Tight Quantum Distance-Bounding with Conditional Terrorist-Fraud Resistance

Kevin Bogner*, Aysajan Abidin*, Dave Singelée†, and Bart Preneel*
*COSIC, KU Leuven, Leuven, Belgium
†DistriNet, KU Leuven, Leuven, Belgium
Emails: {kevin.bogner, aysajan, dave.singelee, bart.preneel}@kuleuven.be

***Abstract*—We present a quantum distance-bounding (QDB) scheme with tight single-round bounds against distance-fraud and mafia-fraud, together with a conditional terrorist-fraud bound for helpers whose assistance does not enable timely prediction of the fresh per-round secret value. QDB authenticates physical proximity through qubit challenge-response exchanges constrained by the speed of light. Yet, prior QDB schemes typically achieve security against distance-fraud and mafia-fraud (relay attacks) but lack provable security against terrorist-fraud. Our main contribution is a novel QDB scheme with three key properties: (i) provable security against distance-fraud and mafia-fraud, together with a conditional terrorist-fraud bound; (ii) tight provable security bounds; and (iii) lightweight quantum prover requirements that reduce hardware overhead and avoid a separate initial slow phase by embedding freshness generation into the fast rounds.**

***Index Terms*—Quantum distance-bounding, distance-bounding, provable security**

I. Introduction

Distance-bounding (DB) enables a verifier to upper-bound the physical distance to a prover by measuring the round-trip time of challenge-response exchanges constrained by the speed of light [1], [2]. DB is used in applications where both *who* and *where* matter, from contactless payments and access control to keyless entry systems [3], [4]. The known security threats in the literature are: (i) *distance-fraud* (DF), in which a lone, far-away prover tries to appear nearby; (ii) *mafia-fraud* (MF), commonly known as relay attacks, in which a man-in-the-middle relays between honest parties; and (iii) *terrorist-fraud* (TF), in which a far-away dishonest prover assists a nearby helper so that the helper can succeed in one protocol session but not in later sessions [2], [5], [6]. Robust DB designs must address all three threats.

Quantum distance-bounding (QDB) explores whether quantum information can strengthen DB guarantees. Intuitively, measurement disturbance and the no-cloning theorem can be leveraged so that responding early or responding without seeing the challenge becomes provably hard [7]–[10]. Unlike classical radio-frequency-based DB, where early-detect/late-commit strategies can exploit partial decoding of challenges, QDB does not allow reliable early decoding of qubits before full reception [11]. Several QDB proposals have appeared [11]–[15], demonstrating how single-qubit or entanglement-based protocols can reduce certain attack surfaces compared to classical schemes.

Despite promising ideas, prior QDB schemes either (i) lack provable security against *all* three threats, in particular against TF; (ii) achieve only non-tight bounds (per-round success $> 1/2$ for some threats); or (iii) place relatively heavy quantum requirements on the prover, such as measure-and-reprepare operations [11]–[15]. This leaves a gap for QDB schemes that require only simple quantum functionality at the prover, yet are tightly secure against all three threats in the baseline QDB security framework [16].

In this paper, we present a QDB protocol that partially closes this gap. Building on the unified security framework for QDB protocols established in [16], we define a protocol where the verifier transmits a challenge comprising a classical random value and a BB84 qubit state. Using the classical value, both parties independently derive a fresh one-bit secret per round. The protocol folds freshness generation into the fast phase and is followed by a decision phase. The prover's quantum operation is deliberately simple: it buffers the incoming qubit while receiving the nonce and computing one bit from a pseudorandom function (PRF), applies either a conditional Pauli-Y or the identity gate, and reflects the challenge qubit; no quantum measurement or state preparation is required on the prover side. This design offers four distinct benefits:

- *Conditional TF security.* We prove a TF bound for helpers that cannot predict the fresh per-round secret value at the fast-round response deadline except with negligible advantage, even after all allowed interaction with the dishonest distant prover. This is a conditional guarantee, not general TF resistance; Section VI-C discusses its scope and the prospects for removing the condition.
- *Tight security bounds.* The hidden, uniformly random one-bit secret forces every DF or MF adversary to succeed in a single round with probability at most $1/2 + \mathrm{negl}(\lambda)$. For TF, the same per-round bound holds under this timely-prediction restriction.
- *Lightweight quantum prover operation.* The prover only implements a conditional unitary (Pauli-Y or identity gate I) and reflects the challenge qubit. This eliminates the need for quantum measurement and state-preparation devices from prior QDB protocols, thereby minimizing the prover's quantum hardware requirements.
- *No separate slow phase.* Unlike prior QDB protocols that require a *separate* preliminary slow phase for nonce

exchange or pre-computation, our protocol folds freshness generation into the fast rounds themselves by sending a fresh nonce in every round. This removes a distinct slow phase, but it does not reduce the total classical communication.

The paper is organized as follows: Section II discusses related work. Section III introduces the notation and the security framework. Section IV specifies the QDB protocol. Section V presents the security analysis. Section VI discusses the scope of the guarantees and the practical costs. Finally, Section VII concludes and outlines future work.

II. Related Work

DB was introduced by Brands and Chaum as a cryptographic technique for upper-bounding a prover's physical distance through timed rapid challenge-response rounds [1]. Since then, classical DB has developed into a mature line of work spanning lightweight constructions for constrained devices, practical applications such as RFID and contactless access control, and systematic analyses of the main fraud notions [2]–[4]. In this literature, DF, MF, and TF are the canonical attack types, with the latter already appearing in early work on identification and relay attacks [5], [6]. A recurring lesson from the classical line is that TF resistance is subtle and highly definition-dependent; in particular, extractor-style formulations and helper-disclosure restrictions capture different assumptions about what a dishonest prover may disclose to a nearby helper without giving away reusable authentication power [2], [17].

QDB explores whether the fast phase can be strengthened by incorporating quantum communication and exploiting quantum-mechanical constraints such as measurement disturbance and no-cloning. Abidin et al. initiated this direction with a BB84-style QDB that adapts the fast phase to qubit transmission and discusses timing, hardware-delay, and implementation aspects together with informal analyses of DF, MF, and TF [11]. Abidin later refined this line with an updated single-qubit QDB that addresses photon-number-splitting concerns and removes prover-to-verifier communication in the final authentication phase [12]. Subsequent work broadened the design space further: entanglement-based protocols use entangled qubits in the fast phase, including a mutual-QDB construction that reduces the communication cost relative to two separate one-way executions, while later proposals also explore Bell-inequality-based and continuous-variable formulations [13]–[15].

Despite this diversity of protocol ideas, QDB security analysis remains less mature than in the classical setting. Existing QDB proposals are typically analyzed in protocol-specific or informal models and do not provide a common game-based treatment of DF, MF, and TF [16]. This has practical security consequences, not just methodological ones. Later re-analysis identified stronger MF strategies against previously published QDB protocols than originally reported and highlighted that claimed TF resistance can depend sensitively on the exact attack model and the form of helper assistance that is allowed [16], [18]. The recent framework of Bogner et al. therefore fills an important gap by formalizing shared system assumptions, timing constraints, and experiments for DF, MF, and TF by quantum-capable adversaries [16].

In this context, the present paper focuses on a combination of properties that prior QDB work has not jointly achieved: tight DF and MF bounds, a conditional TF bound tied to timely prediction of the per-round PRF value, and lightweight prover-side quantum requirements. In particular, prior QDB proposals often have the prover measure incoming quantum states and prepare an outgoing quantum response [11]–[13], whereas our prover uses a timed qubit buffer, applies a simple conditional unitary, and reflects the incoming qubit. This places our protocol at the intersection of the classical literature on formally modeled TF resistance and the QDB literature on quantum fast phases, while using the unified QDB framework as the basis for precise security claims.

Table I summarizes the subset of prior QDB protocols for which the literature provides explicit, directly comparable fraud bounds, and positions our QDB protocol relative to them. The bounds for our QDB protocol are proved later in Section V. We exclude QDB protocols whose analyses are qualitative or whose security claims are not stated as concrete DF/MF/TF success probabilities; in particular, to the best of our knowledge, this applies to the E91-style protocol [14] and the continuous-variable QDB protocol [15].

III. Preliminaries

In this section, we fix notation and recall the QDB security framework from [16].

