Post-Quantum Security of Practical Correlation-Robust Hashing

細山田 光倫 (Akinori Hosoyamada)[1,2], 小菅 悠久 (Haruhisa Kosuge)[1], and 草川 恵太 (Keita Xagawa)[3]

[1] NTT Social Informatics Laboratories, Japan
{hrhs.kosuge,akinori.hosoyamada}@ntt.com
[2] NTT Research Center for Theoretical Quantum Information, Japan
[3] Technology Innovation Institute, UAE
keita.xagawa@tii.ae

September 18, 2026

Abstract. Correlation-robust (CR) hashing and its variants are central components in efficient secure-computation protocols, including OT extension, garbled-circuit optimizations such as Free-XOR and half-gates, and GGM-style tree constructions. In practice, these hashes are typically instantiated from block ciphers such as AES. Prominent constructions include the Matyas-Meyer-Oseas (MMO) construction and its variants such as $\widehat{\mathsf{MMO}}$. Existing analyses, however, are classical and do not cover quantum adversaries with superposition access to the underlying random permutation or ideal cipher.

We establish post-quantum security for these practical block-cipher-based CR hashes. In the quantum ideal cipher model (QICM), we prove multi-user tweakable correlation robustness with leakage (mTCRL) for the MMO construction, and multi-user tweakable circular correlation robustness with leakage (mTCCRL) for $\widehat{\mathsf{MMO}}$ and the Encryption-with-Feed-Forward (EncFF) construction. These results also imply the corresponding leakage-free and single-user guarantees: CR and TCR for MMO, and CR, CCR, TCR, and TCCR for $\widehat{\mathsf{MMO}}$ and EncFF. Security in the quantum random permutation model (QRPM) follows as a special case.

We further verify that, in the QICM/QRPM, the above MMO-type constructions securely instantiate CR-type hashes used in various existing protocol analyses. This covers several representative analyses involving OT extension, half-gates garbling, (correlated) GGM trees, and certain distributed point/comparison function constructions. When the remaining components are post-quantum secure or modeled as ideal functionalities, the resulting protocol instantiations are post-quantum secure in the corresponding model.

Technically, our proof relies on the reprogramming-and-resampling technique of Alagic et al. (Eurocrypt 2022). To handle adaptive key leakage, as needed for some malicious OT-extension analyses, we introduce the conditional min-entropy with leakage (cmel) advantage, which separates leakage-induced entropy loss from the quantum ideal-cipher analysis. Without leakage, our bounds guarantee security up to roughly $2^{\rho/3}$ queries, where ρ is the min-entropy of the secret shift.

Keywords: correlation robustness · post-quantum security · quantum ideal cipher model · quantum random permutation model · block-cipher-based hashing · secure computation

Table of Contents

1 Introduction 2
2 Preliminaries 6
3 Correlation Robustness in the Post-Quantum Setting 9
4 Applications 12
5 Proof of Theorem 3.1 15
A Parameter Index 24
B Proof of Proposition 5.1 25
C The Remaining Part of the Proof for Proposition 5.2 25
D Proof of Proposition 3.1 35
E Proof of Proposition 3.2 36
F Tightness of the Bound in Proposition 3.2 . . . 37
G Comparison of Hybrids 38

1 Introduction

Hash functions are fundamental cryptographic primitives, and their security is traditionally formalized through notions such as preimage resistance, second-preimage resistance, and collision resistance. In many cryptographic constructions, however, these standard notions do not capture all the properties needed for security, motivating stronger or more specialized notions.

Correlation Robustness and More. Correlation robustness (CR) [IKNP03] was introduced to capture a security requirement arising in an extension of oblivious transfer (OT) [Rab81][4]. Roughly speaking, we say a hash function H is CR if $(H(k \oplus x_1), \ldots, H(k \oplus x_m))$ is pseudorandom, where k is a randomly chosen secret key, sometimes called the shift, and $x_1, \ldots, x_m$ can be adversarially chosen. This notion has since found applications beyond OT extension, most notably in the context of the Free-XOR technique [KS08] for garbled circuits (GC) [Yao86], pseudorandom correlation generators (PCG) (e.g., [BCG+19]), as well as in other cryptographic protocols.

Unfortunately, CR alone is insufficient in some settings. Choi et al. [CKKZ12] showed that the Free-XOR technique can be insecure even when the hash function satisfies (a natural variant of) CR, due to a circularity caused by the global offset. To capture the missing requirement, they introduced *circular correlation robustness (CCR)*, where a distinguisher can receive either $H(k \oplus x_i)$ or $H(k \oplus x_i) \oplus k$. Later works extended these notions to the tweakable setting, yielding *tweakable correlation robustness (TCR)* and *tweakable circular correlation robustness (TCCR)* [GKWY20], and then to the multi-instance setting, yielding *multi-instance TCCR (mTCCR)* [GKW+20].

The notions above assume that the secret shift remains hidden throughout the experiment. However, this assumption is too restrictive for some applications. In particular, in malicious-security analyses of OT extension, selective-abort behavior can leak partial information about the shift. This motivates CR-type notions that explicitly account for leakage about the secret shift. Chen and Tessaro [CT21] introduced TCR*, where the entire shift is revealed after the adversary's oracle interaction. Later formulations explicitly modeled partial leakage: Roy [Roy22] studied a TCR-like notion for SoftSpokenOT that incorporates affine selective-abort leakage, while Guo et al. [GWYY25] formulated mTCCRL, a leakage variant of mTCCR that allows predicate-based leakage queries on the secret shifts.

Constructions. A random oracle (RO) immediately satisfies CR and its stronger variants, including CCR, TCR, and TCCR [IKNP03, GKWY20, GKW+20, Roy22, GWYY25], since its outputs on distinct inputs are independent and uniform. This observation underlies the original ROM analysis of OT extension by Ishai et al. [IKNP03], and subsequent works on garbled circuits and correlated primitives have often relied on RO-based arguments [KS08, CKKZ12, ZRE15].

