combination in the signature protocol into account). The bit security levels of cryptographic primitives and their comparison have thus also entered recommendations for cryptographic primitives for security agencies [2,6].

### 1.1 Defining Bit Security

Lenstra [17] informally defined a cryptographic system to offer $\lambda'$-bit security if any attacker is expected to require the effort of at least $2^{\lambda'}$ to break the system. A more rigorous and common approach is to measure bit security by the logarithm of time divided by the success probability, $\log T/\epsilon$, and taking the minimum over all adversaries. More specifically, in the setting of search problems, a scheme has bit security $\lambda'$ if no adversary can exceed a success probability of $\epsilon$ in less than $\epsilon \cdot 2^{\lambda'}$ steps for any $\epsilon$.

Recently, there were several proposals how to capture bit security more formally and more broadly. Micciancio and Walter [20] proposed definitions for search and decision games based on entropy notions. Bit security in this case is defined as

$$\min_{\mathcal{A}} \log\left(T_{\mathcal{A}}/\mathsf{adv}(\mathcal{A})\right),$$

where $\mathsf{adv}(\mathcal{A})$ is the adversary's advantage in term of entropy. Building on this, Watanabe and Yasunaga [34] presented a new framework for defining bit security, focusing on demonstrating the advantage of the adversary. This demonstrating advantage measures how often an outer adversary $\mathcal{B}$ needs to invoke an inner adversary $\mathcal{A}$ (with some success probability $\epsilon_{\mathcal{A}}$) to find a correct answer with overwhelming probability. The bit security for search problems is then defined as

$$\min_{\mathcal{A},\mathcal{B}} \log\left(N_{\mathcal{B}} \cdot T_{\mathcal{A}}\right),$$

where $N_{\mathcal{B}}$ is the number of invocations of $\mathcal{A}$. Note that the success probability $\epsilon_{\mathcal{A}}$ is implicit in $N_{\mathcal{B}}$: For any $\epsilon_{\mathcal{A}}$ one usually requires $N_{\mathcal{B}} = 1/\epsilon_{\mathcal{A}}$ many runs of adversary $\mathcal{A}$ to find a solution. The approaches were shown to be equivalent [19].

More recently, Lee [16] conducted a comparative analysis of the frameworks proposed by Micciancio and Walter [20] and by Watanabe and Yasunaga [34]. Lee gave a more comprehensive framework for capturing search and decision games simultaneously, with the help of a so-called dummy adversary $\mathcal{A}^{\text{dummy}}$ to define a baseline probability. The dummy adversary can be considered the best non-adaptive strategy to win the game in the same time as the adversary $\mathcal{A}$. As in [34], Lee then defines bit security via an observational advantage of an adversary $\mathcal{B}$, amplifying the distinguishing probability between the adversary $\mathcal{A}$ and the dummy adversary $\mathcal{A}^{\text{dummy}}$ of the same complexity $T_{\mathcal{A}}$ as $\mathcal{A}$. This also results in the definition of the bit security level of $\min_{\mathcal{A},\mathcal{B}} \log(N_{\mathcal{B}} \cdot T_{\mathcal{A}})$, but the definition is now independent of the structure of the security game.

Both the work of Watanabe and Yasunaga [34] and of Lee [16], referring to an earlier work by Yasunaga [11,35], also pointed out the importance of the Hellinger distance in determining the observational advantage. Namely, the number $N_{\mathcal{B}}$ of invocations of the adversary $\mathcal{A}$ is roughly in the order of the square of the Hellinger distance (between the probability of $\mathcal{A}$ winning and the probability of the dummy adversary $\mathcal{A}^{\text{dummy}}$ winning), such that bit security can be well approximated by

$$\min_{\mathcal{A}} \log \left( T_{\mathcal{A}} / (d_{\text{Hell}}(\mathcal{A}, \mathcal{A}^{\text{dummy}}))^2 \right) .$$

### 1.2 From Statistical Distance to Bit Security

We are interested in the following question:

> *Suppose that one aims to derive an AES key of* 256 *bits to protect against key search of a quantum adversary. The key generation is not perfectly uniform but has a slight statistical distance from the uniform distribution. What is the reasonable range for this statistical distance?*

The question appears naturally when the key is derived via quantum key distribution. In this case, privacy amplification as the final step of the protocol is usually implemented via randomness extractors [38,12], generating statistically close to uniform keys. Making the statistical distance arbitrarily small is inefficient because the smaller the distance, the more bits have to be truncated. Specifically, for a statistical distance $\Delta$, one must roughly sacrifice $2 \log 1/\Delta$ bits [12].

The problem also appears in any modern hybrid key exchange setting combining classical and quantum-safe keys. Any protocol like TLS 1.3 [27] is based on Krawczyk's extract-and-expand paradigm [14], implemeneted either via HKDF [15] or some AES-based function [7]. The quality of the extraction steps in terms of statistical distance from uniform depends on properties of the underlying cryptographic primitive [10,9]. This is also true for multi-input key derivation functions, such as the hybrid version of TLS 1.3 [28].

We note that in the classical setting, with the interpretation of bit security $\lambda'$ that no adversary can exceed a success probability of $\epsilon$ in less than $\epsilon \cdot 2^{\lambda'}$ steps for any $\epsilon$, one can easily compute an upper bound on the bit security in terms of the statistical distance. For this, consider an adversary $\mathcal{A}$ testing at most $T_{\mathcal{A}}$ potential key values of $\lambda$ bits. The best strategy of $\mathcal{A}$ is to sort the keys in decreasing likelihood, i.e., testing the most likely keys first, where we assume that determining the order is for free. The situation is depicted in Figure 1.

Then we can upper bound the success probability of the adversary $\mathcal{A}$ when testing the most likely $T_{\mathcal{A}}$ keys first by

$$\epsilon \leq T_{\mathcal{A}} \cdot 2^{-\lambda} + \Delta,$$

where $\Delta$ is the statistical distance of the key distribution from uniform (the shaded area in Figure 1). The reason is that the probability of finding the right value for a uniform key would be $T_{\mathcal{A}} \cdot 2^{-\lambda}$, also taking into account the increased probability via the statistical distance in the first steps. Using

$$T_{\mathcal{A}} \cdot 2^{-\lambda} + \Delta \leq T_{\mathcal{A}} \cdot (2^{-\lambda} + 2^{-\log 1/\Delta}) \leq T_{\mathcal{A}} \cdot 2^{\max\{-\lambda, -\log 1/\Delta\}+1},$$

we hence get a bit security of $\lambda' = \min\{\lambda, \log 1/\Delta\} - 1$. This bound is quasi-tight since, in general, the statistical distance may concentrate on a single key only, such that an adversary outputting this key would succeed with probability $2^{-\lambda} + \Delta \geq 2 \cdot \min\{2^{-\lambda}, 2^{-\log 1/\Delta}\}$ with a single guess.

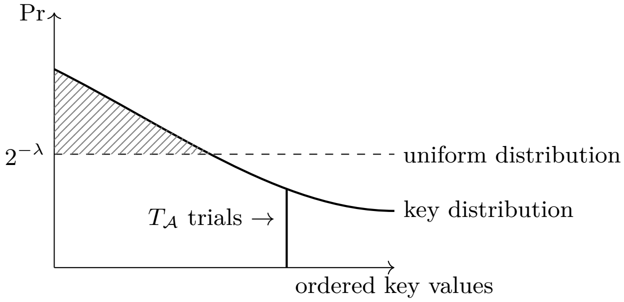

Fig. 1: Optimal search strategy in the classical setting (neglecting ordering effort): Check the most probable keys first.

Setting the statistical distance to $\Delta \leq 2^{-\lambda}$ maximizes the above bound $\min\{\lambda, \log 1/\Delta\} - 1$ for bit security, meaning that the classical setting with the straightforward interpretation of bit security suggests using a statistical distance close to $2^{-\lambda}$. However, it is unclear if this is also true in the quantum setting. For one, the approach here uses a simple approach to define bit security. Then, when a quantum adversary searches for a key, we cannot easily argue about the optimal strategy to order keys in sequential order.

We note that Micciancio and Walter [20] as well as Yasunaga [35] already discuss the relationship of closeness of distributions and bit security for classical games. As mentioned by Yasunaga [35], if the statistical distance of two distributions is $2^{-\lambda}$, e.g., the distance of a uniform key vs. an almost uniform key as in our setting, then the game preserves $\lambda$-bit security (up to additive terms). Micciancio and Walter [20] showed that this also holds if the max-log distance of the distributions is at most $2^{-\lambda/2}$, and Yasunaga [35] extended this to a difference in the Hellinger distance of $2^{-\lambda/2}$. He also proves, by presenting a specific game, that this generally does not hold if the distributions have a statistical distance of $2^{-\lambda/2}$. In contrast, we investigate here quantum key search games and the precise effect of the more common notion of statistical distance on the actual bit security for this game.

### 1.3 Our Results

We start with the bit security notion of Lee [16] according to observational advantages. This model is independent of the game's structure (but not the run time of the adversary and the dummy adversary). Yet, the structural simplicity allows us to transfer the setting to the case of quantum adversaries. We model, for example, the interaction of the adversary with the challenger in the game as a joint quantum computation, such that the dummy adversary, implementing the best non-adaptive strategy, will only act in the first stage of the interaction. We note that the observational distinguisher $\mathcal{B}$, receiving information about the effectiveness of the quantum or dummy adversary, will remain classical. Its task is to distinguish two worlds and amplify the advantage; quantum power is only used to search for the key.

We then define a quantum key search game in which we randomly select a secret key and allow the adversary to test candidates. This corresponds to a black-box search on, say, an AES key space $\{0,1\}^\lambda$ for $\lambda = 256$ where the adversary can trial decrypt a given ciphertext and verify its guess. The distribution of keys will not be perfectly random but instead have a statistical distance $\Delta$ from the uniform distribution. Using an upper bound for quantum searches for skewed input distributions by Montanaro [22], in a version by He et al. [13], we present an upper bound on the adversary's success probability in terms of its run time $T_{\mathcal{A}}$ in our setting. This bound is roughly $T_{\mathcal{A}}^2 \cdot 2^{-\lambda} + \Delta$ for any quantum adversary with run time $T_{\mathcal{A}}$.

It turns out that the optimal strategy of the dummy adversary in our setting is to guess the most likely key and win with probability at most $2^{-\lambda} + \Delta$, independently of its run time. We can then use the two probabilities, the one for $\mathcal{A}$ and the one for $\mathcal{A}^{\text{dummy}}$, to compute the Hellinger distance and use it to estimate the post-quantum bit security of the key search game. In doing so, we derive the following results:

1. Statistical distance smaller than $2^{-\lambda}$ will not yield an improved security level. This is shown by upper and lower bounds for the bit security of key search, which match up to additive terms.
2. Statistical distances in the range $2^{-\lambda}$ and $2^{-\lambda/2}$ yield the optimal bound of bit security $\lambda' = \lambda/2$, again shown via corresponding upper and lower bounds. Here, $2^{-\lambda}$ is the conservative choice, giving tight bounds also in cases where, for instance, quantum adversaries cannot rely on arbitrary run times but are bounded in computing power.
3. Statistical distance larger than $2^{-\lambda/2}$ give worse security guarantees than $\lambda' = \lambda/2$.

These conclusions match common intuition but now rely on a proper formal foundation. We present the technical statements in Section 4 and discuss their meaning in Section 5.

We also discuss in Section 6 our findings in light of quantum key distributions and suggestions for parameter choices in this area. Specifically, there appears to be a general tendency in the literature [25,23,30,21,31,5,24,18,26] to choose the security parameter $\varepsilon$ close to $10^{-10} \approx 2^{-33}$, where $\varepsilon$ can be roughly viewed as our value for statistical distance.[1] However, this choice of $\varepsilon = 10^{-10}$ is not well

[1] The situation is more complex than outlined here, partly because it is common to capture both correctness and security with this single parameter $\varepsilon$, where the correctness error could be the dominating part. See Section 6 for details.

motivated from our viewpoint. In fact, our results indicate that if the quantum key distribution should be used to generate 256-bit AES keys, then setting $\varepsilon = 10^{-10}$ may be too optimistic.