A. Notation for quantum states and measurements

Let $a \in \{0,1\}$ index the basis ($a = 0$ is the computational Z basis, $a = 1$ is the diagonal X basis). Define the four BB84 states [7]

$$|0\rangle_0 = |0\rangle\,, \quad |1\rangle_0 = |1\rangle\,, \qquad |0\rangle_1 = |+\rangle\,, \quad |1\rangle_1 = |-\rangle\,,$$

where

$$|\pm\rangle = \tfrac{1}{\sqrt{2}}(|0\rangle \pm |1\rangle).$$

We denote these BB84 states by $|r\rangle_a$. For $r \in \{0,1\}$ we write the rank-one projectors as

$$\Pi_{a,r} := |r\rangle_a \langle r|_a\,.$$

We write $\langle r|_a := (|r\rangle_a)^\dagger$ for the associated bra. We also write $\Pi_\psi := |\psi\rangle\langle\psi|$ for the projector onto a pure state $|\psi\rangle$. Measuring $|\psi\rangle$ in basis a yields outcome r with probability

$$\Pr[r] = \left\|\Pi_{a,r}|\psi\rangle\right\|^2 = \langle\psi|\,\Pi_{a,r}\,|\psi\rangle\,.$$

Additionally, we use the following property of Pauli-Y: in the computational basis, $Y|0\rangle = i|1\rangle$ and $Y|1\rangle = -i|0\rangle$. In the diagonal basis, $|\pm\rangle = (|0\rangle \pm |1\rangle)/\sqrt{2}$ and, consequently, $Y|+\rangle = -i|-\rangle$, $Y|-\rangle = i|+\rangle$. In all cases, the state is mapped to the other BB84 state in the same basis up to a global phase.

TABLE I

COMPARISON OF THE PROPOSED QDB PROTOCOL WITH PRIOR QDB PROTOCOLS FOR WHICH THE LITERATURE REPORTS DIRECTLY COMPARABLE PER-ROUND FRAUD SUCCESS BOUNDS. DF (DISTANCE-FRAUD), MF (MAFIA-FRAUD), AND TF (TERRORIST-FRAUD) DENOTE THE REPORTED PER-ROUND SUCCESS PROBABILITIES. FORMAL PROOF INDICATES WHETHER THE CITED LITERATURE PROVIDES A FORMAL PROOF OF THE LISTED BOUNDS. AN ENTRY OF 1 UNDER TF MEANS THAT THE PROTOCOL PROVIDES NO REPORTED TF-RESISTANCE MECHANISM FOR THE LISTED HELPER-ASSISTANCE MODEL. SLOW PHASE INDICATES WHETHER THE PROTOCOL REQUIRES A SEPARATE SLOW PHASE IN ADDITION TO THE FAST ROUNDS; THIS WORK INSTEAD SENDS A FRESH λ-BIT NONCE WITHIN EACH FAST ROUND (SECTION VI-D). BOUNDS FOR THIS WORK ARE PROVED IN SECTION V. * TF BOUND IS CONDITIONAL ON THE HELPER NOT BEING ABLE TO PREDICT THE FRESH PER-ROUND PRF VALUE BEFORE THE RESPONSE DEADLINE; THIS IS STRONGER THAN FINAL-STATE NON-TRANSFERABILITY.

Protocol	*Year*	*Formal proof*	*DF*	*MF*	*TF*	*Resource*	*Slow phase*	*Noise analysis*
Abidin et al. [11]	2017	×	$1/2$	$3/4$	1	Single qubits	✓	×
Abidin [12], [16], [18]	2019	✓	$1/2$	$7/8$	1	Single qubits	✓	✓
Abidin et al. [13]	2024	×	$1/2$	$5/8$	1	Entangled pairs	✓	×
This work	*2026*	✓	$1/2$	$1/2 + \mathrm{negl}(\lambda)$	$1/2 + \mathrm{negl}(\lambda)^*$	*Single qubits*	×	✓

B. Concentration bounds

For our security analysis, we first analyze the single-round success probability and then extend this to a multi-round setting. To perform this extension, we employ the Azuma-Hoeffding inequality [19], [20] for supermartingales. This provides a unified framework for analyzing both the completeness of the QDB protocol and its soundness against an adaptive adversary.

Lemma 1 (Azuma-Hoeffding inequality for filtrations)**.** *Let $(\mathcal{F}_i)_{i=0}^n$ be a filtration. Let $I_i \in [0,1]$ be $\mathcal{F}_i$-measurable and let $S = \sum_{i=1}^n I_i$.*

1) Lower tail: *If for all i, $\mathbb{E}[I_i \mid \mathcal{F}_{i-1}] \geq \mu$, then for any $\delta > 0$:*

$$\Pr[S \leq n(\mu - \delta)] \leq \exp(-2n\delta^2).$$

2) Upper tail: *If for all i, $\mathbb{E}[I_i \mid \mathcal{F}_{i-1}] \leq \mu$, then for any $\delta > 0$:*

$$\Pr[S \geq n(\mu + \delta)] \leq \exp(-2n\delta^2).$$

C. Security model

We adopt the QDB framework [16] as our baseline model for completeness, DF, MF, TF, and soundness. The protocol-level timing convention is stated in Section IV-C; throughout the experiments below, adversarial strategies are required to be *time-respecting*, meaning that all operations and messages are causally consistent with that convention. To make this work self-contained, we restate the relevant formal definitions below.

Definition 1 (Completeness)**.** If an honest prover $\mathcal{P}$ is located within the distance bound at a distance $d \leq B$ from the verifier $\mathcal{V}$, and r denotes the verifier's local randomness, then

$$\Pr\Big[(\mathcal{V}(x,r) \leftrightarrow \mathcal{P}(x)) \text{ accepts}\Big] \geq 1 - \mathrm{negl}(\lambda).$$

Definition 2 (Distance-fraud experiment)**.**

1) **Setup.** Verifier $\mathcal{V}$ and dishonest distant prover $\mathcal{P}^\star$ share a long-term secret key x. The dishonest prover $\mathcal{P}^\star$ is located outside the distance bound at a distance $d > B$ from $\mathcal{V}$.
2) **Challenge session.** $\mathcal{V}$ and $\mathcal{P}^\star$ engage in a complete execution of the QDB protocol, which includes n fast rounds subject to the distance bound check enforced by $\mathcal{V}$.
3) **Output.** $\mathcal{V}$ outputs $\mathrm{Out}_\mathcal{V} \in \{0,1\}$.

Definition 3 (Mafia-fraud experiment)**.**

1) **Setup.** Verifier $\mathcal{V}$ and honest prover $\mathcal{P}$ share a long-term secret key x. Two quantum polynomial-time (QPT) adversaries, $\mathcal{A}_1$ (co-located with $\mathcal{V}$) and $\mathcal{A}_2$ (co-located with $\mathcal{P}$), share an authenticated classical and quantum channel.
2) **Learning phase.** The adversaries may initiate and control any polynomial number of auxiliary executions of the QDB protocol between $\mathcal{V}$ and $\mathcal{P}$: in each such execution, every classical and quantum message between the honest parties passes through $(\mathcal{A}_1, \mathcal{A}_2)$, who may relay, delay, drop, modify, or inject messages arbitrarily (subject only to the timing constraints). At the end of this phase, $(\mathcal{A}_1, \mathcal{A}_2)$ retain the entire classical transcript and their joint quantum state.
3) **Challenge session.** The adversaries $(\mathcal{A}_1, \mathcal{A}_2)$ interact with $\mathcal{V}$ in a full execution of the QDB protocol. Simultaneously, they may interact with the honest $\mathcal{P}$ (who uses the correct long-term key x and is at distance $d > B$). $\mathcal{V}$ enforces the distance bound B.
4) **Output.** $\mathcal{V}$ outputs $\mathrm{Out}_\mathcal{V} \in \{0,1\}$.

Several different definitions of TF resistance appear in the DB literature [17]. We use the TF experiment from the QDB security framework as the baseline attack experiment [16].