While the ROM provides a clean abstraction, practical instantiations of correlation-robust hashing almost exclusively rely on simple block-cipher-based constructions tailored specifically to CR-type notions, due to the efficient hardware-assisted implementations of block ciphers, most notably AES-NI for AES [Nat01]. Their security is usually justified in the ideal cipher model (ICM) or the random permutation model (RPM): the former models the entire block cipher E as an ideal cipher (IC), whereas the latter models a fixed-key instance $E(k, \cdot)$, where k is public, as a random permutation (RP).

In this line of work, Guo et al. [GKWY20] showed that the Matyas–Meyer–Oseas (MMO) construction and its variants, when instantiated with a random permutation, satisfy CR, CCR, and TCCR in the RPM. Subsequently, Guo et al. [GKW+20] strengthened this analysis by showing that a variant of MMO instantiated in the ICM satisfies mTCCR. These results provide a rigorous foundation for the widespread use of AES-based correlation-robust hashing in practical secure-computation systems. Table 1 summarizes the major tailored block-cipher-based constructions, together with their known security guarantees.

Post-Quantum Security. The transition to post-quantum cryptography [Nat24a] motivates revisiting not only public-key primitives but also the symmetric-key components used inside cryptographic protocols. Block-cipher-based hash constructions tailored to CR-type notions are a particularly relevant class of such components: they appear in performance-critical parts of secure-computation protocols, including OT extension, garbled circuits, and GGM tree constructions.

However, the existing analyses of these hashes in the RPM/ICM are classical and do not automatically justify security against quantum adversaries with superposition access to the underlying ideal cipher or random

[4] OT extension [Bea96] is a technique that amplifies a small number of base OTs, typically realized using public-key cryptography, into a large number of OTs using only symmetric-key operations or efficient computations.

Table 1. Summary of correlation robust hash functions. $\pi : \{0,1\}^n \to \{0,1\}^n$ is a random permutation. $E : \{0,1\}^m \times \{0,1\}^n \to \{0,1\}^n$ is an ideal cipher. $\sigma : \{0,1\}^n \to \{0,1\}^n$ is a linear orthomorphism. In H^{π}_{CT21}, the symbol $\otimes$ denotes the multiplication in $\mathbb{F}_{2^n}$. In H^{E}_{Roy22}, $\iota : \mathbb{F}_q^{n_C} \to \{0,1\}^n$ is an injection, where n_C denotes the block length of the linear code used in the construction. The checkmarks indicate that this work shows the post-quantum security of the construction in the corresponding model (either QICM or QRPM), including the multi-user setting with adaptive leakage. The checkmark for $H^{E}_{\mathrm{Roy22}}(y,t)$ means that mTCRL is proven for MMO^E, thereby covering the binary one-dimensional repetition-code case of Roy [Roy22]; Roy's general TCR-like notion for arbitrary codes is out of our scope. For EncFF^E, we prove mTCCRL without onewayness.

Name	Construction	Model	Security	PQ Security (This Work)
$\mathsf{MMO}^{\pi}(x)$	$\pi(x) \oplus x$	RPM	CR [GKWY20]	✓
$\widehat{\mathsf{MMO}}^{\pi}(x)$	$\pi(\sigma(x)) \oplus \sigma(x)$	RPM	CCR [GKWY20]	✓
$\mathsf{TMMO}^{\pi}(x,t)$	$\pi(\pi(x) \oplus t) \oplus \pi(x)$	RPM	TCCR [GKWY20]	–
$\mathsf{FPTP}^{\pi}(x,t)$	$\pi(\pi(\sigma(x)) \oplus t) \oplus \sigma(x)$	RPM	TCCR [CT21]	–
$H^{\pi}_{\mathrm{CT21}}(x,t)$	$\pi(x \otimes t) \oplus (x \otimes t)$	RPM	TCCR [CT21]	–
$\widehat{\mathsf{MMO}}^{E}(x,t)$	$E(t,\sigma(x)) \oplus \sigma(x)$	ICM	mTCCR [GKW+20], mTCCRL [GWYY25]	✓
$H^{E}_{\mathrm{Roy22}}(y,t)$	$E(t,\iota(y)) \oplus \iota(y)$ $(= \mathsf{MMO}^E(\iota(y),t))$	ICM	TCR-like with leakage [Roy22]	✓* (mTCRL for MMO^E)
$\mathsf{EncFF}^E(x,t)$	$E(t,x) \oplus \sigma(x)$	ICM	TCCR with onewayness [KX26]	✓(mTCCRL)
$H^{E}_{\mathrm{CGW+24}}(x,t)$	$\widehat{\mathsf{MMO}}^{E}$-based construction	ICM	extractable-mTCCR [CGW+24]	–

permutation. The difficulty is not merely that the adversary is quantum: the classical RPM/ICM analyses of these tailored constructions typically rely on H-coefficient [Pat09] or transcript-based arguments, which classify classical oracle transcripts and compare their probabilities. Such arguments do not directly carry over to the quantum ideal cipher model (QICM) [HY18] and quantum random permutation model (QRPM), where both forward and inverse queries to the underlying primitive may be made in quantum superposition. This raises the following question:

Do block-cipher-based hashes admit post-quantum CR-type security proofs?

1.1 Our Contributions

This paper establishes post-quantum CR-type security for the major single-call block-cipher-based constructions studied in prior work, namely MMO and $\widehat{\mathsf{MMO}}$. In the QICM, we prove that MMO^E satisfies multi-user TCR with leakage (mTCRL), while $\widehat{\mathsf{MMO}}^E$ satisfies multi-user TCCR with leakage (mTCCRL). Our proof for $\widehat{\mathsf{MMO}}^E$ is stated in a general form that also establishes mTCCRL security of the Encryption-with-Feed-Forward construction (EncFF^E) [BBFR24]. These results also imply the corresponding leakage-free and single-user guarantees: CR and TCR for MMO, and CR, CCR, TCR, and TCCR for $\widehat{\mathsf{MMO}}$ and EncFF. The corresponding guarantees (without tweak) in the QRPM follow as a special case. See also Table 1.