# 2 Preliminaries

## 2.1 Notation

We consider a classical adversary $\mathcal{A}$ that interacts in a game with a challenger $\mathsf{X}$. All parties in a game receive the security parameter $\lambda \in \mathbb{N}$ in unary as input. The symbols $\mathsf{win}$ or $\mathsf{lose}$, eventually output by the challenger, indicate if the adversary has won or lost the game. When considering quantum parties, we usually use a subscript $Q$ and write $\mathcal{A}_Q$ and $\mathsf{X}_Q$. We denote by $T_{\mathcal{A}_Q}$ the run time of adversary $\mathcal{A}_Q$. The run time usually depends on the security parameter $\lambda$. When calling a quantum oracle $\mathsf{O}$, it is understood that we use the common approach to make it a unitary operation by considering the map $|x\rangle|y\rangle \mapsto |x\rangle|y \oplus \mathsf{O}(x)\rangle$, where $x \in \{0,1\}^n$ and $y \in \{0,1\}^m$.

Calligraphic letters like $\mathcal{P}$ or $\mathcal{Q}$ usually represent probability distributions or sets. For a distribution $\mathcal{P}$ we write $\mathcal{P}(x)$ for the probability that the distribution outputs $x$. The notation $x \leftarrow\!\!\$\ \mathbf{S}$ means that $x$ is chosen uniformly from a finite set $\mathbf{S}$. The key space is denoted as $\mathcal{K}$ such that sampling a random key is denoted as $\mathsf{k} \leftarrow\!\!\$\ \mathcal{K}$. Unless mentioned differently, all logarithms log are to base 2.

## 2.2 Probability Metrics

This section delineates the primary definitions and properties of the two metrics employed to formulate our concepts: *total variation distance* and *Hellinger distance.*

**Definition 1 (Statistical/Total Variation distance).** *The total variation (aka. statistical) distance between two discrete distributions $\mathcal{P}$ and $\mathcal{Q}$ on the same domain $\Omega$ is defined as*

$$\Delta(\mathcal{P},\mathcal{Q}) = \frac{1}{2}\sum_{x\in\Omega} \mid \mathcal{P}(x) - \mathcal{Q}(x) \mid .$$

The Hellinger distance is frequently employed in information theory to measure the closeness between distributions. It is the probabilistic equivalent of the Euclidean distance. It has turned out to be an instrumental notion in cryptography [29,11,35,34,16,8].

**Definition 2 (Hellinger distance).** *The Hellinger distance $d_{Hell}$ between two discrete distributions $\mathcal{P}$ and $\mathcal{Q}$ on the same domain $\boldsymbol{\Omega}$ is defined as*

$$d_{Hell}(\mathcal{P},\mathcal{Q}) = \frac{1}{\sqrt{2}}\sqrt{\sum_{x\in\Omega}(\sqrt{\mathcal{P}(x)} - \sqrt{\mathcal{Q}(x)})^2}.$$

A different formulation associates the Hellinger distance with the Bhattacharyya coefficient (BC):

$$d_{\mathrm{Hell}}(\mathcal{P}, \mathcal{Q}) = \sqrt{1 - \mathrm{BC}(\mathcal{P}, \mathcal{Q})},$$

where

$$\mathrm{BC}(\mathcal{P}, \mathcal{Q}) = \sum_{x \in \Omega} \sqrt{\mathcal{P}(x) \cdot \mathcal{Q}(x)}.$$

The Bhattacharyya coefficient quantifies the similarity of two random statistical samples. Furthermore, the Bhattacharyya coefficient is closely related to the concept of fidelity, which is known from quantum information theory to measure the closeness of density matrices.

For binary distributions $\mathcal{P}, \mathcal{Q}$ with $\epsilon_{\mathcal{P}} = \mathcal{P}(1)$ and $\epsilon_{\mathcal{Q}} = \mathcal{Q}(1)$ the Hellinger distance is thus given by

$$d_{\mathrm{Hell}}(\mathcal{P}, \mathcal{Q}) = \sqrt{1 - \sqrt{\epsilon_{\mathcal{P}} \cdot \epsilon_{\mathcal{Q}}} - \sqrt{(1 - \epsilon_{\mathcal{P}}) \cdot (1 - \epsilon_{\mathcal{Q}})}}.$$

Both distances serve as metrics for assessing the similarity between distributions. They adhere to the classical analogue of the Fuchs–van de Graaf inequalities:

$$d_{\mathrm{Hell}}(\mathcal{P}, \mathcal{Q})^2 \leq \mathit{\Delta}(\mathcal{P}, \mathcal{Q}) \leq \sqrt{2} \cdot d_{\mathrm{Hell}}(\mathcal{P}, \mathcal{Q}).$$

For additional properties and proofs pertaining to statistical distances, we refer the reader to [1],[29] and [35].

## 3 Definitions

### 3.1 Quantum Security Game

We consider a cryptographic game $\mathsf{G} = (\mathsf{X}, \mathsf{W})$ interacting with an adversary $\mathcal{A}$ as in [16] abstractly as an interaction between the challenger $\mathsf{X}$ and the adversary $\mathcal{A}$, where both parties receive the security parameter $\lambda$ as unary input. At the end of the interaction, the challenger outputs a bit $\mathsf{win}$ or $\mathsf{lose}$ to indicate if the adversary has won or not. Formally, in [16], the decision is captured through the condition predicate $\mathsf{W}$, which evaluates the challenger's view, consisting of the input, the internal randomness, and the incoming messages from $\mathcal{A}$, to $\mathsf{win}$ or $\mathsf{lose}$. This predicate is only used for a better comparison with previous approaches, and we keep it here as well for compatibility reasons.

First, we introduce the quantum variants of the classical definitions provided by Lee [16] to capture security games. We envision the interaction of the quantum adversary $\mathcal{A}_Q$ with the quantum challenger $\mathsf{X}_Q$ in the security game as a joint computation in a Hilbert space $\mathcal{H}_{\mathcal{A}} \otimes \mathcal{H}_{\mathcal{A}\mathsf{X}} \otimes \mathcal{H}_{\mathsf{X}}$, where $\mathcal{H}_{\mathcal{A}}$ is attributed to algorithm $\mathcal{A}_Q$, $\mathcal{H}_{\mathsf{X}}$ to $\mathsf{X}_Q$, and $\mathcal{H}_{\mathcal{A}\mathsf{X}}$ is a joint space for communication. Formally, we thus decompose algorithm $\mathcal{A}_Q$ (and likewise $\mathsf{X}_Q$) into sequences $\mathcal{A}_Q^{(1)}, \mathcal{A}_Q^{(2)}, \mathcal{A}_Q^{(3)}, \ldots$ operating on $\mathcal{H}_{\mathcal{A}_Q} \otimes \mathcal{H}_{\mathcal{A}_Q \mathsf{X}_Q}$. See Figure 2. Note that we let the challenger access the following query of $\mathcal{A}_Q$ in each round for clarity, whereas we can assume that $\mathcal{A}_Q$ may already access all joint registers from the beginning.

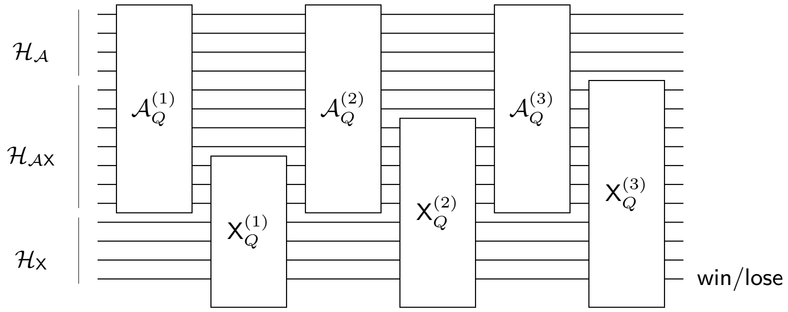

Fig. 2: Interactive computation of $\mathcal{A}_Q$ and $\mathsf{X}_Q$. Both algorithms receive $1^\lambda$ as input in their corresponding registers.

**Definition 3 (Quantum Security Game).** *A* quantum security game $\mathsf{G}_Q = (\mathsf{X}_Q, \mathsf{W}_Q)$ *against quantum adversary* $\mathcal{A}_Q$ *consists of a quantum challenger* $\mathsf{X}_Q$ *and a winning condition* $\mathsf{W}_Q$*. The game is parameterized by the security parameter* $\lambda$*. In this game, the quantum adversary* $\mathcal{A}_Q(1^\lambda)$ *interacts with the quantum challenger* $\mathsf{X}_Q(1^\lambda)$ *as depicted in Figure 2. The challenger eventually outputs a bit* win *or* lose*. We say that the adversary wins (loses) the game for security parameter* $\lambda$ *if the challenger's output is* win *(*lose*) when executing the game for parameter* $\lambda$*. We write*

$$\Pr^{\mathsf{G}_Q}_{\mathcal{A}_Q}(\lambda) := \Pr[\mathcal{A}_Q \textit{ wins game } \mathsf{G}_Q \textit{ for security parameter } \lambda].$$

In Figure 3, we give a protocol-based view of the interactive game between the parties. This notation will be supportive of understanding observational games.

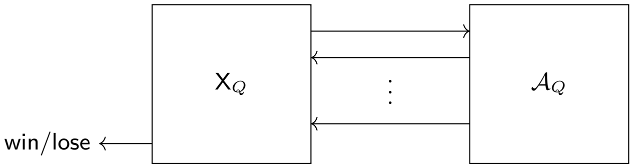

Fig. 3: Protocol-based view on game between $\mathcal{A}_Q$ and $\mathsf{X}_Q$. Both algorithms receive $1^\lambda$ as input (omitted here in the presentation).

### 3.2 Baseline Probabilities

We define the quantum version of dummy adversary following [16]. There, classical dummy adversaries are defined as being independent of the challenger's responses. We note that such adversaries can be considered even more simplistic than trivial adversaries whose output may depend on the challenger's behavior.

The example in [16] is a security game where the challenger sends a random string $s$ to the adversary, and the adversary wins the game if it responds with that string $s$. A trivial adversary duplicating the incoming message $s$ can easily win this game, while a dummy adversary —whose output does not depend on the challenger's message— only has a success probability of $2^{-|s|}$.

For the quantum case, we can translate the independence of the dummy adversary of the challenger by demanding that $\mathcal{A}_Q$ creates all game queries at the outset. The situation is depicted in Figure 4. Note that the challenger's processing steps remain unchanged. This is necessary as we consider fixed games with identical challengers; only the adversary may change.

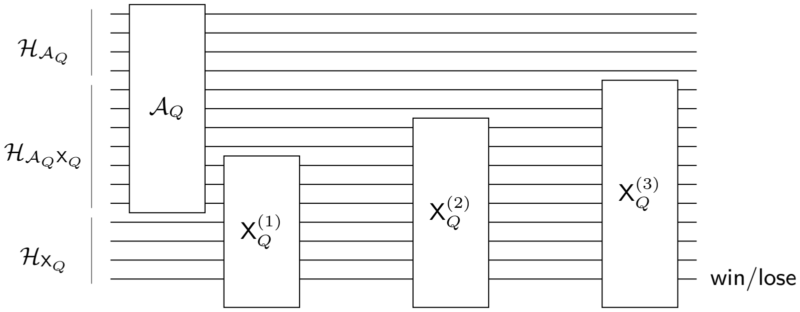
$\mathcal{H}_{\mathcal{A}_Q}$
$\mathcal{H}_{\mathcal{A}_Q \mathsf{X}_Q}$
$\mathcal{H}_{\mathsf{X}_Q}$
$\mathcal{A}_Q$
$\mathsf{X}_Q^{(1)}$
$\mathsf{X}_Q^{(2)}$
$\mathsf{X}_Q^{(3)}$
win/lose

Fig. 4: Interactive computation of dummy adversary $\mathcal{A}_Q$ with $\mathsf{X}_Q$

**Definition 4 (Quantum Dummy Adversary).** *In a quantum security game $\mathsf{G}_Q$ against quantum adversary $\mathcal{A}_Q$, the adversary is called a* quantum dummy adversary *if $\mathcal{A}_Q^{(i)}$ is the identity operation for all $i \geq 2$. That is, $\mathcal{A}_Q$ prepares all queries in advance, and the challenger processes all queries sequentially, eventually outputting* win *or* lose *(see Figure 4).*

By construction, every quantum dummy adversary is a quantum adversary, but in general, the converse does not hold.