Definition 4 (Terrorist-fraud experiment)**.**

1) **Setup.** Verifier $\mathcal{V}$ and dishonest prover $\mathcal{P}^\star$ share a long-term secret key x. $\mathcal{P}^\star$ is located outside the distance bound at a distance $d > B$ from $\mathcal{V}$. A helper $\mathcal{A}$ is co-located with $\mathcal{V}$. The pair $(\mathcal{P}^\star, \mathcal{A})$ share an authenticated classical and quantum channel.
2) **Learning phase.** $(\mathcal{P}^\star, \mathcal{A})$ may engage in any polynomial number of auxiliary executions of the QDB protocol with $\mathcal{V}$. In each such execution, $\mathcal{A}$ is co-located with $\mathcal{V}$ and may relay, modify, or inject messages, while $\mathcal{P}^\star$ participates at its true distant location using key x. The pair may also exchange arbitrary classical and quantum messages over their channel, and all information gathered in this phase is available later.

3) **Challenge session.** $\mathcal{A}$ impersonates $\mathcal{P}$ in a full execution of the QDB protocol with $\mathcal{V}$, subject to the same distance bound B as in an honest run.
4) **Output.** $\mathcal{V}$ outputs $\mathrm{Out}_{\mathcal{V}} \in \{0,1\}$.

Without further restriction, Definition 4 would be trivial: $\mathcal{P}^\star$ could reveal reusable secret information, namely the long-term key x. TF therefore allows help in the current QDB session but must prevent reusable impersonation in later sessions.

Definition 5 (Non-transferable assistance)**.** Let $(\mathcal{P}^\star, \mathcal{A})$ be a QPT pair as in Definition 4, and fix a long-term key x. Run the TF experiment once and consider the final (classical and quantum) state of $\mathcal{A}$ at the end of this execution. Now let $\mathcal{A}$, starting from this state and with no further interaction with $\mathcal{P}^\star$, engage alone in a second execution of the QDB protocol with $\mathcal{V}$, where the long-term key is still x and all nonces and verifier randomness are freshly sampled. We say that the assistance of $(\mathcal{P}^\star, \mathcal{A})$ is *non-transferable* if the probability that $\mathcal{V}$ accepts in this second execution is negligible in the security parameter λ.

This final-state notion is useful for terminology, but the TF theorem below is not stated under this condition alone. The proved TF bound is conditional on the stronger timely-prediction condition introduced in Section V-E.

Definition 6 (Soundness w.r.t. a set of threats)**.** Let $\mathcal{T} \subseteq \{\mathrm{DF}, \mathrm{MF}, \mathrm{TF}\}$. A QDB protocol is $\mathcal{T}$-sound if it is secure against all threats in $\mathcal{T}$: for each $T \in \mathcal{T}$, there exists a negligible function $\varepsilon_T(\lambda)$ such that the verifier accepts in the corresponding T-experiment with probability at most $\varepsilon_T(\lambda)$, maximized over all time-respecting QPT adversaries of the permitted type.

IV. Quantum distance-bounding protocol

Let λ be the security parameter and n the number of fast rounds. The verifier $\mathcal{V}$ and prover $\mathcal{P}$ share a long-term key $x \in \{0,1\}^\lambda$ generated via $\mathsf{KeyGen}(1^\lambda)$. The protocol consists of a fast phase followed by a decision phase.

A. Fast phase

For each round index $i = 1, \ldots, n$, the parties execute the following steps:

1) *Verifier challenge.* $\mathcal{V}$ samples a basis bit $a_i \xleftarrow{\$} \{0,1\}$, a challenge bit $c_i \xleftarrow{\$} \{0,1\}$, and a classical string $\gamma_i \xleftarrow{\$} \{0,1\}^\lambda$ uniformly at random. The verifier records the time-of-departure ToD_i, defined as the time at which the round-i hybrid challenge is emitted, and issues this challenge, which consists of the classical string γ_i and the quantum state $|c_i\rangle_{a_i}$, to $\mathcal{P}$.
2) *Compute one-bit secret.* Using a one-bit quantum-secure pseudorandom function (QPRF) [21] $\mathsf{F}_x : \{0,1\}^* \to \{0,1\}$, both parties compute the one-bit secret
$$b_i \;:=\; \mathsf{F}_x\big(i \,\|\, \gamma_i\big).$$
3) *Prover response.* Upon receiving the challenge, the prover buffers the qubit $|c_i\rangle_{a_i}$ (e.g., using a passive optical delay line) for a duration sufficient to receive γ_i and compute b_i. The prover then applies the unitary Y^{b_i} to the challenge qubit, which acts as a bit flip in both the computational and diagonal bases when $b_i = 1$ and as the identity gate when $b_i = 0$. The resulting response state $|c_i \oplus b_i\rangle_{a_i}$ is then reflected back to $\mathcal{V}$. Thus the prover avoids measurement and state preparation, but it does require a timed quantum buffer or delay line long enough to cover classical nonce reception and PRF-computation latency.
4) *Verifier check.* Upon reception of the response state at time-of-arrival ToA_i, $\mathcal{V}$ records the round-trip time
$$\Delta t_i := \mathrm{ToA}_i - \mathrm{ToD}_i,$$
and measures the response in the secret basis a_i to obtain the outcome c_i'. The round is marked *accepted* by $\mathcal{V}$ if and only if
$$\Delta t_i \le \Delta t_{\max} \quad \text{and} \quad c_i' = c_i \oplus b_i \;=:\; h_i.$$

B. Decision phase

Let $I_i \in \{0,1\}$ be the indicator variable for acceptance in round i, and let $S = \sum_{i=1}^n I_i$ denote the total number of accepted rounds. Fix a constant security slack $\varepsilon' \in (0, \frac{1}{2})$, independent of λ. The session is accepted if and only if $S \ge \tau$, where the threshold is defined as
$$\tau = \left\lceil n\left(\tfrac{1}{2} + \varepsilon'\right)\right\rceil.$$

The complete protocol execution is depicted in Figure 1.

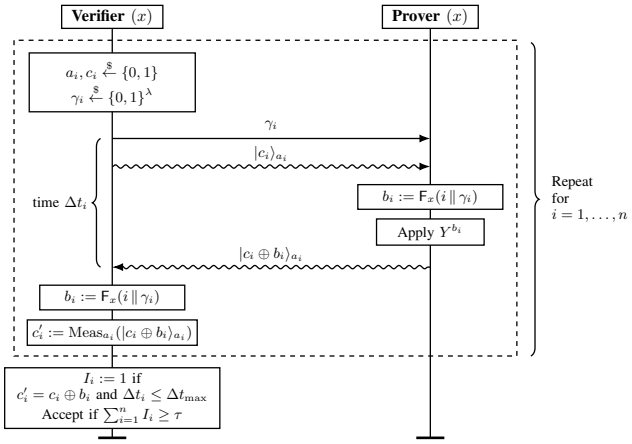

Fig. 1. Schematic of the proposed quantum distance-bounding protocol. Each fast round uses fresh verifier randomness (a_i, c_i, γ_i). The prover applies Y^{b_i} according to the one-bit secret b_i, and the verifier checks timing and response correctness.

C. Timing model and design rationale

In every fast round i, the verifier emits the whole hybrid challenge $(\gamma_i, |c_i\rangle_{a_i})$ as one atomic challenge at time ToD_i. Adversarial classical and quantum channels obey the same maximum signal speed as the honest channels, with adversarial systems co-located as specified in the DF, MF, or TF experiment. Thus, message serialization is idealized away: no party outside the distance bound can obtain $|c_i\rangle_{a_i}$, γ_i, or even

a prefix of γ_i before it must emit a response that can arrive by the deadline. This is an idealized timing assumption; a practical implementation would need sufficiently parallel, high-bandwidth, and calibrated transmission of the classical and quantum challenge components.

The classical string γ_i is a public nonce whose only purpose is to salt the keyed computation $b_i = \mathsf{F}_x(i\|\gamma_i)$, making b_i a fresh one-bit mask for c_i that is *computationally* unpredictable (except with negligible advantage) to any entity that does not possess the long-term key x at the fast-round response deadline. This protects against the pre-ask attack in MF and also underlies the TF analysis of Section V, where successful fast-round impersonation is reduced to predicting the fresh bit b_i on input $i\|\gamma_i$. Embedding γ_i into the fast phase removes a separate slow phase: freshness is obtained on the fly, but each round now carries a fresh λ-bit classical nonce. Finally, the acceptance test itself reduces to the hidden bit h_i, but the challenge bit c_i travels only as a BB84 state in a basis known to nobody except the verifier: while in transit, it cannot be read out reliably, partially decoded, or copied for deferred measurement [7]–[10]. Section VI-A makes precise what the quantum challenge does and does not contribute to the proved bounds, and Section VI-B shows how the guarantees degrade when atomic release is relaxed.