Our security model corresponds to the so-called Q1 setting [KLLN16]: the adversary has quantum superposition access to the underlying ideal cipher (or random permutation), while its queries to the CR-type hash oracle and the leakage oracle are classical. This Q1 setting captures the way CR-type hash functions are invoked in the applications considered later, where the protocol evaluates the hash function on classical values and uses classical outputs. Our results do not establish Q2 security [KLLN16], where quantum superposition queries to the hash oracle itself would also be allowed.

When leakage is absent, our bound guarantees security as long as $q_E, q_C \ll 2^{\rho/3}$, where q_E and q_C denote the numbers of quantum queries to the ideal cipher $E^{\pm}$ and classical queries to the hash oracle, respectively, and ρ is the min-entropy of the key (secret value). In the practically relevant setting of $q_E \gg q_C$, the dominant q_E-dependent term matches known attacks for constant distinguishing advantage, although a quadratic gap remains at lower advantages and matching attacks are not known for the remaining terms.

Table 2. Applicability of mTCCRL to existing constructions. A check mark indicates that the corresponding CR-type hash-replacement step can be justified in the QICM/QRPM when the hash is instantiated by one of the MMO-type constructions covered by our result. For Roy [Roy22], our result applies to the binary one-dimensional repetition-code case.

Application	Hash Assumption	QRPM/ QICM
OT extension (semi-honest) [GKWY20]	CR	✓
OT extension (malicious) [GWYY25, Roy22]	mTCCRL/TCRL	✓
OT extension (half-tree/simulation-based) [GYW+23, CGG+25]	–	×
Garbled circuits (half-gates [ZRE15, GKW+20])	(m)TCCR	✓
Garbled circuits (three-halves [RR21])	RTCCR	×
GGM / PPRF in [GKWY20, BCC+24]	(T)CR	✓
Correlated GGM / PPRF in [GYW+23]	(T)CCR	✓
Vector commitment from GGM [BCD25, WKK+25]	(T)CCR	×
DPF / DCF (half-tree [GYW+23])	CCR	✓

For leakage, we follow the models of [Roy22, GWYY25], where an adversary may adaptively query predicates P on the secret key k and receive $P(k) \in \{0, 1\}$, and the hash oracle associated with k stops responding once a predicate query returns 0.

We stress that the leakage considered in this paper should not be understood as physical side-channel leakage. Its main motivation in our setting comes from maliciously secure OT extension, where selective-abort or selective-failure behavior can reveal partial information about the secret shift [Roy22, GWYY25]. Consequently, obtaining post-quantum guarantees for such applications naturally requires handling quantum access to the underlying ideal primitive and this protocol-induced leakage within the same security analysis.

Applications. Our results can be used to instantiate CR-type hash functions appearing in many practical secure-computation protocols by the MMO-type constructions analyzed in this work, while retaining security against quantum adversaries in the QRPM/QICM. More precisely, suppose that a classical security proof for a protocol uses the random permutation or ideal cipher only through CR-type hash invocations, and uses the CR-type property only to justify replacing the corresponding hash outputs with random strings. Then, when these hash invocations are instantiated by one of our covered MMO-type constructions, our results justify the same hash-replacement step in the QRPM/QICM. If the remaining components of the protocol are post-quantum secure or are treated as ideal functionalities, this yields post-quantum security of the resulting protocol instantiation under the corresponding model and composition theorem.

These conditions are satisfied by the security proofs of several representative applications. For semi-honest OT extension, our results justify the CR-dependent hash-replacement steps used in the security analyses of [GKWY20]. For malicious security of OT extension, our results can be combined with the mTCCRL-based analysis of Guo et al. [GWYY25] and the TCR-like analysis of Roy [Roy22]. In the latter case, our result applies to the binary one-dimensional repetition-code setting, where Roy's notion matches our TCRL notion. Our results also apply to the mTCCR-based analyses of half-gates garbling [ZRE15, GKW+20], GGM-tree-based puncturable pseudorandom function (PPRF) upon (T)CR hashes [GKWY20, BCC+24], correlated-GGM-tree-based PPRF upon (T)CCR hashes [GYW+23], and CCR-based DPF/DCF components in the half-tree framework [GI14, BGI15, BCG+21, GYW+23].

We also identify the limits of this implication. For example, our results do not cover proofs that rely on additional RPM/ICM-specific simulation arguments [GYW+23, CGG+25], full subspace-VOLE analyses for arbitrary linear codes [Roy22], RTCCR-based three-halves garbling [RR21], or full signature security claims that additionally require binding or extractable-binding properties of commitments [BCD25, WKK+25]. Thus, the contribution is not a black-box post-quantum composition theorem, but a post-quantum justification for the use of CR-type hash constructions in a wide range of practical protocols. Table 2 summarizes the covered and uncovered cases, and Section 4 provides more detailed discussions.

Proof Techniques. At a high level, our proof proceeds in two steps. First, we reduce the correlation robustness of the hash constructions to the multi-key security of an Even-Mansour-like tweakable block cipher. We then

analyze this tweakable block cipher in the QICM using the reprogramming-and-resampling paradigm introduced by Alagic et al. [ABKM22].

This analysis is not a black-box application of existing post-quantum analyses of Even-Mansour [ABKM22], tweakable Even-Mansour [ABK+24, Hos25], or FX-type constructions [SI26a, ABMS26]. Our setting simultaneously involves an ideal cipher, possibly low-entropy and non-identically distributed keys, multi-key security, the query restrictions inherited from TCCR security, and adaptive leakage. The main additional difficulty comes from leakage: leaking queries adaptively change the conditional distributions of the secret keys and may occur between quantum primitive queries.

To handle this difficulty, we separate the information-theoretic effect of leakage from the quantum ideal-cipher analysis. Since the amount of information leaked depends on the key distributions and on the allowed predicates, we isolate this effect through a quantity that we call the conditional min-entropy with leakage (cmel) advantage. The cmel advantage controls the probability that the posterior min-entropy of some key becomes too small, and our main security bounds for mTCRL and mTCCRL are stated in terms of this quantity.