We next define the baseline probability as the success probability of dummy adversaries of a particular time complexity bound $T_Q$.

**Definition 5 (Bounded Quantum Baseline Probability).** *A quantum baseline probability for time complexity $T_Q$ for a quantum security game $\mathsf{G}_Q$ is defined as*

$$\Pr^{\mathsf{G}_Q}_{\mathbf{D}[T_Q]}(\lambda) \; = \max_{\textit{dummy } \mathcal{A}_Q : T_{\mathcal{A}_Q} \leq T_Q} \Pr^{\mathsf{G}_Q}_{\mathcal{A}_Q}(\lambda),$$

*where the maximum is over all quantum dummy adversaries of time complexity $T_Q$. A quantum dummy adversary $\mathcal{A}_Q$ against a quantum security game $\mathsf{G}_Q$ is a quantum baseline adversary if it achieves the maximum, $\Pr^{\mathsf{G}_Q}_{\mathcal{A}_Q}(\lambda) = \Pr^{\mathsf{G}_Q}_{\mathbf{D}[T_Q]}(\lambda)$.*

We note that the maximum in the definition is a point-wise maximum over the security parameter $\lambda$. If we assume a fixed representation of quantum circuits, e.g., described via some universal gate set, then for given $\lambda$, run time bound $T_Q$ and game $\mathsf{G}$, the number of dummy adversaries $\mathcal{A}_Q$ with run time $T_{\mathcal{A}_Q} \leq T_Q$ is finite and the maximum thus reached. In the following, we use a similar line of reasoning for the point-wise minimum over quantum search algorithms, which we implicitly bound in run time by $T_Q \leq 2^{\lambda/2}$.

To normalize the advantage of an adversary $\mathcal{A}_Q$, running in time $T_{\mathcal{A}_Q}$, over the dummy adversary of the same complexity, we follow [16] and use the notion of the conventional advantage:

**Definition 6 (Quantum Conventional Advantage).** *We denote the quantum conventional advantage of $\mathcal{A}_Q$ for the game $\mathsf{G}_Q$ as follows:*

$$QAdv_{\mathcal{A}_Q}^{\mathsf{G}_Q}(\lambda) = \max\left\{\Pr_{\mathcal{A}_Q}^{\mathsf{G}_Q}(\lambda) - \Pr_{\mathbf{D}[T_Q]}^{\mathsf{G}_Q}(\lambda), 0\right\}.$$

### 3.3 Hybrid Observation Game

The *(advantage) observation game* in [16] captures the advantage of distinguishing the success of an adversary over a baseline adversary. The idea of the observation game is to have a meta-game between an adversary $\mathcal{B}$ and a "meta-challenger" $\hat{\mathsf{X}}$. Adversary $\mathcal{B}$ can ping the meta-challenger multiple times. Each time $\hat{\mathsf{X}}$ runs the game $\mathsf{G}$ with either the adversary $\mathcal{A}$ or with the baseline adversary $\mathcal{A}^{\text{dummy}}$, the choice depending on a secret random bit $b$ held by $\hat{\mathsf{X}}$. As in [16], we stipulate that the adversary $\mathcal{A}_Q$ has a success probability that is at least as large as the one of the baseline adversary (of the same complexity). The meta-challenger returns the game's outcome $\mathsf{win}$ or $\mathsf{lose}$ to adversary $\mathcal{B}$. Adversary $\mathcal{B}$ eventually needs to predict $b$ almost with certainty.

Based on the conceptual framework in [16], we propose a modification of the observation game to the quantum setting. In this setting, the game and the adversary may be quantum, and since we assume the execution of the game to be part of the meta-challenger, we also make the meta-challenger quantum. The observation algorithm $\mathcal{B}$, however, is still classical. The reason is that we are interested in the quantum adversary $\mathcal{A}_Q$ finding the classical key, i.e., the search game is supposed to use quantum power to solve a classical problem. We, therefore, call this type of observation game a hybrid game.

**Definition 7 (Hybrid Observation Game).** *Let $\mathcal{A}_Q$ be a quantum adversary of run-time $T_{\mathcal{A}_Q}$ for a quantum game $\mathsf{G}_Q = (\mathsf{X}_Q, \mathsf{W}_Q)$.*

*Consider a classical adversary $\mathcal{B}$ and a quantum meta-challenger $\hat{\mathsf{X}}_Q$ in the following hybrid observation game (see Figure 5):*

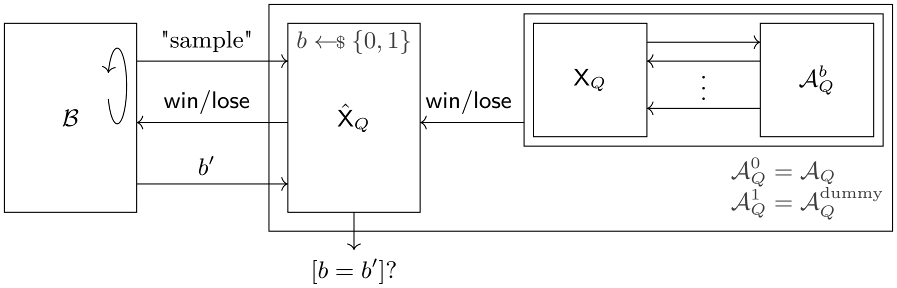
"sample"
win/lose
$b'$
$\mathcal{B}$
$b \leftarrow\!\!\$ \{0,1\}$
$\hat{\mathsf{X}}_Q$
win/lose
$\mathsf{X}_Q$
$\mathcal{A}_Q^b$
$\mathcal{A}_Q^0 = \mathcal{A}_Q$
$\mathcal{A}_Q^1 = \mathcal{A}_Q^{\text{dummy}}$
$[b = b']?$

Fig. 5: Protocol-based view on hybrid observation game between $\mathcal{B}$ and $\hat{\mathsf{X}}_Q$. Both algorithms receive $1^\lambda$ as input (omitted here in the presentation).

*Hybrid Observation Game*

**Challenge:** *At the beginning of the game the quantum meta-challenger $\hat{\mathsf{X}}_Q(1^\lambda)$ chooses a secret bit $b \leftarrow\!\!\$ \{0,1\}$. This secret bit $b$ determines in the following if we execute the game $\mathsf{G}_Q$ with the adversary $\mathcal{A}_Q^0 = \mathcal{A}_Q$ or with a quantum baseline adversary $\mathcal{A}_Q^1 = \mathcal{A}_Q^{dummy}$ of complexity $T_{\mathcal{A}_Q}$ for the game.*

**Sample Queries:** *Adversary $\mathcal{B}(1^\lambda)$ may repeatedly send the special symbol "sample" to $\hat{\mathsf{X}}_Q$. The meta-challenger then invokes the quantum game $\mathsf{G}_Q$ with $\mathcal{A}_Q^b$ for security parameter $\lambda$. When the game invocation eventually returns* win *or* lose*, then algorithm $\hat{\mathsf{X}}_Q$ returns the value to $\mathcal{B}$.*

**Decision:** *The adversary $\mathcal{B}$ eventually returns a prediction $b'$ for the secret bit. The meta-challenger $\hat{\mathsf{X}}_Q$ outputs* win *if $b = b'$ and $\Pr^{\mathsf{G}_Q}_{\mathcal{A}_Q}(\lambda) \geq \Pr^{\mathsf{G}_Q}_{\mathbf{D}[T_{\mathcal{A}_Q}]}(\lambda)$, and* lose *otherwise.*

*Remark 1.* The hybrid observation game can be viewed as a game $\hat{\mathsf{G}}_Q$ between the meta-challenger $\hat{\mathsf{X}}_Q$ and adversary $\mathcal{B}$, such that we can, for example, immediately transfer the notation $\Pr^{\hat{\mathsf{G}}_Q}_{\mathcal{B}}(\lambda)$ to describe the success probability of $\mathcal{B}$ in the observation game.

For the definition of bit security in our setting, we use the approach proposed by [34] which is also used in [16]. The starting point is Definition 7 of the hybrid observation game. We define the bit security as the cost to demonstrate the advantage of the $\mathcal{A}_Q$ against the quantum game $\mathsf{G}_Q$. This cost reflects the overall effort needed to repeatedly invoke $\mathcal{A}_Q$, enabling the classical algorithm $\mathcal{B}$ to win the hybrid observation game with sufficiently large probability. Precisely, we measure the total cost $T_{\mathcal{A}_Q}$ of $\mathcal{A}_Q$ multiplied by the query complexity $N_{\mathcal{B}}$ of $\mathcal{B}$ against the hybrid observation game. To distinguish this setting from the classical scenario and because it corresponds to quantum attacks on otherwise classical key derivation schemes, we call this the post-quantum bit security (PQBS).

**Definition 8 (Post-Quantum Bit Security as Cost to Demonstrate Advantage).** *For any quantum security game $\mathsf{G}_Q$ we define the* Post-Quantum Bit Security *(PQBS) with error probability* $0 < \delta(\lambda) < \frac{1}{2}$ *as follows:*

$$PQBS^{\mathsf{G}_Q,\delta}_{Dem}(\lambda) \;=\min_{\mathcal{A}_Q,\mathcal{B}}\{\log\big(T_{\mathcal{A}_Q}\cdot N_{\mathcal{B}}\big) \;:\; \Pr^{\hat{\mathsf{G}}_Q}_{\mathcal{B}}(\lambda) \geq 1-\delta(\lambda)\}.$$

### 3.4 Estimating the Post-Quantum Bit Security

Lee [16] relates the bit security, as defined above, via the Hellinger distance. Namely, the number of (classical) samples $N_\delta$ to distinguish two probability distributions $\mathcal{P}$ and $\mathcal{Q}$ with success probability $1-\delta$ for $0 < \delta < \frac{1}{2}$ is tightly bounded by the inverse square of the Hellinger distance of the two distributions:

$$\frac{1}{4\ln 2}\cdot\frac{\ln\Big(\frac{1}{4\delta(1-\delta)}\Big)}{d_{\mathrm{Hell}}(\mathcal{P},\mathcal{Q})^2} \leq N_\delta(\mathcal{P},\mathcal{Q}) \leq \frac{\ln\big(\frac{1}{2\delta}\big)}{d_{\mathrm{Hell}}(\mathcal{P},\mathcal{Q})^2}.$$

Here, the upper bound holds unconditionally and the lower bound holds if $d_{\mathrm{Hell}}(\mathcal{P},\mathcal{Q})^2 \leq \frac{1}{2}$. Since we consider a classical adversary $\mathcal{B}$ in our hybrid observation game, we can also apply this bound and derive the same conclusion as in [16].

**Theorem 1 (Estimation of Post-Quantum Bit Security [16]).** *For* $0 < \delta < \frac{(2-\sqrt{3})}{4}$*, we have the following estimation of the quantum demonstration bit security, up to a small additive error* $\alpha$*, satisfying* $0 \leq \alpha \leq 1 + \log\big(\ln\big(\frac{1}{2\delta}\big)\big)$ *:*

$$PQBS^{\mathsf{G}_Q,\delta}_{Dem}(\lambda) = \min_{\mathcal{A}_Q}\log\left(\ln\left(\frac{1}{2\delta}\right)\cdot\frac{T_{\mathcal{A}_Q}}{d_{Hell}(\Pr^{\mathsf{G}_Q}_{\mathcal{A}_Q}(\lambda),\Pr^{\mathsf{G}_Q}_{\mathbf{D}[T_{\mathcal{A}_Q}]}(\lambda))^2}\right) - \alpha,$$

*where the minimum is over all quantum adversaries* $\mathcal{A}_Q$ *satisfying* $\Pr^{\mathsf{G}_Q}_{\mathcal{A}_Q}(\lambda) \geq \Pr^{\mathsf{G}_Q}_{\mathbf{D}[T_{\mathcal{A}_Q}]}(\lambda)$.