V. Security analysis

In this section, we prove the security properties of the QDB protocol. We first establish completeness, followed by a formal analysis of security against DF, MF, and TF in the baseline security framework, concluding with DF soundness, PRF-based MF soundness, and TF soundness relative to helpers that cannot predict fresh per-round PRF values in time. While our completeness analysis incorporates channel noise to reflect realistic conditions, the subsequent adversarial proofs assume a noiseless setting. This assumption represents a worst-case scenario for the verifier and is without loss of generality, as stochastic noise inherently increases the adversary's probability of failure. Throughout, timing arguments are interpreted in the idealized model.

A. Completeness

We first establish completeness (Definition 1) for the QDB protocol defined in Section IV. Recall that round i is *accepted* if and only if the verifier's timing constraints are met and the measured bit satisfies $c_i' = c_i \oplus b_i$. We assume the classical nonce is received correctly, for example through authenticated error-corrected classical transmission; classical-channel errors can be treated as an additional completeness loss.

Lemma 2 (Single-round completeness (noiseless))**.** *For every round i, an honest prover $\mathcal{P}$ is accepted with probability 1 in the noiseless model:*

$$\Pr[I_i = 1] = 1.$$

Proof. Timing. By the definition of $\Delta t_{\max}$, any signal exchange with an honest prover located within the distance bound arrives within the acceptance window. Any exchange with round-trip time greater than $\Delta t_{\max}$ is rejected. Therefore, for an honest prover $\mathcal{P}$ within the distance bound, the round-trip time is bounded by $\Delta t_{\max}$, satisfying the timing constraint.

Value. For each round i, the verifier $\mathcal{V}$ prepares and transmits the challenge consisting of the classical string γ_i and the quantum state $|c_i\rangle_{a_i}$. Because the honest prover $\mathcal{P}$ possesses the long-term key x and the round index i, they can correctly compute the one-bit secret b_i upon receiving γ_i. The prover then applies the unitary Y^{b_i} and returns the state $Y^{b_i}|c_i\rangle_{a_i}$ to the verifier. As shown in Section III-A, $Y^{b_i}|c_i\rangle_{a_i}$ is equivalent to $|c_i \oplus b_i\rangle_{a_i}$ up to a global phase, which does not affect measurement statistics. Consequently, measuring in basis a_i yields the outcome $c_i' = c_i \oplus b_i$ with probability 1 in the noiseless setting, satisfying the value test. Thus, $I_i = 1$. □

Noise model. We model the transmission of each qubit as an independent passage through a single-qubit depolarizing channel [10] $\mathcal{D}_\eta$ with parameter $\eta \in [0,1]$:

$$\mathcal{D}_\eta(\rho) \;:=\; (1-\eta)\,\rho \;+\; \eta\,\tfrac{\mathbb{I}}{2}.$$

Both the quantum forward (challenge) and backward (response) channels are subject to $\mathcal{D}_\eta$, applied independently to each hop and all rounds. Any noise introduced by the prover's timed quantum buffer or delay line is absorbed into this effective depolarizing parameter. Photon loss is not captured by this model; Section VI-D discusses a loss-tolerant variant of the acceptance rule. The next theorem quantifies how this noise affects the honest acceptance probability over n rounds.

Theorem 3 (Completeness with noise)**.** *Under the defined noise model, the acceptance probability for an honest prover $\mathcal{P}$ in a single round is $p(\eta) = 1 - \eta + \frac{1}{2}\eta^2$. For any threshold $\tau < n\,p(\eta)$,*

$$\Pr\big[S < \tau\big] \;\le\; \exp\Big(-2n\big(p(\eta) - \tfrac{\tau}{n}\big)^2\Big).$$

In the specific case where $\eta = 0$, we recover $\Pr[S < \tau] = 0$ as per Lemma 2.

Proof. In an honest execution, the prover performs the unitary operation Y^{b_i} on the received state ρ. For the depolarizing channel, the order of this unitary operation and the noise does not matter:

$$\mathcal{D}_\eta\big(Y^{b_i}\rho(Y^{b_i})^\dagger\big) = Y^{b_i}\,\mathcal{D}_\eta(\rho)\,(Y^{b_i})^\dagger.$$

Since the qubit passes through independent depolarizing channels on the forward and backward paths, the effective channel is:

$$\mathcal{D}_\eta \circ \mathcal{D}_\eta = \mathcal{D}_{\eta'}, \qquad \text{where } \eta' = 2\eta - \eta^2.$$

Applying the depolarizing channel twice shrinks the Bloch vector by a factor of $(1-\eta)^2$, since each hop shrinks it by a factor of $(1-\eta)$. Thus, initializing with the pure state $|c_i \oplus b_i\rangle_{a_i}\langle c_i \oplus b_i|_{a_i}$, the state received by the verifier is

$$\rho_{\text{hon}} = (1-\eta')\Pi_{a_i, c_i\oplus b_i} + \eta'\,\tfrac{\mathbb{I}}{2}.$$

Thus, the received state is the correct state with weight $1-\eta'$ and uniform noise with weight η'. The verifier measures the

received qubit in basis a_i, yielding the correct outcome with probability

$$p(\eta) = \mathrm{tr}\big(\Pi_{a_i, c_i \oplus b_i}\, \rho_{\mathrm{hon}}\big) = (1-\eta') + \tfrac{\eta'}{2} = 1 - \eta + \tfrac{1}{2}\eta^2.$$

We derive the tail bound for the total number of successful rounds $S = \sum_{i=1}^{n} I_i$. Since the rounds are independent with success probability $p(\eta)$, they form a trivial supermartingale satisfying the condition of Lemma 1 (lower tail). Thus, for any $\tau < np(\eta)$:

$$\Pr[S < \tau] \le \exp\Big(-2n\big(p(\eta) - \tfrac{\tau}{n}\big)^2\Big).$$

□

For constant noise η and constant slack $p(\eta) - \tau/n > 0$, the completeness failure bound is negligible in λ when $n = n(\lambda) = \omega(\log \lambda)$. If n is fixed, the same expression should instead be read as a concrete finite-session failure bound.

B. Reduction to hidden-bit prediction

In this subsection, we reduce an adversary's probability of making the verifier accept to a classical hidden-bit prediction problem. This reduction is specific to the QDB protocol defined in Section IV. It allows us to bound the acceptance probability of any strategy by the adversary's information-theoretic or computational ability to predict the value $h_i := c_i \oplus b_i$.

Lemma 4 (Hidden-bit Lemma)**.** *Let $\mathcal{F}_{i-1}$ be any admissible history before round i. Let ρ_i be the reduced single-qubit state received by the verifier within the timing window, after tracing out any private adversarial systems with which it may be entangled. For any realized ρ_i,*

$$\Pr[I_i = 1 \mid \mathcal{F}_{i-1}, a_i, c_i, b_i, \rho_i] = \mathrm{tr}\left(\Pi_{a_i, c_i \oplus b_i} \rho_i\right).$$

Suppose that, conditioned on $\mathcal{F}_{i-1}$, no QPT predictor given the adversary's full classical and quantum view at the round-i response deadline, including any response system emitted to the verifier, and the basis a_i can predict $h_i = c_i \oplus b_i$ with probability greater than $\frac{1}{2} + \alpha$. Then every time-respecting response strategy satisfies

$$\mathbb{E}[I_i \mid \mathcal{F}_{i-1}] \le \tfrac{1}{2} + \alpha.$$

Proof. The displayed identity is Born's rule applied to the verifier's basis-a_i measurement of the reduced response state. For the bound, construct a predictor that runs the adversarial strategy through the response deadline, measures the emitted response qubit in basis a_i, and outputs the measurement result as its guess for h_i. If no timely response qubit is emitted, the predictor outputs an arbitrary bit. Whenever the verifier accepts, this predictor guesses h_i correctly. Therefore the predictor's success probability is at least $\mathbb{E}[I_i \mid \mathcal{F}_{i-1}]$. The assumed upper bound on all such predictors gives the claim. □

C. Distance-fraud security

Recall the DF experiment from Definition 2. The dishonest distant prover $\mathcal{P}^\star$ is located outside the distance bound at a distance $d > B$ and shares the long-term key x with the verifier $\mathcal{V}$. Unlike in existing QDB protocols [11]–[15], the basis a_i is withheld from $\mathcal{P}^\star$, as it is generated locally by $\mathcal{V}$. However, the prover can still perform the Y^{b_i} operation if it learns the fresh nonce in time.