Crucially, however, this does not mean that one can simply combine a leakage-free QICM proof with a separate analysis of leakage. Leaking queries and quantum primitive queries may be interleaved, so the posterior key distributions can change after key-dependent reprogramming or resampling has already been introduced. To handle this interaction, we introduce additional sub-hybrids that move the reprogramming and resampling operations across intervening leaking queries. Together with the cmel advantage, these sub-hybrids allow us to incorporate adaptive leakage into the reprogramming-and-resampling technique.

Finally, we give a concrete upper bound on the cmel advantage for a practically relevant setting: keys uniformly distributed over affine subspaces and linear predicates. These choices are particularly motivated by the binary one-dimensional repetition-code case of SoftSpokenOT [Roy22], which is the practically relevant case covered by our result. We also show a general upper bound that is looser but applicable to any key distributions and any set of predicates.

1.2 Related Works

Definitions. Starting from correlation robustness (CR) [IKNP03], several stronger variants have been introduced for different applications, including CCR, TCR, TCCR, their multi-user and leakage variants, and simulation-based notions [CKKZ12, GKWY20, GKW+20, CT21, GWYY25, CGG+25]. Extensions to arithmetic settings and related notions such as correlated-input and related-key security have also been studied [KKKS15, BCG+19, Roy22, GOR11, BK03].

On Conditional Min-Entropy under Leakage. The use of conditional or residual min-entropy to quantify information leaked through side information is not new. Related ideas appear in the study of fuzzy extractors and secure sketches [DORS08], and are widely used in leakage-resilient cryptography (in, e.g., [DP08, Pie09, ADW09, NS12, BK12]). Our cmel advantage tailors this common entropy-loss viewpoint to the predicate-leakage setting of mTCCRL: it isolates the information-theoretic effect of adaptive leakage as a separate adversarial quantity, which is then incorporated into the mTCCRL security bound.

Other Constructions of Correlation-Robust Hashing. Standard-model constructions of correlation-robust hashing are known under strong assumptions such as multilinear maps and indistinguishability obfuscation [ABP19, AAB+25], but are substantially less efficient than the symmetric-key constructions considered here.

Another generic approach is to replace the random oracle in an RO-based analysis by a hash construction that is indifferentiable from an RO [MRH04, RSS11, CDMP05, BDPV08, CLL19, BN18, GBJ+23, GBK16]. These constructions offer a generic route to build CR-type hashes from RO-based analyses, but they are not the constructions typically used in practical applications of correlation-robust hashing. Compared with the tailored block-cipher-based constructions in Table 1, such generic instantiations typically require larger parameters or more primitive calls to achieve comparable concrete security for the CR-type properties relevant to major applications.

Post-Quantum Security of Two-Party Secure-Computation Protocols. Büscher et al. [BDK+20] studied secure two-party computation against quantum adversaries. They developed a lattice-based OT protocol, analyzed

the post-quantum security of the OT extension of Ishai et al. [IKNP03] in the quantum random oracle model (QROM) [BDF+11], and proved that Yao's garbled-circuit protocol [Yao86] remains post-quantum secure when instantiated with post-quantum-secure building blocks. A more recent work by Nakamura and Kawachi [NK26] announced an improvement of this QROM security bound on the receiver's distinguishing advantage for the OT extension to a tight bound. To show the security in the QROM, Büscher et al. [BDK+20] used the one-way to hiding lemma [Unr15, AHU19], while Nakamura and Kawachi [NK26] used a purified-oracle analysis. Note that these QROM analyses treat hash functions as monolithic ideal primitives. In contrast, our focus is on instantiating those hash functions with concrete block-cipher-based constructions in the QICM/QRPM. Thus, our results are complementary: they provide post-quantum justification for practical MMO-type hash instantiations used in the CR-type hash-replacement steps of OT extension, garbling, and related secure-computation protocols.

On Post-Quantum Security Proofs of Keyed Symmetric Cryptosystems in the QRPM and QICM. Jaeger, Song, and Tessaro [JST21] studied the post-quantum security of the FX construction in the QICM, and proved the security against non-adaptive adversaries, using the one-way to hiding lemma [AHU19].

Regarding the reprogramming-and-resampling proof method, Alagic et al. [ABKM22] proved the post-quantum security of the Even-Mansour (EM) cipher [EM93] in the quantum random permutation model (QRPM) when the key has full entropy, i.e., entropy equal to the block length. Later, Alagic et al. [ABK+24] proved the post-quantum security of the Tweakable Even-Mansour (TEM) cipher [CLS15] in the QRPM. Their result covers both the full-entropy-key setting and the setting where a low-entropy key is first expanded by a permutation before being used; the latter construction is called TEM-KX. They also proved the post-quantum security of several modes, including Chaskey [MMV+14, Int19], in the QRPM by reducing the security to TEM(-KX). Hosoyamada [Hos25] observed that the post-quantum security of TEM with low-entropy keys does not follow directly from the security of TEM-KX, and gave a dedicated proof for TEM with low-entropy keys in the QRPM. In the same paper, the post-quantum security of keyed sponge-based constructions such as KMAC [KCP16] and Ascon [TMK+25] is also shown in the QRPM by reducing their security to TEM with low-entropy keys. More recently, Shiba and Iwata [SI26a] proved post-quantum multi-key security of TEM-KX in the QRPM and of the FX construction [KR96] in the QICM. Concurrently and independently, Alagic et al. [ABMS26] developed a general framework for post-quantum security proofs of block-cipher constructions. In particular, they proved QICM security of FX, some tweakable block ciphers, and a broad class of block-cipher modes. A key technical ingredient of their work is an ideal-cipher resampling lemma, which is used in our proof. The proof method is also used to prove the security of (tweakable) key-alternating Feistel ciphers [BBC+25, SI26b]. These results do not directly cover our setting, which simultaneously involves possibly low-entropy keys, multi-user security, (T)CCR-specific query restrictions, and adaptive leakage.