For small $\delta$ this bound can be simplified in terms of the quantum Hellinger advantage:

**Definition 9 (Quantum Hellinger Advantage).** *In a quantum security game* $\mathsf{G}_Q$ *with a quantum adversary* $\mathcal{A}_Q$ *we denote the quantum Hellinger advantage as:*

$$QAdv^{\mathsf{G}_Q,\mathcal{A}_Q}_{Hell^2}(\lambda) = \begin{cases} d_{Hell}\big(\Pr^{\mathsf{G}_Q}_{\mathcal{A}_Q}(\lambda),\Pr^{\mathsf{G}_Q}_{\mathbf{D}[T_{\mathcal{A}_Q}]}(\lambda)\big)^2 & \text{if } \Pr^{\mathsf{G}_Q}_{\mathcal{A}_Q}(\lambda) \geq \Pr^{\mathsf{G}_Q}_{\mathbf{D}[T_{\mathcal{A}_Q}]}(\lambda), \\ 0 & \text{if } \Pr^{\mathsf{G}_Q}_{\mathcal{A}_Q}(\lambda) < \Pr^{\mathsf{G}_Q}_{\mathbf{D}[T_{\mathcal{A}_Q}]}(\lambda). \end{cases}$$

Then, we get a simplified approximation of post-quantum bit security:

**Definition 10 (Hellinger Post-Quantum Bit Security).** *Let $\mathsf{G}_Q$ be a quantum security game. Then, the Hellinger Post-Quantum Bit Security is defined as:*

$$PQBS^{\mathsf{G}_Q}_{Hell^2}(\lambda) = \min_{\mathcal{A}_Q} \log \left( \frac{T_{\mathcal{A}_Q}}{QAdv^{\mathsf{G}_Q,\mathcal{A}_Q}_{Hell^2}(\lambda)} \right).$$

# 4 Bit Security of Quantum Key Search

In this section, we examine the bit security within the context of the quantum key search game, where the adversary endeavors to find an unknown key for a distribution with a small statistical distance from the uniform distribution. One can view this game as reflecting the best adversarial strategy in settings where the deployment of the keys does not introduce collateral weaknesses. For example, one may think of a key used in a secure AES-based authenticated encryption scheme, such that the adversary can check a key guess via trial decryption. Still, neither the usage of AES nor of the authenticated encryption scheme is known to facilitate the key search.

We further evaluate the upper and lower bounds for the success probabilities associated with any adversary participating in the quantum key search game and we adopt here the definition 10. We adopt here the simplified version of bit security, defined via quantum Hellinger bit security:

$$\mathrm{PQBS}^{\mathsf{G}_Q}_{\mathrm{Hell}^2}(\lambda) = \min_{\mathcal{A}^{\mathrm{ks}}_Q} \log \left( \frac{T_{\mathcal{A}_Q}}{\mathrm{QAdv}^{\mathsf{G}_Q}_{\mathrm{Hell}^2}\mathcal{A}_Q(1^\lambda)} \right),$$

where the quantum Hellinger advantage is articulated through the Hellinger distance $d_{\mathrm{Hell}}\big(\mathrm{Pr}^{\mathsf{G}_Q}_{\mathcal{A}^{\mathrm{ks}}_Q}(\lambda), \mathrm{Pr}^{\mathsf{G}_Q}_{\mathbf{D}[T_{\mathcal{A}^{\mathrm{ks}}_Q}]}(\lambda)\big)^2$ for successful adversaries $\mathcal{A}_Q$. Within this framework, the derivation of an *upper* bound on post-quantum bit security requires a *lower* bound on the Hellinger distance, whereas deriving a *lower* bound on post-quantum bit security requires an *upper* bound on the Hellinger distance.

## 4.1 Quantum Key Search Game and Baseline Probability

The quantum key search game is parameterized by some distribution $\mu$, which assigns a probability to each of the possible $\lambda$-bit keys $\mathsf{k} \leftarrow\!\!\$\ \mathcal{K} = \{0,1\}^\lambda$. The distribution $\mu$ has a statistical distance $\Delta$ from the uniform distribution. The goal of the adversary is to find a randomly sampled key. For this, it can query the challenger in superposition about the candidate keys. Finally, it outputs a classical prediction for the secret key.

**Definition 11 (Quantum Key Search Game).** *The quantum key search game $\mathsf{G}^{ks,\mu,\Delta}_Q = (\mathsf{X}^{ks}_Q, \mathsf{W}^{ks}_Q)$ with quantum key search adversary $\mathcal{A}^{ks}_Q$ is defined as follows:*

*Quantum Key Search Game*

**Challenge:** *At the beginning of the game the quantum challenger* $\mathsf{X}_Q^{ks}(1^\lambda)$ *chooses a secret key* $\mathsf{k} \in \{0,1\}^\lambda$ *according to distribution* $\mu(\lambda)$.

**Queries:** *Adversary* $\mathcal{A}_Q^{ks}$ *can repeatedly query the challenger about any quantum state* $\sum \alpha_{xy} |x\rangle |y\rangle$ *to receive the answer* $\sum \alpha_{xy} |x\rangle |y \oplus I_\mathsf{k}(x)\rangle$, *where* $I_\mathsf{k}(x)$ *is the indicator function returning* $1$ *if and only if* $x = \mathsf{k}$.

**Decision:** *The adversary* $\mathcal{A}_Q^{ks}$ *eventually sends a prediction* $\mathsf{k}'$ *to the challenger* $\mathsf{X}_Q^{ks}$. *The challenger outputs* win *if* $\mathsf{k} = \mathsf{k}'$, *and* lose *otherwise.*

We consider the baseline probability for the key search game. Recall that dummy adversaries against the key search game must prepare all queries and their output at the beginning of the game. In particular, this means for our search game that the dummy adversary must compute its guess $\mathsf{k}'$ for the challenger's secret key $\mathsf{k}$ before it learns any information about the actual secret. This implies that the best strategy is to pick the most likely key, whose probability is bounded from above by $2^{-\lambda} + \mathit{\Delta}$. At the same time, at least one key must be hit with probability at least $2^{-\lambda}$, giving us a lower bound for the optimal dummy adversary. We conclude that

$$2^{-\lambda} \;\le\; \Pr^{\mathsf{G}_Q^{\mathrm{ks},\mu,\mathit{\Delta}}}_{\mathbf{D}[T_Q]}(\lambda) \;\le\; 2^{-\lambda} + \mathit{\Delta}.$$

This holds independently of the run time of the dummy adversary.

### 4.2 Bounds for Quantum Key Search Adversaries

We now turn to lower and upper bounds on the success probabilities of adversaries $\mathcal{A}_Q^{\mathrm{ks}}$ against the quantum key search game $\mathsf{G}_Q^{\mathrm{ks},\mu,\mathit{\Delta}}$. He et al. [13], optimizing Montanaro's algorithm [22], give a quantum search algorithm for *known* input distribution $p_i = \mathcal{P}(i)$ for $i = 1, 2, \ldots, N$. The idea of their search algorithm is to run Grover's algorithm but use an initial state $|s\rangle = \sum_{i=1}^{N} \sqrt{q_i}\,|i\rangle + \sqrt{q_0}\,|0\rangle$ for non-negative values $q_i$ and $q_0 = 1 - \sum_{i=1}^{N} q_i \ge 0$. The values $q_i$ can be chosen arbitrarily, yielding different bounds on the success probability after $T_Q$ steps of the search algorithm. Specifically, He et al. [13] show that the expected success probability $\mathsf{ESP}_{T_Q}(p,q)$ of finding the key in $T_Q$ steps when using $q$ as a starting vector as

$$\mathsf{ESP}_{T_Q}(p,q) = \sum_{i=1}^{N} p_i \cdot \sin^2((2T_Q + 1) \arcsin \sqrt{q_i}).$$

Each term $\sin^2((2T_Q+1) \cdot \arcsin \sqrt{q_i})$ reaches its maximum 1 for $q_i = \sin^2 \frac{\pi}{2(2T_Q+1)}$.

For our quantum key search problem, we assume that the probabilities $p_i$ are given in non-increasing order, i.e., that the adversary has already ordered the keys according to their likelihood (see Figure 1 on page 4).[2] Given a bound $T_Q$ on the run time, we now set the $q_i$ as

$$q_i = \begin{cases} \sin^2 \frac{\pi}{2(2T_Q+1)} & \text{if } i \leq T_Q^2 \\ 0 & \text{otherwise} \end{cases},$$

such that the sum over all $q_i$'s is bounded by 1 as required. Then, we get a lower bound for the expected success probability of

$$\mathsf{ESP}_{T_Q}(p,q) \geq \sum_{i=1}^{T_Q^2} p_i \cdot \sin^2((2T_Q+1)\arcsin\sqrt{q_i}) = \sum_{i=1}^{T_Q^2} p_i.$$

Since the $p_i$-values are in non-increasing order, the sum over the first $T_Q^2$ values can never be smaller than $T_Q^2 \cdot 2^{-\lambda}$. To see this, let $I$ be the index where the probability $p_i$ drops below $2^{-\lambda}$, i.e., where keys are less likely than uniformly distributed keys, and the curve of the actual key distribution lies below the curve of the uniform distribution in Figure 1. Then we can split the sum of the probabilities in the values up to index $I$ and after $I$ (if $I$ exceeds $T_Q^2$):

$$\sum_{i=1}^{T_Q^2} p_i \;=\; \sum_{i=1}^{\max\{I,T_Q^2\}} p_i \;+\; \sum_{i=I+1}^{T_Q^2} p_i.$$

Noting that the first sum equals the "uniform contribution" plus the statistical distance, and the second sum at least the "uniform contribution" minus the statistical distance, we conclude:

$$\begin{aligned} &\geq \max\{I,T_Q^2\} \cdot 2^{-\lambda} + \Delta \;+\; \min\{0, T_Q^2 - I\} \cdot 2^{-\lambda} - \Delta \\ &\geq T_Q^2 \cdot 2^{-\lambda}. \end{aligned}$$

He et al. [13] also show that their algorithm is optimal. That is, they show that any quantum search algorithm has a success probability of at most

$$\max_q \mathsf{ESP}_{T_Q}(p,q) \quad \text{ where } q_i \geq 0 \text{ and } \sum q_i \leq 1.$$

Letting $q^*(p)$ denote the optimal choice for $q$ given a specific distribution $p$,

$$q^*(p) = \arg\max_q \mathsf{ESP}_{T_Q}(p,q),$$

they argue that for any other distribution $\hat{p}$ we have

$$\mathsf{ESP}_{T_Q}(\hat{p}, q^*(\hat{p})) \geq \mathsf{ESP}_{T_Q}(\hat{p}, q^*(p))$$

[2] Computing this ordering may nonetheless be quite expensive in practice and significantly increase the run time of the adversary [33]. This depends on the concrete key generation algorithm and its specific manifestation of the statistical distance. We assume here, in favor of the adversary, that this step is inexpensive.

because $q^*(p)$ cannot yield a better probability than the maximizing value $q^*(\hat{p})$. Exploiting linearity in the probabilities, we conclude further

$$= \mathsf{ESP}_{T_Q}(p, q^*(p)) + \mathsf{ESP}_{T_Q}(\hat{p} - p, q^*(p))$$

Noting that the squared sin-values lie between 0 and 1 and that $\hat{p}_i - p_i \geq -|\hat{p}_i - p_i|$, we obtain the lower bound

$$\begin{aligned} &\geq \mathsf{ESP}_{T_Q}(p, q^*(p)) - \sum_{i=1}^{N} |\hat{p}_i - p_i| \\ &= \mathsf{ESP}_{T_Q}(p, q^*(p)) - 2 \cdot \mathit{\Delta} \end{aligned}$$

for the statistical distance $\mathit{\Delta}$ between $\hat{p}$ and $p$. It follows that we can bound the value $\mathsf{ESP}_{T_Q}(p, q^*(p))$ via the expected success probability for the *uniform* distribution $\hat{p}$ on the values and (twice) the statistical distance from uniform.