Crucially, under the atomic-release timing model of Section IV-C, any response that arrives at $\mathcal{V}$ within the acceptance window must have been emitted by $\mathcal{P}^\star$ before any information about the fresh hybrid challenge $(\gamma_i, |c_i\rangle_{a_i})$ could have reached the prover. This rules out partial-nonce strategies in which a distant key holder learns a prefix of γ_i before responding. Consequently, the response state available to $\mathcal{V}$ within the acceptance window is independent of the fresh challenge. This independence property is formalized in [16], [22].

Lemma 5 (Distance-fraud single-round bound)**.** *For every round i and every (possibly adversarially generated) history H_{i-1} consisting of the full view up to the end of round $i-1$, a distant prover's time-respecting strategy satisfies*

$$\Pr[I_i = 1 \mid H_{i-1}] \;\le\; \tfrac{1}{2}.$$

Proof. By the timing model in Section IV-C, the spatial separation $d > B$ ensures that, conditioned on any history H_{i-1}, the distant prover's deadline view is independent of the fresh challenge bit c_i and the fresh nonce γ_i. Since c_i is uniform, the hidden bit $h_i = c_i \oplus b_i$ is uniform from that view even if the verifier's basis a_i is revealed. Thus no predictor given the deadline view and a_i can predict h_i with probability above $1/2$. Applying Lemma 4 with $\mathcal{F}_{i-1} = H_{i-1}$ gives $\Pr[I_i = 1 \mid H_{i-1}] \le 1/2$. □

Remark 1 (Robustness of the distance-fraud bound). The bound of Lemma 5 does not depend on the secrecy or the timing of the classical components. Suppose the nonce γ_i, the secret bit b_i, and even the basis a_i are revealed to the distant prover arbitrarily early. Any response state ρ_i emitted before the challenge bit c_i can have reached the prover is still independent of c_i, so its acceptance probability, averaged over the uniform c_i, is

$$\tfrac{1}{2} \sum_{c \in \{0,1\}} \mathrm{tr}\big(\Pi_{a_i,\, c \oplus b_i}\, \rho_i\big) \;=\; \tfrac{1}{2}\,\mathrm{tr}(\rho_i) \;=\; \tfrac{1}{2},$$

since $\Pi_{a_i,0} + \Pi_{a_i,1} = \mathbb{I}$. The distance-fraud bound therefore rests on the quantum challenge bit alone. Sections VI-A and VI-B build on this observation.

D. Mafia-fraud security

Recall the MF experiment from Definition 3. In our protocol, the verifier's round-i acceptance condition depends only on the hidden bit

$$h_i := c_i \oplus b_i, \qquad b_i := \mathsf{F}_x(i \| \gamma_i).$$

The adversary pair $(\mathcal{A}_1, \mathcal{A}_2)$ may obtain polynomially many input-output samples of the keyed function F_x in the learning phase by interacting with the honest prover on chosen nonce/qubit pairs. Because history-conditioned per-round bounds do not transfer cleanly across the PRF game hop, we analyze MF in two steps: first in a random-function game, where fresh inputs yield truly random hidden bits, and then at the *session* level, where the final acceptance event is transferred back to the real PRF world. All PRF inputs in these executions are classical strings $m = i\|\gamma$: the nonce γ is an ordinary classical protocol message. Thus each prover query uses one chosen nonce string, not a quantum state encoding many nonce values at once. During learning, the adversary may nevertheless submit arbitrary quantum states as the challenge qubit; the honest prover applies $Y^{\mathsf{F}_x(m)}$ to the submitted qubit for the corresponding classical input m.

Lemma 6 (Mafia-fraud single-round bound in the random-function game)**.** *Let Game RF denote the mafia-fraud experiment in which F_x is replaced everywhere by a truly random function*

$$\mathsf{R} : \{0,1\}^* \to \{0,1\}.$$

The function R is sampled once and is consistent across repeated queries. Then, for every time-respecting QPT mafia strategy $(\mathcal{A}_1, \mathcal{A}_2)$, every round i, and every history H_{i-1} up to the end of round $i-1$ in Game RF,

$$\Pr[I_i = 1 \mid H_{i-1}] \le \tfrac{1}{2} + \mathrm{negl}(\lambda).$$

Proof. Fix a round i and a history H_{i-1}. By the start of round i, the adversaries can have determined $\mathsf{R}(m)$ on only polynomially many classical inputs m: these values can only come from the polynomially many auxiliary-session inputs and the first $i-1$ challenge-round inputs. Hence there is a polynomial q such that, conditioned on any fixed H_{i-1}, at most $q(\lambda)$ inputs to R are already known.

The next input is

$$m_i := i\|\gamma_i.$$

Since $\gamma_i \xleftarrow{\$} \{0,1\}^\lambda$ is fresh and independent of H_{i-1}, any fixed previously known input is hit with probability at most $2^{-\lambda}$. A union bound over the at most $q(\lambda)$ previously known inputs gives

$$\Pr[\mathsf{R}(m_i) \text{ is already known} \mid H_{i-1}] \le \frac{q(\lambda)}{2^\lambda},$$

which is $\mathrm{negl}(\lambda)$.

Let Fresh_i be the complementary event, namely that $\mathsf{R}(m_i)$ was not previously determined. Conditioned on Fresh_i, the bit

$$b_i = \mathsf{R}(m_i)$$

is uniform and independent of the adversaries' view at the response deadline. Moreover, although $\mathcal{A}_2$ may start a concurrent query to the honest prover on the fresh input m_i after learning γ_i, the timing convention of Section IV-A prevents the returned bit from reaching $\mathcal{A}_1$ before the verifier's deadline because the honest prover is at distance $d > B$.

Therefore, conditioned on $(H_{i-1}, \mathrm{Fresh}_i)$, the hidden bit

$$h_i := c_i \oplus b_i$$

is uniform and independent of the adversaries' deadline view, and Lemma 4 yields

$$\Pr[I_i = 1 \mid H_{i-1}, \mathrm{Fresh}_i] \le \tfrac{1}{2}.$$

In the complementary event $\neg\mathrm{Fresh}_i$, we use the trivial bound 1. Hence

$$\Pr[I_i = 1 \mid H_{i-1}] \le \tfrac{1}{2} + \Pr[\neg\mathrm{Fresh}_i \mid H_{i-1}] \le \tfrac{1}{2} + \mathrm{negl}(\lambda).$$

□

Theorem 7 (Mafia-fraud session bound in the real game)**.** *Assume that F is secure against QPT distinguishers making classical oracle queries. For every time-respecting QPT mafia strategy $(\mathcal{A}_1, \mathcal{A}_2)$, every polynomially bounded n, every constant $\varepsilon' \in (0, \frac{1}{2})$, and threshold*

$$\tau = \left\lceil n\left(\tfrac{1}{2} + \varepsilon'\right)\right\rceil,$$

the session acceptance probability in the real MF experiment satisfies

$$\Pr[S \ge \tau] \le \exp\!\left(-2n(\varepsilon' - \mathrm{negl}(\lambda))^2\right) + \mathrm{negl}(\lambda)$$

for all sufficiently large λ.

Proof. Game 0 (real). This is the real MF experiment with $b_i = \mathsf{F}_x(i\|\gamma_i)$.

Game 1 (random-function). This is identical except that every use of F_x by the honest parties is replaced by a truly random function R.