2 Preliminaries

2.1 Notations

The set of all permutations over $\{0,1\}^n$ is denoted by $\mathsf{Perm}(n)$. A permutation $\sigma : \{0,1\}^n \to \{0,1\}^n$ is called an *orthomorphism* if the map defined by $x \mapsto \sigma(x) \oplus x$ is also a permutation. We say that a function $E : \{0,1\}^m \times \{0,1\}^n \to \{0,1\}^n$ is an n-bit block cipher with m-bit keys if $E(k,\cdot)$ is a permutation for all $k \in \{0,1\}^m$. The inverse of E is denoted by E^{-1} (i.e., E^{-1} is a function from $\{0,1\}^m \times \{0,1\}^n$ to $\{0,1\}^n$ such that $E^{-1}(k, E(k,x)) = x$ for all k, x). The set of all n-bit block ciphers with m-bit keys is denoted by $\mathsf{BC}(m,n)$. For a (quantum or classical) algorithm $\mathcal{A}$, by $x \leftarrow \mathcal{A}$ we denote the event that an output x is returned when $\mathcal{A}$ is executed. For a distribution D, by $x \leftarrow D$ we denote that x is sampled according to the distribution.

The quantum oracle of a function $f : \{0,1\}^m \to \{0,1\}^n$ is the unitary operator U_f such that $U_f |x\rangle |y\rangle = |x\rangle |y \oplus f(x)\rangle$ for all $(x,y) \in \{0,1\}^m \times \{0,1\}^n$. For $\pi \in \mathsf{Perm}(n)$ (resp., $E \in \mathsf{BC}(m,n)$), we use notations like $\mathcal{A}^{|\pi^{\pm}\rangle}$ (resp., $\mathcal{A}^{|E^{\pm}\rangle}$) to indicate that the quantum oracles of π and π^{-1} (resp., E and E^{-1}) are available to an algorithm $\mathcal{A}$.

Quantum Random Permutation Model and Quantum Ideal Cipher Model. When we discuss a security notion in the quantum random permutation model (QRPM), we assume that a random permutation π is sampled from $\mathsf{Perm}(n)$ uniformly at random, and that the quantum oracles of π and π^{-1} are available to an adversary. Similarly, in the quantum ideal cipher model (QICM), a block cipher E is sampled from $\mathsf{BC}(m,n)$ uniformly at random, and an adversary is given quantum oracle access to E and E^{-1}. (The classical random permutation model (RPM) and ideal cipher model (ICM) are defined by restricting the oracle access to classical.) (Q)RPM can be regarded as the

special case of (Q)ICM when $m = 0$. We say that a predicate P (i.e., a Boolean function) is *oracle-free* if it can be evaluated without making oracle queries to E and E^{-1} in the (Q)ICM (or π and π^{-1} in the (Q)RPM).

Affine Subspaces and Linear Predicates. We often identify the set of n-bit strings $\{0,1\}^n$ with the n-dimensional vector space $\mathbb{F}_2^n$ (the addition is the XOR operation), without any mention. For $a, b \in \{0,1\}^n$, we denote their dot product by $\langle a, b\rangle$. We say that a subset S of $\{0,1\}^n$ is a d-dimensional affine subspace if there is a $\mathbb{F}_2$-vector subspace $V \subset \{0,1\}^n$ and a $c \in \{0,1\}^n$ such that $\dim(V) = d$ and $S = V \oplus c := \{v \oplus c \mid v \in V\}$. In addition, we say that a predicate P over $\{0,1\}^n$ (or equivalently, a Boolean function $P : \{0,1\}^n \to \{0,1\}$) is linear if there exist some $a \in \{0,1\}^n$ and $b \in \{0,1\}$ such that $P(k) = 1$ iff $\langle a, k\rangle = b$.

2.2 Correlation Robustness and Related Notions

Next, we recall security notions related to correlation robustness. We describe first multi-user tweakable circular correlation robustness with leakage (mTCCRL) of Guo et al. [GWYY25], which is the most general notion for which we prove security, and then define others as its special cases.

mTCCRL, Adapted to the QICM. Let H^E be a function from $\{0,1\}^n \times \{0,1\}^m$ to $\{0,1\}^n$, and $\mathbf{R} = (\mathcal{R}_1, \dots, \mathcal{R}_u)$ be a tuple of distributions over $\{0,1\}^n$, and $\mathcal{P}$ be a set of oracle-free predicates defined over $\{0,1\}^n$. Suppose that u keys $\mathbf{k} = (k_1, \dots, k_u)$ are sampled as $k_\alpha \leftarrow \mathcal{R}_\alpha$ for each $\alpha = 1, \dots, u$. Define $\mathcal{O}_{H,\mathbf{k}}^{\mathrm{mTCCRL}} : \{1, \dots, u\} \times \{0,1\}^n \times \{0,1\}^m \times \{0,1\} \to \{0,1\}^n$ and $\mathcal{L}_{\mathbf{k}} : \{1, \dots, u\} \times \mathcal{P} \to \{0,1\}$ by $\mathcal{O}_{H,\mathbf{k}}^{\mathrm{mTCCRL}}(\alpha; x, t, b) := H(x \oplus k_\alpha, t) \oplus b \cdot k_\alpha$ and $\mathcal{L}_{\mathbf{k}}(\alpha; P) := P(k_\alpha)$. For an adversary $\mathcal{A}$, we define its mTCCRL advantage as

$$\mathsf{Adv}_{H,\mathbf{R},\mathcal{P}}^{\mathrm{mTCCRL}}(\mathcal{A}) := \left| \Pr_{\mathbf{k},E}\left[1 \leftarrow \mathcal{A}^{\mathcal{O}_{H,\mathbf{k}}^{\mathrm{mTCCRL}}, \mathcal{L}_{\mathbf{k}}, |E^{\pm}\rangle}\right] - \Pr_{f,\mathbf{k},E}\left[1 \leftarrow \mathcal{A}^{f, \mathcal{L}_{\mathbf{k}}, |E^{\pm}\rangle}\right] \right|. \tag{1}$$

Here, E is an ideal cipher ($\in \mathsf{BC}(m,n)$), and $f : \{1, \dots, u\} \times \{0,1\}^n \times \{0,1\}^m \times \{0,1\} \to \{0,1\}^n$ is a random function. The first two oracles are classical, while the oracle of E is quantum. In addition, the following two conditions are assumed.