For the uniform distribution over all inputs and a single positive entry, we can apply the common upper bounds for any search algorithm [4] in the slightly optimized version of [36]: The probability $\rho$ that any quantum search algorithm for the uniform distribution on a set of $N$ elements (with a single matching element) is found after $T_Q$ steps, satisfies

$$4T_Q^2 \geq 2N - 2\sqrt{N}\sqrt{\rho} - 2\sqrt{N}\sqrt{N-1}\sqrt{1-\rho}.$$

Using $\sqrt{1-x} \leq 1 - \frac{x}{2}$, we obtain:

$$\begin{aligned} 4T_Q^2 &\geq 2N - 2\sqrt{N}\sqrt{\rho} - 2\sqrt{N}\sqrt{N-1}\sqrt{1-\rho} \\ &\geq 2N - 2\sqrt{N\rho} - 2N(1 - \frac{\rho}{2}) \\ &\geq N\rho - 2\sqrt{N\rho} \\ &= (\sqrt{N\rho} - 1)^2 - 1. \end{aligned}$$

Since $T_Q \geq 1$ we can write this as

$$5T_Q^2 \geq 4T_Q^2 + 1 \geq (\sqrt{N\rho} - 1)^2$$

Since we can assume that $\rho \geq \frac{1}{N}$ to surpass the dummy adversary, we have $N\rho \geq 1$ and the value $\sqrt{N\rho} - 1$ is thus non-negative. On the left-hand side, we also have $T_Q \geq 1$ such that we can take roots to obtain:

$$\sqrt{5}T_Q \geq \sqrt{N\rho} - 1$$

and thus $4T_Q \geq \sqrt{5}T_Q + 1 \geq \sqrt{N\rho}$ which implies

$$\frac{16T_Q^2}{N} \geq \rho.$$

We can now put the bounds together to obtain a sharp bound (up to constants and additive terms) on the adversary's success probability in relation to the number of steps $T_{\mathcal{A}_Q^{\mathrm{ks}}}$, the search space $N = 2^\lambda$, and the statistical distance $\Delta$ of the key distribution from uniform:

$$T^2_{\mathcal{A}_Q^{\mathrm{ks}}} \cdot 2^{-\lambda} \;\le\; \Pr{}^{\mathsf{G}_Q^{\mathrm{ks},\mu,\Delta}}_{\mathcal{A}_Q^{\mathrm{ks}}}(\lambda) \;\le\; 16 T^2_{\mathcal{A}_Q^{\mathrm{ks}}} \cdot 2^{-\lambda} + 2 \cdot \Delta$$

The upper bound holds for adversaries $\mathcal{A}_Q^{\mathrm{ks}}$ with success probability at least $2^{-\lambda}$, which equals the dummy adversary's success probability.

### 4.3 Upper Bound for the Post-Quantum Bit Security

We give the upper bound of

$$\mathrm{PQBS}^{\mathsf{G}_Q}_{\mathrm{Hell}^2}(\lambda) = \min_{\mathcal{A}_Q^{\mathrm{ks}}} \log\left(\frac{T_{\mathcal{A}_Q}}{\mathrm{QAdv}^{\mathsf{G}_Q,\mathcal{A}_Q}_{\mathrm{Hell}^2}(\lambda)}\right)$$

in terms of the given statistical distance $\Delta$. Since the bit security is defined over the minimum over all adversaries, we may pick a specific adversary and its run time in dependency of $\Delta$ to get an upper bound. Specifically, we assume that $T_{\mathcal{A}_Q^{\mathrm{ks}}} \ge 48$. For the bounds, we also need that $T_{\mathcal{A}_Q^{\mathrm{ks}}} \le 2^{\lambda/2}$, which matches the known upper bound on the run time of Grover's search algorithm.

For the statistical distance, we stipulate for technical reasons that $\Delta \le \frac{1}{48^2}$. In other words, we do not consider constant statistical distances $\Delta > \frac{1}{48^2}$. This is not a major restriction since constant distances are considered insecure in the first place, because then a small set of keys, even a single key, can have a constant probability of being chosen by the key generation algorithm. Technically, this may also imply that the success probabilities $2^{-\lambda} + \Delta$ of the dummy adversary can be quite very large. With the assumption about $\Delta$, we can conclude that $\Delta \le \frac{1}{48^2} \cdot T^2_{\mathcal{A}_Q^{\mathrm{ks}}} \cdot 2^{-\lambda}$. This can always be accomplished by picking $T_{\mathcal{A}_Q^{\mathrm{ks}}}$ as its maximum $2^{\lambda/2}$. We can therefore write the distance $\Delta$ for factor $\gamma \le \frac{1}{48^2}$ as

$$\Delta = \gamma \cdot T^2_{\mathcal{A}_Q^{\mathrm{ks}}} \cdot 2^{-\lambda}.$$

To derive an upper bound on the bit security we need a lower bound on the Hellinger distance. Recall that for binary probability distributions $\mathcal{P}$, $\mathcal{Q}$, as in the case of our key search game, the square of the Hellinger distance can be written as

$$d_{\mathrm{Hell}}(\mathcal{P},\mathcal{Q})^2 = 1 - \sqrt{\epsilon_{\mathcal{P}} \cdot \epsilon_{\mathcal{Q}}} - \sqrt{(1-\epsilon_{\mathcal{P}}) \cdot (1-\epsilon_{\mathcal{Q}})},$$

where $\epsilon_{\mathcal{P}} = \mathcal{P}(1)$ and $\epsilon_{\mathcal{Q}} = \mathcal{Q}(1)$ denote the success probabilities. In our setting, the distributions are given by an arbitrary quantum key search adversary $\mathcal{A}_Q^{\mathrm{ks}}$ playing the game $\mathsf{G}_Q^{\mathrm{ks},\mu,\Delta}$, and a baseline dummy adversary $\mathcal{A}_Q^{\mathrm{dummy,ks}}$ in the game. Hence, to give an upper bound on the bit security—and therefore a lower bound on the squared Hellinger distance—we plug in upper bounds for $\epsilon_{\mathcal{P}}, \epsilon_{\mathcal{Q}}$ in the first term (taking into account the negative sign) and lower bounds for $\epsilon_{\mathcal{P}}, \epsilon_{\mathcal{Q}}$ in the second term (taking into account the inner and outer negative sign).

Recall that we have $\epsilon_{\mathcal{P}} \leq\ 16T^2_{\mathcal{A}_Q^{\mathrm{ks}}} \cdot 2^{-\lambda} + 2\mathit{\Delta}$ for the key search adversaries and $\epsilon_{\mathcal{Q}} \leq 2^{-\lambda} + \mathit{\Delta}$ for any dummy adversary. Similarly, we have $\epsilon_{\mathcal{P}} \geq T^2_{\mathcal{A}_Q^{\mathrm{ks}}} \cdot 2^{-\lambda}$ for the key search adversaries and $\epsilon_{\mathcal{Q}} \geq 2^{-\lambda}$ for dummy adversaries. This yields:

$$\begin{aligned}
&d_{\mathrm{Hell}}\left(\Pr^{\mathsf{G}_Q}_{\mathcal{A}_Q}(\lambda), \Pr^{\mathsf{G}_Q}_{\mathbf{D}[T_{\mathcal{A}_Q^{\mathrm{ks}}}]}(\lambda)\right)^2 \\
&\quad \geq 1 - \sqrt{(16T^2_{\mathcal{A}_Q^{\mathrm{ks}}} \cdot 2^{-\lambda} + 2 \cdot \mathit{\Delta}) \cdot (2^{-\lambda} + \mathit{\Delta})} - \sqrt{(1 - T^2_{\mathcal{A}_Q^{\mathrm{ks}}} \cdot 2^{-\lambda}) \cdot (1 - 2^{-\lambda})} \\
&\quad \geq 1 - \sqrt{16T^2_{\mathcal{A}_Q^{\mathrm{ks}}} \cdot 2^{-2\lambda} + 2 \cdot \mathit{\Delta}^2 + 16T^2_{\mathcal{A}_Q^{\mathrm{ks}}} \cdot 2^{-\lambda} \cdot \mathit{\Delta} + 2^{-\lambda+1} \cdot \mathit{\Delta}} \\
&\qquad\qquad - \sqrt{1 - 2^{-\lambda} - T^2_{\mathcal{A}_Q^{\mathrm{ks}}} \cdot 2^{-\lambda} + T^2_{\mathcal{A}_Q^{\mathrm{ks}}} \cdot 2^{-2\lambda}}.
\end{aligned}$$

We start to simplify the second square root by dropping the negative term $-2^{-\lambda}$ and bounding the term $T^2_{\mathcal{A}_Q^{\mathrm{ks}}} \cdot 2^{-2\lambda} \leq \frac{1}{2} \cdot T^2_{\mathcal{A}_Q^{\mathrm{ks}}} \cdot 2^{-\lambda}$ for $\lambda \geq 1$. The gives:

$$\sqrt{1 - 2^{-\lambda} - T^2_{\mathcal{A}_Q^{\mathrm{ks}}} \cdot 2^{-\lambda} + T^2_{\mathcal{A}_Q^{\mathrm{ks}}} \cdot 2^{-2\lambda}} \leq \sqrt{1 - \tfrac{1}{2}T^2_{\mathcal{A}_Q^{\mathrm{ks}}} \cdot 2^{-\lambda}}.$$

Using the bound $(1+x)^r \leq 1 + rx$ (Bernoulli's inequality) for $r \in [0,1]$ and for $x \geq -1$ (which holds in our case because we presume $T_{\mathcal{A}_Q^{\mathrm{ks}}} \leq 2^{\lambda/2}$) we thus obtain

$$\sqrt{1 - \tfrac{1}{2}T^2_{\mathcal{A}_Q^{\mathrm{ks}}} \cdot 2^{-\lambda}} \leq 1 - \tfrac{1}{4}T^2_{\mathcal{A}_Q^{\mathrm{ks}}} \cdot 2^{-\lambda}.$$

Hence, the Hellinger distance is lower bounded by

$$\begin{aligned}
&d_{\mathrm{Hell}}\left(\Pr^{\mathsf{G}_Q}_{\mathcal{A}_Q}(\lambda), \Pr^{\mathsf{G}_Q}_{\mathbf{D}[T_{\mathcal{A}_Q^{\mathrm{ks}}}]}(\lambda)\right)^2 \\
&\quad \geq \tfrac{1}{4}T^2_{\mathcal{A}_Q^{\mathrm{ks}}} \cdot 2^{-\lambda} - \sqrt{16T^2_{\mathcal{A}_Q^{\mathrm{ks}}} \cdot 2^{-2\lambda} + 2 \cdot \mathit{\Delta}^2 + 16T^2_{\mathcal{A}_Q^{\mathrm{ks}}} \cdot 2^{-\lambda} \cdot \mathit{\Delta} + 2^{-\lambda+1} \cdot \mathit{\Delta}}.
\end{aligned}$$

Next, we bound the remaining square root. We make a case distinction:

- Case $\mathit{\Delta} \leq 2^{-\lambda}$. In this case, we can upper-bound the square root expression as

$$\begin{aligned}
&\sqrt{16T^2_{\mathcal{A}_Q^{\mathrm{ks}}} \cdot 2^{-2\lambda} + 2 \cdot \mathit{\Delta}^2 + 16T^2_{\mathcal{A}_Q^{\mathrm{ks}}} \cdot 2^{-\lambda} \cdot \mathit{\Delta} + 2^{-\lambda+1} \cdot \mathit{\Delta}} \\
&\quad \leq \sqrt{16T^2_{\mathcal{A}_Q^{\mathrm{ks}}} \cdot 2^{-2\lambda} + 2 \cdot 2^{-2\lambda} + 16T^2_{\mathcal{A}_Q^{\mathrm{ks}}} \cdot 2^{-2\lambda} + 2^{-2\lambda+1}} \\
&\quad \leq \sqrt{32T^2_{\mathcal{A}_Q^{\mathrm{ks}}} \cdot 2^{-2\lambda} + 4 \cdot 2^{-2\lambda}} \\
&\quad \leq \sqrt{36T^2_{\mathcal{A}_Q^{\mathrm{ks}}} \cdot 2^{-2\lambda}} \\
&\quad = 6T_{\mathcal{A}_Q^{\mathrm{ks}}} \cdot 2^{-\lambda}
\end{aligned}$$