We build a QPT distinguisher $\mathcal{D}$ for the PRF by simulating the entire MF experiment using its oracle $\mathcal{O}$ in place of F_x. Whenever an honest prover or verifier needs the bit on a classical input $m = i\|\gamma$, $\mathcal{D}$ queries $\mathcal{O}(m)$ and uses the returned bit to simulate the honest protocol exactly; for honest-prover learning sessions, it applies $Y^{\mathcal{O}(m)}$ to the arbitrary submitted challenge qubit. If $\mathcal{O} = \mathsf{F}_x$, the simulation is Game 0; if $\mathcal{O} = \mathsf{R}$, it is Game 1. Hence, by the assumed classical-query post-quantum PRF security,

$$\left|\Pr_{\mathrm{G0}}[S \ge \tau] - \Pr_{\mathrm{G1}}[S \ge \tau]\right| \le \mathrm{negl}(\lambda).$$

In Game 1, Lemma 6 gives, for every round i and every history H_{i-1},

$$\mathbb{E}[I_i \mid H_{i-1}] \le \tfrac{1}{2} + \mathrm{negl}(\lambda).$$

Applying Lemma 1 (upper tail) to this per-round bound yields

$$\Pr_{\mathrm{G1}}[S \ge \tau] \le \exp\!\left(-2n(\varepsilon' - \mathrm{negl}(\lambda))^2\right)$$

for all sufficiently large λ.

Combining the two games,

$$\Pr_{\mathrm{G0}}[S \ge \tau] \le \exp\!\left(-2n(\varepsilon' - \mathrm{negl}(\lambda))^2\right) + \mathrm{negl}(\lambda).$$

This proves the claim. □

E. Terrorist-fraud security

We now consider TF in the experiment of Definition 4. The dishonest prover $\mathcal{P}^\star$ knows the long-term key x but is outside the distance bound, while the helper $\mathcal{A}$ is close to the verifier. By the timing model of Section IV-C, once the verifier samples the fresh nonce γ_i in round i, the distant prover cannot receive γ_i, compute $b_i = \mathsf{F}_x(i\|\gamma_i)$, and send useful information back to $\mathcal{A}$ before the verifier's response deadline.

Thus, in this protocol, a timely TF helper can beat the random-guessing bound only if the assistance already available to the helper before the response deadline predicts the fresh round bit b_i. We make this condition explicit.

We say that a TF strategy has fresh-bit-unpredictable assistance if the following holds. For every round index i, every admissible history H_{i-1} before round i, and every fresh nonce $\gamma_i \xleftarrow{\$} \{0,1\}^\lambda$ whose input $i\|\gamma_i$ has not appeared in the learning phase or in earlier challenge rounds, every QPT algorithm that is given γ_i, the verifier's sampled a_i, c_i, and the helper's entire timely pre-response view predicts

$$b_i = \mathsf{F}_x(i\|\gamma_i)$$

with probability at most $\frac{1}{2} + \mathrm{negl}(\lambda)$.

The values a_i and c_i are given to the predictor only for the security reduction; they are not revealed to the helper in the actual protocol. This is the precise conditional TF assumption used below. It is stronger than merely requiring that the helper's final state after one completed session cannot be used to impersonate the prover later: for example, a self-erasing one-session oracle for F_x could satisfy such a final-state condition while still allowing the helper to answer the current timed session. Theorem 8 below is therefore a conditional guarantee, not general TF resistance: a helper equipped with such a session-bound evaluator falls outside its hypothesis and would pass every fast round. Section VI-C discusses whether this gap can be closed.

Theorem 8 (TF bound under fresh-bit-unpredictable assistance)**.** *For every time-respecting QPT TF pair $(\mathcal{P}^\star, \mathcal{A})$ with fresh-bit-unpredictable assistance, every round i, and every history H_{i-1},*

$$\Pr[I_i = 1 \mid H_{i-1}] \le \tfrac{1}{2} + \mathrm{negl}(\lambda).$$

Consequently, for

$$\tau = \left\lceil n\left(\tfrac{1}{2} + \varepsilon'\right)\right\rceil$$

with constant $\varepsilon' \in (0, \frac{1}{2})$,

$$\Pr[S \ge \tau] \le \exp\!\left(-2n(\varepsilon' - \mathrm{negl}(\lambda))^2\right).$$

Proof. Fix a round i and a history H_{i-1}. Let $m_i = i\|\gamma_i$. Before round i, only polynomially many PRF inputs can have appeared in the learning phase and in earlier challenge rounds. Since γ_i is sampled uniformly from $\{0,1\}^\lambda$, the probability that m_i collides with one of these previous inputs is negligible.

Condition on the complementary event that m_i is fresh. Suppose that the helper makes the verifier accept in round i with probability p_i. We build a predictor for b_i. The predictor continues from the helper's timely pre-response view. If the helper emits a response qubit in time, the predictor measures this qubit in the verifier's basis a_i, obtains an outcome c_i', and outputs

$$\hat{b}_i := c_i \oplus c_i'.$$

If no timely response qubit is emitted, the predictor outputs an arbitrary bit.

Whenever the verifier accepts, its measurement outcome satisfies

$$c_i' = c_i \oplus b_i,$$

and therefore the predictor outputs $\hat{b}_i = b_i$. Hence the predictor guesses b_i with probability at least p_i. By fresh-bit-unpredictability, this probability is at most $\frac{1}{2} + \mathrm{negl}(\lambda)$. Adding the negligible probability of a repeated input gives

$$\Pr[I_i = 1 \mid H_{i-1}] \le \tfrac{1}{2} + \mathrm{negl}(\lambda).$$

The session bound follows by applying the upper-tail part of Lemma 1 to the indexed sequence of indicators $(I_i)_{i=1}^n$. □

F. Soundness

We now combine the preceding analyses into a session-level soundness theorem in the sense of Definition 6. Recall that a session is accepted if and only if at least $\tau = \lceil n(\frac{1}{2} + \varepsilon')\rceil$ rounds are accepted, where $\varepsilon' \in (0, \frac{1}{2})$ is a fixed constant. Let $I_i \in \{0,1\}$ denote acceptance in round i and $S = \sum_{i=1}^n I_i$.

Theorem 9 (Soundness)**.** *Fix a constant $\varepsilon' \in (0, \frac{1}{2})$ and let*

$$\tau = \left\lceil n\left(\tfrac{1}{2} + \varepsilon'\right)\right\rceil.$$

Assume that $n = n(\lambda)$ is polynomially bounded and satisfies $n(\lambda) = \omega(\log \lambda)$; in particular, $n = \lambda$ is allowed.

Then:

(i) For every time-respecting QPT adversary in the DF experiment,

$$\Pr[S \ge \tau] \le \exp(-2n\varepsilon'^2).$$

(ii) Assuming F is secure against QPT distinguishers making classical oracle queries, for every time-respecting QPT adversary pair in the MF experiment, there exists a negligible function $\varepsilon_{\mathrm{MF}}(\lambda)$ such that

$$\Pr[S \ge \tau] \le \exp\!\left(-2n(\varepsilon' - \varepsilon_{\mathrm{MF}}(\lambda))^2\right) + \varepsilon_{\mathrm{MF}}(\lambda).$$

(iii) For every time-respecting QPT pair $(\mathcal{P}^\star, \mathcal{A})$ with fresh-bit-unpredictable assistance in the TF experiment, there exists a negligible function $\varepsilon_{\mathrm{TF}}(\lambda)$ such that

$$\Pr[S \ge \tau] \le \exp\!\left(-2n(\varepsilon' - \varepsilon_{\mathrm{TF}}(\lambda))^2\right).$$

Consequently, the protocol has negligible DF, MF, and conditional TF false-accept probability under the stated assumptions.

Proof. For DF, Lemma 5 and Lemma 1 give

$$\Pr[S \ge \tau] \le \exp(-2n\varepsilon'^2),$$

as claimed.

For MF, the claim follows from Theorem 7.

For TF, the claim follows directly from Theorem 8.

Since ε' is a fixed positive constant, every negligible function appearing in the MF and TF bounds is at most $\varepsilon'/2$ for all sufficiently large λ. Hence each exponential term is bounded by

$$\exp\bigl(-2n(\varepsilon'/2)^2\bigr) \le \exp(-n\varepsilon'^2/2).$$

Because $n(\lambda) = \omega(\log \lambda)$, this term is negligible in λ, and adding negligible PRF or collision terms preserves negligibility. □

For DF, this is a statistical framework-level soundness guarantee. For MF, it is computational in λ via the PRF assumption. For TF, it is conditional on the fresh-bit-unpredictable assistance condition specified in Section V-E.