1. For any (α, x, t), $\mathcal{A}$ is prohibited from querying both $(\alpha; x, t, 0)$ and $(\alpha; x, t, 1)$ to $\mathcal{O}_{H,\mathbf{k}}^{\mathrm{mTCCRL}}$ or f.
2. Once $\mathcal{L}_{\mathbf{k}}$ returns 0 to a query $(\alpha; P)$, $\mathcal{A}$ is subsequently prohibited from querying $(\alpha; *)$ to $\mathcal{L}_{\mathbf{k}}$ and $(\alpha; *, *, *)$ to $\mathcal{O}_{H,\mathbf{k}}^{\mathrm{mTCCRL}}$ or f.

Following the previous works (especially, [GKW+20], [Roy22], and [GWYY25]), we always assume that predicates are oracle-free, and the keys $k_1, \dots, k_u$ are mutually independent.

We refer to the oracle corresponding to $\mathcal{O}_{H,\mathbf{k}}^{\mathrm{mTCCRL}}/f$ given to $\mathcal{A}$ as the *construction oracle*, and $\mathcal{L}_{\mathbf{k}}$ as the *leaking oracle*. Queries to these oracles are called *construction queries* and *leaking queries*, respectively. In addition, we refer to the quantum oracle corresponding to $|E^{\pm}\rangle$ as the *primitive oracle*, and queries to the primitive oracle as *primitive queries*.

We say that an adversary $\mathcal{A}$ against mTCCRL security is $(q_E, q_C, u, M, q_{L,\max}, Q_{L,\max}, \mu)$-bounded if

- the number of the keys considered is u,
- $\mathcal{A}$ makes at most q_E quantum primitive queries,
- $\mathcal{A}$ makes at most q_C construction queries in total, where queries of the form $(*; *, t, *)$ are made at most μ for any t,
- $\mathcal{A}$ makes at most $q_{L,\max}$ leaking queries in total, involving at most M distinct values of α, and at most $Q_{L,\max}$ leaking queries for any fixed α.

See also Table 4 for an index of main parameters used throughout the paper.

mTCCR. By removing the leaking oracle $\mathcal{L}_{\mathbf{k}}$ (or equivalently, if $\mathcal{P} = \emptyset$), we obtain the notion of multi-user tweakable circular correlation robustness (mTCCR). The mTCCR advantage, denoted as $\mathsf{Adv}_{H,\mathbf{R}}^{\mathrm{mTCCR}}(\mathcal{A})$, is defined as in Eq. (1) by removing $\mathcal{L}_{\mathbf{k}}$. We say that an adversary $\mathcal{A}$ against mTCCR security is (q_E, q_C, u, μ)-bounded if the number of the keys considered is u, $\mathcal{A}$ makes at most q_E quantum primitive queries, and $\mathcal{A}$ makes at most q_C construction queries in total, where queries of the form $(*; *, t, *)$ are made at most μ for any t.

Table 3. Summary of CR, CCR, TCR, and TCCR. TCCR is the strongest, and CR is the weakest. In general, CCR and TCR are incomparable. For each notion, we can also consider multi-user and/or leakage versions.

	CR	CCR	TCR	TCCR
tweak	no ($m = 0$)	no ($m = 0$)	yes ($m \geq 0$)	yes ($m \geq 0$)
bit b	fixed to 0	0 or 1	fixed to 0	0 or 1

mTCRL/mTCR. By fixing the bit b to 0 in construction queries in the definition of mTCCRL (resp., mTCCR), we obtain mTCRL (resp., mTCR). We define $(q_E, q_C, u, M, q_{L,\max}, Q_{L,\max}, \mu)$-bounded adversaries (resp., (q_E, q_C, u, μ)-bounded adversaries) as before.

On Further Weaker Notions. When $u = 1$, mTCCR(L) and mTCR(L) are referred to just as TCCR(L) and TCR(L). In addition, CCR (resp., CR) is the special case of TCCR (resp., TCR) when $m = 0$, i.e., there is no tweak t (see also Table 3). In addition, multi-user and/or leakage versions of CCR and CR can also be defined analogously. We define the advantages $\mathsf{Adv}^{\mathrm{CCR}}_{H,\mathcal{R}}(\mathcal{A})$, $\mathsf{Adv}^{\mathrm{CR}}_{H,\mathcal{R}}(\mathcal{A})$, etc., similarly to the advantages defined above.

So far, we have discussed the security notions in the QICM, but they can similarly be defined in the classical ICM and (Q)RPM.

Remark 2.1. In our "multi-user" security notions, the key distributions $\mathcal{R}_1, \ldots, \mathcal{R}_u$ are not required to be identical. This follows the formulation of [GWYY25]. In contrast, some previous works consider the restricted case where all key distributions are identical, such as the "multi-instance" security notion of [GKW+20]. Proving the security in the latter setting is often sufficient, but our results in the subsequent section are stated for the former, more general setting. Our model allows arbitrary and non-identical marginal key distributions, but does not cover correlations among distinct users' keys.

When the concrete parameters are not important, we simply say that H^E is mTCCRL-secure if its mTCCRL advantage is negligible for every adversary making polynomially many queries, with respect to the class of key distributions and predicates under consideration. We use the same shorthand for the weaker notions (e.g., mTCCR, mTCR, and CR), and sometimes simply say that H^E is secure when the intended notion is clear from the context.

2.3 CR-Type Hash Constructions and Known Results

The hash constructions this paper focuses on are MMO, $\widehat{\mathsf{MMO}}$, and EncFF.

The first function MMO is defined by $\mathsf{MMO}^E(x, t) := E(t, x) \oplus x$. In the classical setting, Guo et al. [GKWY20] proved the following (single-user, leakage-free) CR security bound when $m = 0$ and E is just a single permutation π.

Theorem 2.1 ([GKWY20, Theorem 4]). *Suppose $m = 0$. For any adversary $\mathcal{A}$ making q_C construction queries and q_E primitive queries, and for any key distribution $\mathcal{R}$ with min-entropy at least ρ, $\mathsf{Adv}^{\mathrm{CR}}_{\mathsf{MMO}^\pi,\mathcal{R}}(\mathcal{A}) \leq 2q_C q_E/2^\rho + q_C^2/2^{n+1}$ holds in the classical RPM.*

Roy [Roy22, Proposition 2.7] also proved that MMO^E satisfies a stronger TCR-like notion for $m \geq 0$ in the ICM.