For $T_{\mathcal{A}_Q^{\mathrm{ks}}} \geq 48$ we thus obtain

$$d_{\mathrm{Hell}}\left(\Pr^{\mathsf{G}_Q}_{\mathcal{A}_Q}(\lambda), \Pr^{\mathsf{G}_Q}_{\mathbf{D}[T_{\mathcal{A}_Q^{\mathrm{ks}}}]}(\lambda)\right)^2 \geq \tfrac{1}{4}T^2_{\mathcal{A}_Q^{\mathrm{ks}}} \cdot 2^{-\lambda} - 6T_{\mathcal{A}_Q^{\mathrm{ks}}} \cdot 2^{-\lambda} \geq \tfrac{1}{8}T^2_{\mathcal{A}_Q^{\mathrm{ks}}} \cdot 2^{-\lambda}.$$

Therefore, the upper bound for the post-quantum Hellinger bit security $\mathrm{PQBS}^{\mathsf{G}_Q}_{\mathrm{Hell}^2}(\lambda)$ becomes:

$$\mathrm{PQBS}^{\mathsf{G}_Q}_{\mathrm{Hell}^2}(\lambda) \leq \min_{\mathcal{A}_Q^{\mathrm{ks}}} \log \frac{T_{\mathcal{A}_Q^{\mathrm{ks}}}}{\frac{1}{8}T^2_{\mathcal{A}_Q^{\mathrm{ks}}} \cdot 2^{-\lambda}} \leq \min_{\mathcal{A}_Q^{\mathrm{ks}}}(\lambda - \log T_{\mathcal{A}_Q^{\mathrm{ks}}} + 3).$$

- Case $\Delta > 2^{-\lambda}$. In this case, we can bound the square root expression due to $\Delta > 2^{-\lambda}$ and $\Delta \leq \frac{1}{48^2}T^2_{\mathcal{A}_Q^{\mathrm{ks}}} \cdot 2^{-\lambda}$ as follows:

$$\begin{aligned}
&\sqrt{16T^2_{\mathcal{A}_Q^{\mathrm{ks}}} \cdot 2^{-2\lambda} + 2 \cdot \Delta^2 + 16T^2_{\mathcal{A}_Q^{\mathrm{ks}}} \cdot 2^{-\lambda} \cdot \Delta + 2^{-\lambda+1} \cdot \Delta} \\
&\quad \leq \sqrt{16T^2_{\mathcal{A}_Q^{\mathrm{ks}}} \cdot 2^{-\lambda} \cdot \Delta + 2 \cdot \frac{1}{48^2}T^2_{\mathcal{A}_Q^{\mathrm{ks}}} \cdot 2^{-\lambda} \cdot \Delta + 16T^2_{\mathcal{A}_Q^{\mathrm{ks}}} \cdot 2^{-\lambda} \cdot \Delta + 2 \cdot 2^{-\lambda} \cdot \Delta} \\
&\quad \leq \sqrt{36T^2_{\mathcal{A}_Q^{\mathrm{ks}}} \cdot 2^{-\lambda} \cdot \Delta} \\
&\quad = 6T_{\mathcal{A}_Q^{\mathrm{ks}}} \cdot 2^{-\lambda/2} \cdot \sqrt{\Delta}
\end{aligned}$$

We thus get from the condition $\Delta = \gamma \cdot T^2_{\mathcal{A}_Q^{\mathrm{ks}}} \cdot 2^{-\lambda}$ that this is equal to

$$= 6T_{\mathcal{A}_Q^{\mathrm{ks}}} \cdot 2^{-\lambda/2} \cdot \sqrt{\gamma} \cdot T_{\mathcal{A}_Q^{\mathrm{ks}}} \cdot 2^{-\lambda/2} = 6\sqrt{\gamma} \cdot T^2_{\mathcal{A}_Q^{\mathrm{ks}}} \cdot 2^{-\lambda}.$$

Plugging this in, and using that $\sqrt{\gamma} \leq \frac{1}{48}$, we thus get

$$d_{\mathrm{Hell}}\left(\Pr^{\mathsf{G}_Q}_{\mathcal{A}_Q}(\lambda), \Pr^{\mathsf{G}_Q}_{\mathbf{D}[T_{\mathcal{A}_Q^{\mathrm{ks}}}]}(\lambda)\right)^2 \geq \tfrac{1}{4}T^2_{\mathcal{A}_Q^{\mathrm{ks}}} \cdot 2^{-\lambda} - 6\sqrt{\gamma} \cdot T^2_{\mathcal{A}_Q^{\mathrm{ks}}} \cdot 2^{-\lambda} \geq \tfrac{1}{8}T^2_{\mathcal{A}_Q^{\mathrm{ks}}} \cdot 2^{-\lambda}.$$

This gives us, once more, the upper bound on the post-quantum Hellinger bit security:

$$\mathrm{PQBS}^{\mathsf{G}_Q}_{\mathrm{Hell}^2}(\lambda) \leq \min_{\mathcal{A}_Q^{\mathrm{ks}}} \log \frac{T_{\mathcal{A}_Q^{\mathrm{ks}}}}{\frac{1}{8}T^2_{\mathcal{A}_Q^{\mathrm{ks}}} \cdot 2^{-\lambda}} \leq \min_{\mathcal{A}_Q^{\mathrm{ks}}}(\lambda - \log T_{\mathcal{A}_Q^{\mathrm{ks}}} + 3).$$

### 4.4 Lower Bound for the Post-Quantum Bit Security

To determine a lower bound on post-quantum bit security, yielding an upper bound on the squared Hellinger distance, we use the upper bound of the squared Hellinger distance in terms of the statistical distance:

$$d_{\mathrm{Hell}}(\mathcal{P}, \mathcal{Q})^2 \leq d_{\mathrm{TV}}(\mathcal{P}, \mathcal{Q}).$$

The statistical distance for the binary random variables $\mathcal{P}$ and $\mathcal{Q}$ is given by $|\epsilon_{\mathcal{P}} - \epsilon_{\mathcal{Q}}|$ which, in our case, can be upper bounded by

$$d_{\mathrm{TV}}(\mathcal{P}, \mathcal{Q}) \le 16T^2_{\mathcal{A}_Q^{\mathrm{ks}}} \cdot 2^{-\lambda} + 2 \cdot \mathit{\Delta} - 2^{-\lambda} \le 16T^2_{\mathcal{A}_Q^{\mathrm{ks}}} \cdot 2^{-\lambda} + 2 \cdot \mathit{\Delta}.$$

We make a case distinction:

- Case $\mathit{\Delta} \le T^2_{\mathcal{A}_Q^{\mathrm{ks}}} \cdot 2^{-\lambda}$. Then we conclude that

  $$d_{\mathrm{Hell}}(\mathcal{P}, \mathcal{Q})^2 \le \mathit{\Delta}(\mathcal{P}, \mathcal{Q}) \le 18T^2_{\mathcal{A}_Q^{\mathrm{ks}}} \cdot 2^{-\lambda},$$

  such that we derive a lower bound for the post-quantum bit security as

  $$\mathrm{PQBS}^{\mathsf{G}_Q}_{\mathrm{Hell}^2}(\lambda) \ge \min_{\mathcal{A}_Q^{\mathrm{ks}}} \log \frac{T_{\mathcal{A}_Q^{\mathrm{ks}}}}{18T^2_{\mathcal{A}_Q^{\mathrm{ks}}} \cdot 2^{-\lambda}} \ge \min_{\mathcal{A}_Q^{\mathrm{ks}}}(\lambda - \log T_{\mathcal{A}_Q^{\mathrm{ks}}} - 5).$$

- Case $\mathit{\Delta} > T^2_{\mathcal{A}_Q^{\mathrm{ks}}} \cdot 2^{-\lambda}$. In this case the dominating term in the upper bound $16T^2_{\mathcal{A}_Q^{\mathrm{ks}}} \cdot 2^{-\lambda} + 2 \cdot \mathit{\Delta}$ is the statistical distance $\mathit{\Delta}$ and we can conclude that

  $$\mathrm{PQBS}^{\mathsf{G}_Q}_{\mathrm{Hell}^2}(\lambda) \ge \min_{\mathcal{A}_Q^{\mathrm{ks}}} \log \frac{T_{\mathcal{A}_Q^{\mathrm{ks}}}}{18\mathit{\Delta}} \ge \min_{\mathcal{A}_Q^{\mathrm{ks}}}(\log T_{\mathcal{A}_Q^{\mathrm{ks}}} - \log \mathit{\Delta} - 5).$$

  If we rewrite $\mathit{\Delta}$ as $\mathit{\Delta} = \gamma \cdot T^2_{\mathcal{A}_Q^{\mathrm{ks}}} \cdot 2^{-\lambda}$ for factor $\gamma > 1$ then we derive

  $$\mathrm{PQBS}^{\mathsf{G}_Q}_{\mathrm{Hell}^2}(\lambda) \ge \min_{\mathcal{A}_Q^{\mathrm{ks}}}(\lambda - \log T_{\mathcal{A}_Q^{\mathrm{ks}}} - \log \gamma - 5).$$

  Hence, we only get a slightly worse bound, losing $\log \gamma$ additional bits. Note that $\gamma > 1$ in this case here, such that we indeed reduce the lower bound for post-quantum bit security compared to the other case.

## 5 Interpretation of Results

We discuss here the significance of our results in light of the potential choices of the statistical distance $\mathit{\Delta}$. We note that we technically require that $\mathit{\Delta} \le \frac{1}{48^2}$ as well as $48 \le T_{\mathcal{A}_Q^{\mathrm{ks}}} \le 2^{\lambda/2}$ for the upper bounds. For the lower bounds, we usually also make the assumption that any adversary, over which we minimize, is bounded by run time $2^{\lambda/2}$, as quantum key search already succeeds with probability close to 1 in this case.

### 5.1 The Case $\boldsymbol{\mathit{\Delta} \le 2^{-\lambda}}$

If we choose the statistical to be very small, $\mathit{\Delta} \le 2^{-\lambda}$, we have a matching upper and lower bound. The upper bound tells us

$$\mathrm{PQBS}^{\mathsf{G}_Q}_{\mathrm{Hell}^2}(\lambda) \le \min_{\mathcal{A}_Q^{\mathrm{ks}}}(\lambda - \log T_{\mathcal{A}_Q^{\mathrm{ks}}} + 3),$$

and the lower bound for $\varDelta \leq 2^{-\lambda} \leq T^2_{\mathcal{A}^{\mathrm{ks}}_Q} \cdot 2^{-\lambda}$ says

$$\mathrm{PQBS}^{\mathsf{G}_Q}_{\mathrm{Hell}^2}(\lambda) \geq \min_{\mathcal{A}^{\mathrm{ks}}_Q}(\lambda - \log T_{\mathcal{A}^{\mathrm{ks}}_Q} - 5).$$

Hence, both bounds match up to a constant number of bits. Furthermore, if we consider adversaries with run time $T_{\mathcal{A}^{\mathrm{ks}}_Q} = \varTheta(2^{\lambda/2})$, then we can conclude that we get a bit security of approximately $\lambda - \lambda/2 = \lambda/2$. This matches the known expectations for the quadratic speed-up for search with quantum computers, e.g., a uniform 256-bit AES key gives bit security of 128. Our result here confirms that this is indeed the case according to formal models and still holds if the statistical distance of the key distribution from uniform is $\varDelta \leq 2^{-\lambda}$.

Remarkably, our result also shows that decreasing the statistical distance further, e.g., to $\varDelta \leq 2^{-2\lambda}$, does not yield any advantage in terms of bit security. Our lower bound indicates that this is not known to increase bit security since the bound is independent of the statistical distance $\varDelta$ in case $\varDelta \leq 2^{-\lambda}$. Our upper bound proves that this is not only due to a loose lower bound but that the bit security, indeed, cannot increase.