G. Parameter choices

To ensure both completeness and soundness, the honest acceptance probability $p(\eta)$ from Theorem 3 must exceed the decision threshold fraction of Theorem 9: $p(\eta) > \frac{1}{2} + \varepsilon'$, with ε' the security slack. Substituting $p(\eta) = 1 - \eta + \frac{1}{2}\eta^2$, this is equivalent to $(\eta - 1)^2 > 2\varepsilon'$, i.e., for $0 \le \eta \le 1$ and $\varepsilon' < 1/2$, to the maximal tolerable depolarizing parameter $\eta < 1 - \sqrt{2\varepsilon'}$. Provided η satisfies this bound, a threshold τ ensuring both completeness and soundness exists.

Table II lists round counts n that make the adversary's session success probability at most 2^{-80}, a common target false-accept probability in the literature. The asymptotic negligible guarantee in Theorem 9 additionally relies on choosing $n(\lambda) = \omega(\log \lambda)$ with n polynomially bounded. The table values use the exponential bound of Theorem 9, ignoring the negligible terms, and are slightly conservative relative to exact binomial tails.

TABLE II
ROUND COUNTS FOR TARGET FALSE-ACCEPT PROBABILITY 2^{-80}, COMPUTED FROM HOEFFDING'S INEQUALITY $\Pr[S \ge \tau] \le \exp\bigl(-2(\tau/n - \frac{1}{2})^2 n\bigr)$ WITH $\tau = \lceil n(\frac{1}{2} + \varepsilon')\rceil$, IGNORING THE NEGLIGIBLE TERMS.

ε'	Threshold τ/n	n for 2^{-80}
0.02	0.52	69,315
0.03	0.53	30,807
0.05	0.55	11,091
0.08	0.58	4,333
0.10	0.60	2,773
0.15	0.65	1,233

VI. DISCUSSION

The security analysis rests on a hidden-bit reduction, an idealized timing model, and a conditional TF hypothesis. This section examines these choices in turn.

A. Role of the quantum challenge

All bounds in Theorem 9 are driven by one classical quantity, the hidden bit $h_i = c_i \oplus b_i$. Under the same atomic-release idealization, a classical variant that transmits c_i as a plaintext bit and expects the response $c_i \oplus b_i$ satisfies the same DF, MF, and conditional TF bounds. In this precise sense the acceptance test is classical, and the per-round guarantees follow from the freshness and unpredictability of b_i.

What the quantum challenge changes is the status of the timing idealization. For a classical challenge, atomic release is contradicted in practice: a receiver can decode the early part of a symbol and commit to a response late, a documented attack class on classical DB channels [23], [24]. These early-detect and late-commit strategies fail against a single qubit in a basis known only to the verifier: it cannot be partially decoded, measured reliably without the basis, or copied for deferred measurement [8], [9]. For the quantum component of the challenge, atomicity is enforced by physics; for the classical component, it remains an engineering assumption.

Remark 1 makes the resulting division of labor exact. The DF bound of $\frac{1}{2}$ survives even if the nonce, the secret bit, and the basis all leak arbitrarily early, because the protection against a distant prover rests on the quantum challenge bit alone. A classical variant loses this guarantee as soon as the challenge bit can be decoded early. The quantum challenge is thus load-bearing exactly where the distance claim is made, while the keyed classical mechanism carries the collusion guarantees. The proofs use nothing more: no-cloning and measurement disturbance enter through the timing model and Remark 1, not through the acceptance-test analysis, which proceeds entirely via the hidden-bit reduction of Lemma 4.

B. Sensitivity of the timing assumption

The atomic-release model of Section IV-C idealizes away message serialization. We examine two relaxations: early leakage of the nonce, and the latency of the prover's buffer.

Early nonce leakage. Suppose useful information about γ_i becomes available a time δ earlier than the model allows, through serialization or partial early decoding. The DF bound is unaffected for every δ, by Remark 1. The MF and TF bounds can fail only through one route: a distant key holder (the honest prover used as a Y^{b_i}-oracle in MF, or the dishonest prover in TF) computes b_i from the leaked nonce and relays it to the attacker near the verifier; without the leak, this bit arrives too late from beyond the distance bound. A head start of δ pays for $c\delta$ of extra signal travel, where c is the signal speed, hence $c\delta/2$ of extra distance each way, so all guarantees of Theorem 9 hold with B replaced by $B' = B + c\delta/2$. Degradation is linear in δ, not abrupt. Leaking a strict prefix of γ_i is weaker still: without the key, b_i remains pseudorandom given any proper prefix of the nonce.

Buffer latency. The prover holds the qubit while it receives the λ-bit nonce and computes the PRF, so the verifier must accept round-trip times longer by this buffer time t_{buf}. Read as pure time of flight, the relaxed deadline widens the certified distance to $B + c\,t_{\mathrm{buf}}/2$. Responding without b_i passes any clock but is capped at $\frac{1}{2}$ per round by the hidden-bit analysis and cannot pass the session; passing the session requires receiving the full nonce and evaluating the PRF. Receiving the nonce takes the same time at every distance, so it cancels from the timing comparison; the only advantage left is faster PRF

hardware. For session-passing provers the certified distance is therefore B plus $c/2$ times the gap between the honest and the fastest adversarial PRF latency, the processing-latency margin familiar from classical DB channels [24]. Releasing the nonce ahead of the qubit to shrink the buffer would be the wrong trade: it converts the buffer time into release slack δ and moves the full $c\,t_{\mathrm{buf}}/2$ into the MF/TF margin B', whereas buffering leaves only the hardware margin.

C. Prospects for unconditional terrorist-fraud resistance

The TF guarantee of Theorem 8 is conditional, and the condition sits close to the conclusion: passing a fast round is equivalent to predicting the hidden bit, so assuming that no timely predictor for b_i exists excludes all but one attack class by hypothesis. The theorem's content is the reduction and the session-level lifting; the open question is whether the excluded class can be ruled out rather than assumed away.

The excluded class is exactly the session-bound evaluator of Section V-E: an apparatus, handed to the helper by the distant prover, that predicts b_i on fresh inputs during the timed session yet confers no advantage in later sessions; it would satisfy final-state non-transferability (Definition 5) while breaking every fast round.

In classical DB, TF resistance is typically obtained by extraction: any helper able to answer enough fast rounds thereby obtains reusable key material [22], [25]. This route is structurally unavailable here: the fast-round responses reveal only outputs of F_x on inputs that never repeat, so a successful helper learns nothing it can reuse; the property that yields the tight MF bound also removes the leakage channel on which extraction relies. Within this design family, TF resistance must come from unclonability of the evaluation capability itself.

The question is therefore whether one-session-only evaluation power for a keyed function can exist at all. If every apparatus that predicts b_i on fresh inputs during one timed session can also do so in later sessions, then non-transferability already implies the fresh-bit-unpredictability hypothesis, and Theorem 8 upgrades to TF soundness under Definition 5 alone. Known results point in both directions: one-time programs cannot be built from quantum information alone, even under computational assumptions, yet they can be built from one-time-memory hardware [26]. Neither result settles our setting, since the impossibility does not cover a helper that shares entanglement with the distant prover and interacts with it between rounds, and the constructions have not been analyzed under fast-phase timing constraints. Whether unconditional TF resistance is achievable or formally obstructed in this model is the central open problem raised by this work, and we are pursuing it in ongoing work.

D. Communication cost and photon loss

Communication. Each fast round carries one qubit and λ classical bits from verifier to prover, and one qubit back, so a session transmits $n\lambda$ classical bits, a factor n more than designs whose separate slow phase costs $O(\lambda)$ bits regardless of round count (Table I). For the round counts of Table II at $\lambda = 128$, this is roughly 10^5 to 10^7 classical bits per session, depending on the slack ε'. This is the price of per-round freshness, which removes the slow phase and makes the MF bound tight (Section IV-C); whether it is acceptable depends on the available classical bandwidth.