Note that MMO is not CCR: If a construction query (x, t, b) with $b = 1$ is allowed in the real world, then the oracle returns the value $y := \mathsf{MMO}^E(x \oplus k, t) \oplus k = E(t, x \oplus k) \oplus x$. Once this value is obtained, we can easily recover the key because $E^{-1}(t, y \oplus x) \oplus x = k$.

The second function $\widehat{\mathsf{MMO}}$ is defined by $\widehat{\mathsf{MMO}}^E(x, t) := E(t, \sigma(x)) \oplus \sigma(x)$, where σ is a linear orthomorphism. In the classical setting, Guo et al. [GKWY20] showed a single-user CCR security bound of $\widehat{\mathsf{MMO}}$ when $m = 0$ (i.e., E is a single permutation π, in the random permutation model), and later extended the result into the mTCCR security with $m \geq 0$ in the ICM [GKW+20] as follows.

Theorem 2.2 ([GKW+20, Theorem 2]). *Suppose that σ is a linear orthomorphism, E is an ideal cipher, and that the key distributions $\mathcal{R}_1, \ldots, \mathcal{R}_u$ are identical. Then, for any (q_E, q_C, u, μ)-bounded adversary $\mathcal{A}$, the mTCCR advantage regarding $\widehat{\mathsf{MMO}}^E$ is upper bounded by $\varepsilon := 2\mu q_E/2^\rho + (\mu - 1)q_C/2^\rho$ in the classical ICM, where ρ is the min-entropy of the key distribution.*

Guo et al. [GWYY25] further extended the result in the presence of leakage as follows.

Theorem 2.3 ([GWYY25, Theorem 1]). *Let $H^E := \widehat{\mathsf{MMO}}^E$. Suppose that each of $\mathcal{R}_1, \ldots, \mathcal{R}_u$ is the uniform distribution over a subset of $\{0,1\}^n$ of size $\geq 2^\rho$. Then, in the classical ICM, for any $(q_E, q_C, u, M, q_{L,\max}, Q_{L,\max}, \mu)$-bounded adversary $\mathcal{A}$ and any set $\mathcal{P}$ of oracle-free predicates,*

$$\mathsf{Adv}^{\mathrm{mTCCRL}}_{H,\mathbf{R},\mathcal{P}}(\mathcal{A}) \leq \frac{2q_E(M+1)(q_{L,\max}+1)\mu + q_C(M+2)(q_{L,\max}+1)(\mu-1)}{2^\rho}.$$

The third function EncFF is defined by $\mathsf{EncFF}^E(x,t) := E(t,x) \oplus \sigma(x)$, where σ is again a linear orthomorphism. This function is used in a GGM-tree component of MQOM [BBFR24], a Round-3 candidate in NIST's additional post-quantum digital-signature standardization process [ABC+26]. A classical single-user, leakage-free TCCR guarantee for this construction can be obtained from the TCCR-with-one-wayness result of [KX26, Theorem D.2] by removing the auxiliary one-way-function information.

3 Correlation Robustness in the Post-Quantum Setting

This section shows the mTCR(L) security of MMO and the mTCCR(L) security of $\widehat{\mathsf{MMO}}$ and EncFF in the QICM. These results can be viewed as post-quantum counterparts/extensions of the classical bounds recalled in the previous section, in particular Theorems 2.1–2.3.

Before stating the main theorem, we first introduce a quantity that isolates the information-theoretic effect of leakage on the secret keys. Adaptive leaking queries change the posterior key distributions, and the resulting entropy loss depends on the underlying key distributions and the allowed class of predicates. Treating this entropy loss directly inside the QICM proof, together with quantum queries to the ideal cipher, would lead to substantial technical complications. We therefore introduce the *conditional min-entropy with leakage* (cmel) advantage, which measures the probability that the conditional min-entropy of some key falls below a prescribed threshold. Our main theorem uses this quantity as an interface between application-specific leakage analysis and the quantum ideal-cipher analysis.

This separation does not mean that leakage can simply be analyzed independently and then combined with a leakage-free QICM proof: leaking queries may be interleaved with quantum primitive queries and hence change the key distributions relevant to the quantum analysis. Further details on the role of the cmel advantage in the proof are given in Section 5.2.

In what follows, Section 3.1 formally defines the cmel advantage and gives bounds for relevant classes of key distributions and leakage predicates. Then, Section 3.2 states our main theorem on mTCRL/mTCCRL security with a bound that depends explicitly on this cmel advantage.

3.1 On Conditional Min-Entropy in the Presence of Leakage

Let $\rho > 0$ be a parameter, $\mathbf{R} := (\mathcal{R}_1, \ldots, \mathcal{R}_u)$ be a tuple of distributions over $\{0,1\}^n$, and $\mathcal{P}$ be a set of (oracle-free[5]) predicates over $\{0,1\}^n$. For a key vector $\mathbf{k} = (k_1, \ldots, k_u)$, let $\mathcal{L}_\mathbf{k}$ denote the leakage oracle in the definition of mTCCRL security (Section 2.2), namely, $\mathcal{L}_\mathbf{k}(\alpha; P) := P(k_\alpha)$. For a list $\mathsf{L} := \Big(\big((\alpha_1; P_1), \ell_1 \big), \ldots, \big((\alpha_q; P_q), \ell_q \big) \Big)$ such that $\alpha_j \in \{1, \ldots, u\}$, $P_j \in \mathcal{P}$, $\ell_j \in \{0,1\}$, and $\Pr_{k_1 \leftarrow \mathcal{R}_1, \ldots, k_u \leftarrow \mathcal{R}_u}[\mathcal{L}_\mathbf{k}(\alpha_i; P_i) = \ell_i \text{ for all } i = 1, \ldots, q] > 0$, we define[6] $\mathcal{R}_\alpha|_\mathsf{L}$ as the distribution $\mathcal{R}_\alpha$ conditioned on all leakage observations concerning α in the transcript L. That is,

$$\mathcal{R}_\alpha|_\mathsf{L}(k) := \Pr_{k_* \leftarrow \mathcal{R}_\alpha}\left[k = k_* \mid P_j(k_*) = \ell_j \text{ for } 1 \leq j \leq q \text{ s.t. } \alpha_j = \alpha\right]. \tag{2}$$

Consider the following experiment $\mathsf{Expt}^{\mathrm{cmel}}_{\mathbf{R},\mathcal{P},\rho,\mathcal{A}}$ that involves an adversary $\mathcal{A}$. (Intuitively, $\mathcal{A}$ "wins" if, for some α, $\mathcal{A}$ succeeds in reducing the conditional min-entropy of $\mathcal{R}_\alpha$ to be below the threshold ρ.)