### 5.2 The Case $2^{-\lambda} \leq \varDelta \leq 2^{-\lambda/2}$

The quadratic speed-up in quantum search halves the bit security. A natural idea is then to also allow for a larger statistical distance, say, $\varDelta = 2^{-\lambda/2}$, which reduces, for example, the number of truncated bits for privacy amplification. In this case, the upper bounds remain unchanged. For the lower bound, we then have $2^{-\lambda/2} \geq \varDelta = \gamma \cdot T^2_{\mathcal{A}^{\mathrm{ks}}_Q} \cdot 2^{-\lambda}$ and therefore $\gamma \leq T^{-2}_{\mathcal{A}^{\mathrm{ks}}_Q} \cdot 2^{\lambda/2}$. This gives us

$$\begin{aligned}\mathrm{PQBS}^{\mathsf{G}_Q}_{\mathrm{Hell}^2}(\lambda) &\geq \min_{\mathcal{A}^{\mathrm{ks}}_Q}(\lambda - \log T_{\mathcal{A}^{\mathrm{ks}}_Q} + 2\log T_{\mathcal{A}^{\mathrm{ks}}_Q} - \lambda/2 - 5)\\ &= \min_{\mathcal{A}^{\mathrm{ks}}_Q}(\lambda/2 + \log T_{\mathcal{A}^{\mathrm{ks}}_Q} - 5).\end{aligned}$$

Hence we get a lower bound of at least $\lambda/2$ bits, matching the upper bound, if we consider $T_{\mathcal{A}^{\mathrm{ks}}_Q} = 1$.

Note that the run time $T_{\mathcal{A}^{\mathrm{ks}}_Q}$ for the upper and lower bound may differ. For the upper bound, any choice of $T_{\mathcal{A}^{\mathrm{ks}}_Q}$ gives an upper limit for $\mathrm{PQBS}^{\mathsf{G}_Q}_{\mathrm{Hell}^2}(\lambda)$ which is defined over the minimum over all adversaries. For the lower bound, we therefore have to consider all adversaries with all possible choices of $T_{\mathcal{A}^{\mathrm{ks}}_Q}$ and then take the minimum of all these values. In the example above we would thus take $T_{\mathcal{A}^{\mathrm{ks}}_Q} = 1$ for the minimum, whereas we choose $T_{\mathcal{A}^{\mathrm{ks}}_Q} = 2^{\lambda/2}$ for the upper bound.

It may now seem as if the statistical distance $\varDelta = 2^{-\lambda/2}$ should be preferable over $\varDelta = 2^{-\lambda}$ since it yields the same bound at a relaxed requirement. This interpretation, however, crucially relies on the estimated global upper bound of $T_{\mathcal{A}^{\mathrm{ks}}_Q}$. Above, we assume that this is at most $2^{\lambda/2}$. If one instead assumes that the best quantum algorithm can make at most $2^{\lambda/8}$ steps, e.g., because of engineering constraints, then we obtain a different picture. In this case, the upper bound would be in the order of $\lambda - \lambda/8 = 7\lambda/8$, whereas the lower bound would still be $\lambda/2$. This is because we get the negative term $\log\gamma$ in the lower bound in this case here, unless in the first case of $\Delta \leq 2^{-\lambda}$, and this value $\log\gamma$ can be quite significant if the run time $T_{\mathcal{A}_Q^{\mathrm{ks}}}$ is smaller.

In summary, a conservative choice, which works independently of some upper bound on $T_{\mathcal{A}_Q^{\mathrm{ks}}}$ beyond $2^{\lambda/2}$, is to set $\Delta = 2^{-\lambda}$.

### 5.3 The Case $\boldsymbol{\Delta > 2^{-\lambda/2}}$

Increasing the statistical distance even beyond $2^{-\lambda/2}$ may be tempting. Recall that we have $\Delta = \gamma \cdot T^2_{\mathcal{A}_Q^{\mathrm{ks}}} \cdot 2^{-\lambda}$ and that the lower bound tells us

$$\mathrm{PQBS}^{\mathsf{G}_Q}_{\mathrm{Hell}^2}(\lambda) \geq \min_{\mathcal{A}_Q^{\mathrm{ks}}}(\lambda - \log T_{\mathcal{A}_Q^{\mathrm{ks}}} - \log\gamma - 5).$$

We may pick $T_{\mathcal{A}_Q^{\mathrm{ks}}} = 1$ such that $\gamma = \Delta \cdot 2^{\lambda}$ and $\log\gamma = \lambda + \log\Delta$; we discuss below that this indeed minimizes the bound. Plugging this into the inequality, we derive

$$\mathrm{PQBS}^{\mathsf{G}_Q}_{\mathrm{Hell}^2}(\lambda) \geq -\log\Delta - 5.$$

If the statistical distance $\Delta$ now goes significantly above the bound $2^{-\lambda/2}$, then this yields a notably decreased lower bound. For example, if we choose $\Delta = 2^{-\lambda/4}$, we only get a guarantee of $\lambda/4$ bits of security. In this case, however, our upper bound is not tight. This leaves the possibility that the actual bit security may be higher.

We note that picking $T_{\mathcal{A}_Q^{\mathrm{ks}}} = 1$ and thus adversaries which only test a single key is not surprising in this context. In the worst case, only one key may encompass the entire (large) statistical distance $\Delta$ such that testing only this key may be a valid strategy. In terms of our bound, both negative terms $\log T_{\mathcal{A}_Q^{\mathrm{ks}}}$ and $\log\gamma$ enter linearly into the bound for bit security. Yet, in $\Delta = \gamma \cdot T^2_{\mathcal{A}_Q^{\mathrm{ks}}} \cdot 2^{-\lambda}$ for given $\Delta$ the run time enters quadratically. We thus maximize the loss in bits by setting $T_{\mathcal{A}_Q^{\mathrm{ks}}} = 1$ and picking $\gamma$ as large as possible.

## 6 Implications to Quantum Key Distribution

Several works in the area of quantum key distribution investigate the security bounds from a cryptographic point of view [25,23,30,21,31,5,24,18,26]. These works usually divide the error $\varepsilon = \varepsilon_{\mathrm{correct}} + \varepsilon_{\mathrm{secure}}$ of the overall protocol into an error $\varepsilon_{\mathrm{correct}}$ for the two parties not arriving at the same key, and an error $\varepsilon_{\mathrm{secure}}$ for an adversary learning information about the secret. Our work does not consider the former type of error because we only investigate key generation as a monolithic inner process. The latter error somewhat corresponds to our notion of statistical distance: In particular, the final step in the protocols consists of the privacy amplification step, resulting in a close-to-uniform key. In this step, the works [25,23,31,5,24,26] consider the trace distance as the quantum analog to the statistical distance; still, the trace and statistical distances are tightly related [24].[3]

One can argue if one should combine the errors for correctness and secrecy within a single parameter $\varepsilon = \varepsilon_{\text{correct}} + \varepsilon_{\text{secure}}$. If the two parties do not derive the same key, then it may be easy to check for the parties, e.g., by running a key confirmation protocol before any sensitive data is transmitted. A small but non-negligible error $\varepsilon_{\text{correct}}$ of, say, $10^{-6}$ may be acceptable, forcing the parties to restart the execution occasionally. This issue has also been pointed out in [30], but that work still uses the sum to discuss example figures.

Conversely, secrecy is generally not verifiable, i.e., the parties cannot easily determine that the adversary can learn more information than desired. A smaller value $\varepsilon_{\text{secure}} \ll \varepsilon_{\text{correct}}$ may thus be preferable. In the combined sum, however, the value $\varepsilon_{\text{correct}}$ then overshadows the other term, suggesting that $\varepsilon_{\text{secure}}$ could be chosen close to $10^{-6}$, too. Some practical demonstrations like [3] confirm the use such a choice with $\varepsilon_{\text{secure}} = \varepsilon_{\text{correct}} = 10^{-9}$.

Indeed, the mix-up of the correctness and secrecy parameters also makes it hard to interpret the suggested figures for $\varepsilon$ in [30,21,31,5,18,26] with respect to the choice for $\varepsilon_{\text{secret}}$. The works [30,31,18] give a bound of $\varepsilon = 10^{-10}$. Mizutani et al. [21] set $\varepsilon$ to be $10^{-8}$ or $10^{-10}$. Renner and Wolf [26] list typical values for parameter $\varepsilon$ in the range of $10^{-6}$ to $10^{-12}$. The work by Bunandar et al. [5] states explicitly in their example that $\varepsilon_{\text{correct}} = 10^{-15}$ and $\varepsilon_{\text{secure}} = 10^{-10}$. Müller-Quade and Renner [23] explicitly mention $\varepsilon_{\text{secure}} = 10^{-10}$ as an example instantiation. We are unaware of the origin of these figures, e.g., if they are based on values used in practice, recommendations, or merely provide numerical examples. It is also unclear whether these figures have been evaluated against a concrete security goal, like the key search game in our work here. Remarkably, the work by Zhang et al. [37] nonetheless states for instance that a choice of $\varepsilon = 10^{-5}$ "is considered to be realistic for cryptography applications."

One may now interpret the suggested parameter $\varepsilon = 10^{-10}$ in [30,21,31,18,26] to use $\varepsilon_{\text{secure}} \approx \varepsilon_{\text{correct}}$, approximately also matching the concrete suggestions in [23,5] for $\varepsilon_{\text{secure}}$. However, this parameter selection appears optimistic in light of our results regarding bit security. Even for a short 256-bit AES key of quantum security level 128, choosing, for example, $\varepsilon_{\text{secure}} \approx 10^{-12} \approx 2^{-40}$ may be insufficient. According to our results, a more conservative choice with a significantly smaller value is advisable unless one has further information about the actual distributions.

If, on the other hand, one assumes that the aforementioned works suggest using significantly smaller $\varepsilon_{\text{secure}} \ll \varepsilon_{\text{correct}}$ for $\varepsilon = 10^{-10}$, then they leave open

[3] We note that some works in this domain also define a third criterion, robustness (see, for example, [23,24]). This property can be roughly described as correctness in the absence of an adversary, i.e., noise resilience. In cryptography, one usually uses the two terms correctness and robustness oppositely.

how to choose $\varepsilon_{\text{secure}}$, as the term $\varepsilon_{\text{correct}}$ dominates the sum. This is even more remarkable because the actual choice of $\varepsilon_{\text{secure}}$ can influence the key rate notably, challenging the practical results in this area: Computing, for example, a 256-bit AES key for $\varepsilon_{\text{secure}} = 2^{-40}$ requires to cut approximately $2 \log 1/\varepsilon_{\text{secure}} = 80$ bits for privacy amplification (also in the quantum case [32]), whereas for $\varepsilon_{\text{secure}} = 2^{-256}$ one loses $2 \log 1/\varepsilon_{\text{secure}} = 512$ bits. Hence, in one case, one requires at least 336 reconciled bits to produce an AES key; in the other case, at least 768 reconciled bits. This even neglects the overhead of maintaining the seed for privacy amplification and the extra time needed to evaluate the amplification step for the smaller statistical bound.

## 7 Conclusion

When one implements a key generation procedure for $\lambda$-bit keys close to uniform, our results suggest using a statistical distance $\Delta$ equal to $2^{-\lambda}$. This gives the best bit security level against quantum key search one can hope for. It guarantees a level of approximately $\lambda - \log T_Q$ where $T_Q \leq 2^{\lambda/2}$ is an upper bound on the run time of quantum algorithms. In particular, one achieves the expected bit security bound of at least $\lambda/2$.

Choosing a smaller statistical distance than $2^{-\lambda}$ gives no additional advantage according to our results. If necessary, one may increase $\Delta$ up to $2^{-\lambda/2}$, in which case one still gets the expected security lower bound of $\lambda/2$ bits. However, this lower bound is not known to match potentially improved bounds if the global time bound $T_Q$ of quantum adversaries is actually smaller than $2^{\lambda/2}$. Choosing statistical distances (significantly) larger than $2^{-\lambda/2}$ is generally not recommended. Depending on the concrete distribution, one may, however, achieve better results than via the abstract view on the statistical distance.