Photon loss. The noise model of Section V-A captures depolarization but not loss, in practice the larger effect in single-photon implementations. The natural extension is a loss-tolerant acceptance rule: the verifier counts a round only if a response arrives in time, and accepts the session if at least $m_{\min}$ rounds are counted and at least a fraction $\frac{1}{2} + \varepsilon'$ of them are correct. Post-selection does not help: whether a response arrives is determined by the adversary's view at the response deadline, and the hidden bit is unpredictable from exactly that view, so the per-round bound of Lemma 4 continues to hold conditioned on the round being counted. Applying the argument behind Lemma 1 to the counted indicators yields a session bound with tail exponent proportional to $m_{\min}^2/n$, still exponential in n when $m_{\min}/n$ is a constant; loss and detector inefficiency then enter only the completeness condition. A full loss-calibrated analysis, including dark counts and a calibrated $m_{\min}$, belongs to the experimental validation outlined in Section VII.

VII. Conclusion and Future Work

In this paper, we proposed a conceptually simple QDB protocol that combines a BB84-style qubit challenge with a fresh classical nonce and a one-bit keyed function value. The verifier's fast-round acceptance condition reduces to the hidden bit $h_i = c_i \oplus b_i$, and from this reduction we proved completeness under an i.i.d. effective depolarizing noise model, a single-round DF bound of $1/2$, and MF security from a random-function argument combined with a session-level PRF game hop. For helpers that cannot predict the fresh per-round PRF value at the response deadline, we proved a conditional TF bound of $1/2 + \mathrm{negl}(\lambda)$ per round; Theorem 9 lifts all bounds to the session level. The prover needs only a timed qubit buffer, a conditional Pauli-Y or identity gate, and reflection of the incoming qubit, without any quantum measurements or state preparation. Compared to existing QDB proposals [11]–[13], the protocol thus offers statistical DF security, PRF-based MF security, a conditional TF bound, and lightweight quantum prover operation at the same time.

Several questions remain open: extending the single-verifier, single-prover setting to a multi-party one; a full experimental implementation, validating the noise tolerance and providing concrete latency and throughput figures; whether the conditional TF guarantee can be made unconditional (Section VI-C); and a *fully quantum* protocol that removes the classical γ_i component while retaining tight per-round bounds.

Acknowledgment

This work was supported in part by CyberSecurity Research Flanders with reference number VR20192203, and by the European Commission through the Horizon Europe research and innovation programme under the project *Quantum Security Networks Partnership* (QSNP, grant agreement No. 101114043).

REFERENCES

[1] S. Brands and D. Chaum, "Distance-Bounding Protocols," in *Advances in Cryptology - EUROCRYPT '93*, ser. Lecture Notes in Computer Science, T. Helleseth, Ed., vol. 765. Berlin, Heidelberg: Springer, 1994, pp. 344–359.

[2] G. Avoine, M. A. Bingol, Ş. Kardaş, C. Lauradoux, and B. Martin, "Security of Distance-Bounding: A Survey," *ACM Computing Surveys*, vol. 51, no. 5, 2018.

[3] G. P. Hancke and M. G. Kuhn, "An RFID Distance Bounding Protocol," in *Proceedings of the First International Conference on Security and Privacy for Emerging Areas in Communications Networks (SecureComm 2005)*, 2005, pp. 67–73.

[4] S. Drimer and S. J. Murdoch, "Keep Your Enemies Close: Distance Bounding Against Smartcard Relay Attacks," in *16th USENIX Security Symposium (USENIX Security 07)*. Boston, MA: USENIX Association, Aug. 2007. [Online]. Available: https://www.usenix.org/conference/16th-usenix-security-symposium/keep-your-enemies-close-distance-bounding-against

[5] Y. Desmedt, C. Goutier, and S. Bengio, "Special Uses and Abuses of the Fiat-Shamir Passport Protocol," in *Conference on the theory and application of cryptographic techniques*. Springer, 1987, pp. 21–39.

[6] S. Bengio, G. Brassard, Y. G. Desmedt, C. Goutier, and J.-J. Quisquater, "Secure Implementation of Identification Systems," *Journal of Cryptology*, vol. 4, pp. 175–183, 1991.

[7] C. H. Bennett and G. Brassard, "Quantum Cryptography: Public Key Distribution and Coin Tossing," in *Proceedings of the IEEE International Conference on Computers, Systems and Signal Processing*, Bangalore, India, 1984, pp. 175–179.

[8] W. K. Wootters and W. H. Zurek, "A Single Quantum Cannot Be Cloned," *Nature*, vol. 299, pp. 802–803, 1982.

[9] D. Dieks, "Communication by EPR Devices," *Physics Letters A*, vol. 92, no. 6, pp. 271–272, 1982.

[10] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information: 10th Anniversary Edition*. Cambridge: Cambridge University Press, 2010.

[11] A. Abidin, E. Marin, D. Singelée, and B. Preneel, "Towards Quantum Distance Bounding Protocols," in *Radio Frequency Identification and IoT Security: 12th International Workshop, RFIDSec 2016, Hong Kong, China, November 30–December 2, 2016, Revised Selected Papers 12*. Springer, 2017, pp. 151–162.

[12] A. Abidin, "Quantum Distance Bounding," in *Proceedings of the 12th Conference on Security and Privacy in Wireless and Mobile Networks*, 2019, pp. 233–238.

[13] A. Abidin, K. Eldefrawy, and D. Singelée, "Entanglement-Based Mutual Quantum Distance Bounding," in *International Symposium on Cyber Security, Cryptology, and Machine Learning*. Springer, 2024, pp. 219–235.

[14] K. Bogner, D. Singelée, and A. Abidin, "Entangled States and Bell's Inequality: A New Approach to Quantum Distance Bounding," in *2024 IEEE Symposium on Computers and Communications (ISCC)*. IEEE, 2024, pp. 1–6.

[15] K. Bogner, A. Abidin, and D. Singelée, "Continuous Variable Quantum Distance Bounding," in *IEEE INFOCOM 2025-IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*. IEEE, 2025, pp. 1–6.

[16] K. Bogner, A. Abidin, D. Singelée, and B. Preneel, "Security framework for quantum distance-bounding," *Quantum Information Processing*, vol. 25, no. 5, p. Paper No. 162, 2026.

[17] G. Avoine, X. Bultel, S. Gambs, D. Gérault, P. Lafourcade, C. Onete, and J.-M. Robert, "A Terrorist-Fraud Resistant and Extractor-Free Anonymous Distance-Bounding Protocol," in *Proceedings of the 2017 ACM on Asia conference on computer and communications security*, 2017, pp. 800–814.

[18] S. R. Verschoor, "Quantum Information in Security Protocols," Ph.D. dissertation, University of Waterloo, 2022.

[19] K. Azuma, "Weighted Sums of Certain Dependent Random Variables," *Tohoku Mathematical Journal*, vol. 19, no. 3, pp. 357–367, 1967.

[20] W. Hoeffding, "Probability Inequalities for Sums of Bounded Random Variables," *Journal of the American Statistical Association*, vol. 58, no. 301, pp. 13–30, 1963.

[21] M. Zhandry, "How to Construct Quantum Random Functions," in *2012 IEEE 53rd Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE, 2012, pp. 679–687.

[22] I. Boureanu, A. Mitrokotsa, and S. Vaudenay, "Practical and Provably Secure Distance-Bounding," *Journal of Computer Security*, vol. 23, no. 2, pp. 229–257, 2015.

[23] J. Clulow, G. P. Hancke, M. G. Kuhn, and T. Moore, "So Near and Yet So Far: Distance-Bounding Attacks in Wireless Networks," in *Security and Privacy in Ad-Hoc and Sensor Networks (ESAS 2006)*, ser. Lecture Notes in Computer Science, vol. 4357. Springer, 2006, pp. 83–97.

[24] G. P. Hancke and M. G. Kuhn, "Attacks on Time-of-Flight Distance Bounding Channels," in *Proceedings of the First ACM Conference on Wireless Network Security (WiSec 2008)*. ACM, 2008, pp. 194–202.

[25] M. Fischlin and C. Onete, "Terrorism in Distance Bounding: Modeling Terrorist-Fraud Resistance," in *Applied Cryptography and Network Security (ACNS 2013)*, ser. Lecture Notes in Computer Science, vol. 7954. Springer, 2013, pp. 414–431.

[26] A. Broadbent, G. Gutoski, and D. Stebila, "Quantum One-Time Programs," in *Advances in Cryptology – CRYPTO 2013*, ser. Lecture Notes in Computer Science, vol. 8043. Springer, 2013, pp. 344–360.