We emphasize that our results give conservative bounds for general key distributions for which only the statistical distance from uniform is known. If one has specific information about the distribution of keys, one could, in principle, derive more exact bounds. This may be relevant for the case of quantum key distribution (cf. Section 6) where smaller statistical distances are often preferred for efficiency reasons. Additionally, our results assume that ordering the keys according to the likelihood comes for free. Quantifying the effort for this part may also allow to argue bounds for smaller statistical distances.

## Acknowledgments

We thank the anonymous reviewers for their valuable and constructive feedback. This research has been funded by the German Federal Ministry of Education and Research (BMBF) as part of the program "High-performance components and optimized materials for quantum communication" under project grant CBQD (16KIS1942).

## References

1. Bar-Yossef, Z., Jayram, T.S., Kumar, R., Sivakumar, D.: An information statistics approach to data stream and communication complexity. J. Computer and System Sciences **68**(4), 702–732 (Jun 2004). https://doi.org/10.1016/j.jcss.2003.11.006, (Preliminary Version in *43rd FOCS*, 2002)
2. Barker, E.: Nist sp 800-57 part 1 rev. 5: Recommendation for key management: Part 1 – general. Tech. rep., National Institute of Standards and Technology (2020)
3. Boaron, A., Boso, G., Rusca, D., Vulliez, C., Autebert, C., Caloz, M., Perrenoud, M., Gras, G., Bussières, F., Li, M.J., Nolan, D., Martin, A., Zbinden, H.: Secure quantum key distribution over 421 km of optical fiber. Phys. Rev. Lett. **121**, 190502 (Nov 2018). https://doi.org/10.1103/PhysRevLett.121.190502, https://link.aps.org/doi/10.1103/PhysRevLett.121.190502
4. Boyer, M., Brassard, G., Høyer, P., Tapp, A.: Tight bounds on quantum searching. Fortschritte der Physik **46**(4-5), 493–505 (1998)
5. Bunandar, D., Govia, L., Krovi, H., Englund, D.: Numerical finite-key analysis of quantum key distribution. npj Quantum Information **6** (12 2020). https://doi.org/10.1038/s41534-020-00322-w
6. Bundesamt für Sicherheit in der Informationstechnik: Bsi tr-02102-1: Cryptographic mechanisms: Recommendations and key lengths. Tech. rep., Bundesamt für Sicherheit in der Informationstechnik (2024)
7. Chen, L.: Recommendation for key derivation using pseudorandom functions. Tech. Rep. NIST Special Publication (SP) 800-108-r1-upd1, National Institute of Standards and Technology (NIST), Gaithersburg, MD (Feb 2024). https://doi.org/10.6028/NIST.SP.800-108r1-upd1
8. Cogliati, B., Fouque, P.A., Goubin, L., Minaud, B.: New security proofs and techniques for hash-and-sign with retry signature schemes. Cryptology ePrint Archive, Paper 2024/609 (2024), https://eprint.iacr.org/2024/609
9. Dodis, Y., Gennaro, R., Håstad, J., Krawczyk, H., Rabin, T.: Randomness extraction and key derivation using the cbc, cascade and HMAC modes. In: Franklin, M.K. (ed.) Advances in Cryptology - CRYPTO 2004, 24th Annual International CryptologyConference, Santa Barbara, California, USA, August 15-19, 2004, Proceedings. Lecture Notes in Computer Science, vol. 3152, pp. 494–510. Springer (2004). https://doi.org/10.1007/978-3-540-28628-8_30, https://doi.org/10.1007/978-3-540-28628-8_30
10. Fouque, P.A., Pointcheval, D., Zimmer, S.: HMAC is a randomness extractor and applications to TLS. In: Abe, M., Gligor, V. (eds.) ASIACCS 08: 3rd ACM Symposium on Information, Computer and Communications Security. pp. 21–32. ACM Press, Tokyo, Japan (Mar 18–20, 2008). https://doi.org/10.1145/1368310.1368317
11. Gentry, C., Wichs, D.: Separating succinct non-interactive arguments from all falsifiable assumptions. Cryptology ePrint Archive, Paper 2010/610 (2010), https://eprint.iacr.org/2010/610, https://eprint.iacr.org/2010/610
12. Håstad, J., Impagliazzo, R., Levin, L.A., Luby, M.: A pseudorandom generator from any one-way function. SIAM Journal on Computing **28**(4), 1364–1396 (1999)
13. He, X., Sun, X., Zhang, J.: Quantum search with prior knowledge. Science China Information Sciences **67**(9), 192503 (2024). https://doi.org/10.1007/s11432-023-3972-y, https://doi.org/10.1007/s11432-023-3972-y
14. Krawczyk, H.: Cryptographic extraction and key derivation: The HKDF scheme. In: Rabin, T. (ed.) Advances in Cryptology – CRYPTO 2010. Lecture Notes in

Computer Science, vol. 6223, pp. 631–648. Santa Barbara, CA, USA (Aug 15–19, 2010). https://doi.org/10.1007/978-3-642-14623-7_34

15. Krawczyk, H., Eronen, P.: HMAC-based Extract-and-Expand Key Derivation Function (HKDF). RFC 5869 (May 2010). https://doi.org/10.17487/RFC5869, https://www.rfc-editor.org/info/rfc5869
16. Lee, K.: Bit security as cost to demonstrate advantage **1**(1), 1 (2024). https://doi.org/10.62056/an5txol7
17. Lenstra, A.K.: Key lengths. In: Bidgoli, H. (ed.) The Handbook of Information Security, pp. 617–635. John Wiley & Sons, Hoboken, NJ (2005), contribution to The Handbook of Information Security
18. Liu, H., Yin, Z., Wang, R., Wang, Z.H., Wang, S., Chen, W., Guo, G.C., Han, Z.F.: Tight finite-key analysis for quantum key distribution without monitoring signal disturbance. npj Quantum Information **7** (12 2021). https://doi.org/10.1038/s41534-021-00428-9
19. Micciancio, D., Schultz-Wu, M.: Bit security: Optimal adversaries, equivalence results, and a toolbox for computational-statistical security analysis. In: Boyle, E., Mahmoody, M. (eds.) Theory of Cryptography - 22nd International Conference, TCC 2024, Milan, Italy, December 2-6, 2024, Proceedings, Part II. Lecture Notes in Computer Science, vol. 15365, pp. 224–254. Springer (2024). https://doi.org/10.1007/978-3-031-78017-2_8, https://doi.org/10.1007/978-3-031-78017-2_8
20. Micciancio, D., Walter, M.: On the bit security of cryptographic primitives. In: Nielsen, J.B., Rijmen, V. (eds.) Advances in Cryptology – EUROCRYPT 2018, Part I. Lecture Notes in Computer Science, vol. 10820, pp. 3–28. Tel Aviv, Israel (Apr 29 – May 3, 2018). https://doi.org/10.1007/978-3-319-78381-9_1
21. Mizutani, A., Curty, M., Imoto, N., Tamaki, K.: Finite-key security analysis of quantum key distribution with imperfect light sources. New Journal of Physics **17** (09 2015). https://doi.org/10.1088/1367-2630/17/9/093011
22. Montanaro, A.: Quantum search with advice. In: van Dam, W., Kendon, V.M., Severini, S. (eds.) Theory of Quantum Computation, Communication, and Cryptography. pp. 77–93. Springer Berlin Heidelberg, Berlin, Heidelberg (2011)
23. Müller-Quade, J., Renner, R.: Composability in quantum cryptography. New Journal of Physics **11**, 085006 (2009), https://api.semanticscholar.org/CorpusID:7735175
24. Portmann, C., Renner, R.: Security in quantum cryptography. Rev. Mod. Phys. **94**, 025008 (Jun 2022). https://doi.org/10.1103/RevModPhys.94.025008, https://link.aps.org/doi/10.1103/RevModPhys.94.025008
25. Renner, R., König, R.: Universally composable privacy amplification against quantum adversaries. In: Kilian, J. (ed.) Theory of Cryptography, Second Theory of Cryptography Conference, TCC 2005, Cambridge, MA, USA, February 10-12, 2005, Proceedings. Lecture Notes in Computer Science, vol. 3378, pp. 407–425. Springer (2005). https://doi.org/10.1007/978-3-540-30576-7_22, https://doi.org/10.1007/978-3-540-30576-7_22
26. Renner, R., Wolf, R.: Quantum advantage in cryptography. AIAA Journal **61**(5), 1895–1910 (2023). https://doi.org/10.2514/1.J062267, https://doi.org/10.2514/1.J062267
27. Rescorla, E.: The Transport Layer Security (TLS) Protocol Version 1.3. RFC 8446 (Aug 2018). https://doi.org/10.17487/RFC8446, https://www.rfc-editor.org/info/rfc8446
28. Stebila, D., Fluhrer, S., Gueron, S.: Hybrid key exchange in TLS 1.3. Internet-Draft draft-ietf-tls-hybrid-design-11, Internet Engineering Task Force (Oct 2024),

https://datatracker.ietf.org/doc/draft-ietf-tls-hybrid-design/11/, work in Progress

29. Steinberger, J.: Improved security bounds for key-alternating ciphers via hellinger distance. Cryptology ePrint Archive, Paper 2012/481 (2012), https://eprint.iacr.org/2012/481
30. Tomamichel, M., Gisin, N., Renner, R.: Tight finite-key analysis for quantum cryptography. Nature communications **3**, 634 (01 2012). https://doi.org/10.1038/ncomms1631
31. Tomamichel, M., Leverrier, A.: A largely self-contained and complete security proof for quantum key distribution. Quantum **1**, 14 (07 2017). https://doi.org/10.22331/q-2017-07-14-14
32. Tomamichel, M., Schaffner, C., Smith, A., Renner, R.: Leftover hashing against quantum side information. Information Theory, IEEE Transactions on **57**, 5524 – 5535 (09 2011). https://doi.org/10.1109/TIT.2011.2158473
33. Viamontes, G.F., Markov, I.L., Hayes, J.P.: Is quantum search practical? Comput. Sci. Eng. **7**(3), 62–70 (2005). https://doi.org/10.1109/MCSE.2005.53, https://doi.org/10.1109/MCSE.2005.53
34. Watanabe, S., Yasunaga, K.: Bit security as computational cost for winning games with high probability. In: Tibouchi, M., Wang, H. (eds.) Advances in Cryptology – ASIACRYPT 2021, Part III. Lecture Notes in Computer Science, vol. 13092, pp. 161–188. Singapore (Dec 6–10, 2021). https://doi.org/10.1007/978-3-030-92078-4_6
35. Yasunaga, K.: Replacing Probability Distributions in Security Games via Hellinger Distance. In: Tessaro, S. (ed.) 2nd Conference on Information-Theoretic Cryptography (ITC 2021). Leibniz International Proceedings in Informatics (LIPIcs), vol. 199, pp. 17:1–17:15. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl, Germany (2021). https://doi.org/10.4230/LIPIcs.ITC.2021.17, https://drops.dagstuhl.de/entities/document/10.4230/LIPIcs.ITC.2021.17
36. Zalka, C.: Grover's quantum searching algorithm is optimal. Physical Review A **60**(4), 2746–2751 (Oct 1999). https://doi.org/10.1103/physreva.60.2746, http://dx.doi.org/10.1103/PhysRevA.60.2746
37. Zhang, W., Leent, T., Redeker, K., Garthoff, R., Schwonnek, R., Fertig, F., Eppelt, S., Rosenfeld, W., Scarani, V., Lim, C., Weinfurter, H.: A device-independent quantum key distribution system for distant users. Nature **607**, 687–691 (07 2022). https://doi.org/10.1038/s41586-022-04891-y
38. Zuckerman, D.: General weak random sources. In: 31st Annual Symposium on Foundations of Computer Science. pp. 534–543. IEEE Computer Society Press, St. Louis, MO, USA (Oct 22–24, 1990). https://doi.org/10.1109/FSCS.1990.89574