1. The keys $\mathbf{k} = (k_1, \ldots, k_u)$ are sampled independently, with $k_\alpha \leftarrow \mathcal{R}_\alpha$ for each $\alpha = 1, \ldots, u$.
2. $\mathcal{A}$ performs some computation, making queries to the $\mathcal{L}_\mathbf{k}$ oracle. As in the definition of mTCCRL security, if $\mathcal{L}_\mathbf{k}$ returns 0 in response to a query $(\alpha; P)$, then $\mathcal{A}$ is prohibited from querying $(\alpha; *)$ thereafter.

[5] No ideal-cipher or random-permutation oracle appears in the definition of the cmel advantage itself. We nevertheless restrict to oracle-free predicates when applying cmel bounds in the QICM/QRPM analysis.

[6] We will use this notation $\mathcal{R}_\alpha|_\mathsf{L}$ not only here but also in later sections.

3. After $\mathcal{A}$ stops, define $\mathsf{L} := \Big(\big((\alpha_1; P_1), \ell_1\big), \ldots, \big((\alpha_q; P_q), \ell_q\big)\Big)$ as the transcript of all the queries made by $\mathcal{A}$ and the responses (i.e., $\mathcal{L}_{\mathbf{k}}(\alpha_i; P_i) = \ell_i$ for $i = 1, \ldots, q$). Define ρ_α as the min-entropy of $\mathcal{R}_\alpha|_{\mathsf{L}}$, i.e., $\rho_\alpha := -\log_2 \max_{k \in \{0,1\}^n} \mathcal{R}_\alpha|_{\mathsf{L}}(k)$. Finally, the value $\rho_* := \min_{\alpha \in \{1,\ldots,u\}} \rho_\alpha$ is returned as the output of the experiment.

We define the conditional min-entropy with leakage (cmel) advantage of $\mathcal{A}$ as $\mathsf{Adv}^{\mathrm{cmel}}_{\mathbf{R},\mathcal{P},\rho}(\mathcal{A}) := \Pr\left[\rho_* \leftarrow \mathsf{Expt}^{\mathrm{cmel}}_{\mathbf{R},\mathcal{P},\rho,\mathcal{A}} : \rho_* < \rho\right]$. We further define $\mathsf{Adv}^{\mathrm{cmel}}_{\mathbf{R},\mathcal{P},\rho}(M, q_{L,\max}, Q_{L,\max}) := \max_{\mathcal{A}} \mathsf{Adv}^{\mathrm{cmel}}_{\mathbf{R},\mathcal{P},\rho}(\mathcal{A})$, where the maximum is taken over all possible $\mathcal{A}$ making at most $q_{L,\max}$ queries to $\mathcal{L}_{\mathbf{k}}$ in total, involving at most M distinct values of α, and at most $Q_{L,\max}$ queries for any fixed α.

Bounding the Advantage for a Practically Relevant Structured Setting. The cmel advantage depends on the key distributions and on the class of allowed leakage predicates. We first consider a setting that naturally arises in practical applications, where keys are uniformly distributed over affine subspaces and leakage predicates are given by affine linear equations. The following proposition gives an upper bound in this setting (see Section D for a proof).

Proposition 3.1. *Let d be a positive integer and $\mathcal{P}$ be a set of (oracle-free) linear predicates. For each $\alpha = 1, \ldots, u$, suppose that $\mathcal{R}_\alpha$ is the uniform distribution over an affine subspace of dimension at least d. Then, for any non-negative integers w, M, $q_{L,\max}$, and $Q_{L,\max}$ such that $w < d$, we have $\mathsf{Adv}^{\mathrm{cmel}}_{\mathbf{R},\mathcal{P},(d-w)}(M, q_{L,\max}, Q_{L,\max}) \le M/2^w$.*

The assumptions in Proposition 3.1 are motivated by the affine selective-abort leakage considered by Roy [Roy22]. In Roy's work, selective-abort leakage is modeled by testing whether the secret shift belongs to an affine subspace. Membership in an affine subspace is equivalent to satisfying a system of affine linear equations. Thus, by treating these equations as individual linear predicates, the linear predicate class can be viewed as an information-theoretic abstraction of such leakage. Moreover, conditioning a uniform distribution over an affine subspace on the outcomes of linear predicates preserves this form. Therefore, Proposition 3.1 provides a conservative and easy-to-use bound on the entropy loss caused by such leakage[7].

Bounding the Advantage for General Settings. The previous proposition exploits structure in both the key distributions and the leakage predicates. In general, the following proposition holds. The bound is worse than Proposition 3.1 by a factor of $Q_{L,\max}$, but is applicable to any distributions $\mathcal{R}_\alpha$ with sufficient min-entropy and any set of (oracle-free) predicates. See Section E for a proof[8].

Proposition 3.2. *Let d be a positive integer and $\mathcal{P}$ be any set of (oracle-free) predicates. For each $\alpha = 1, \ldots, u$, suppose that $\mathcal{R}_\alpha$ is a distribution with min-entropy at least d. Then, for any non-negative integers w, M, $q_{L,\max}$, and $Q_{L,\max}$ such that $w < d$, we have*

$$\mathsf{Adv}^{\mathrm{cmel}}_{\mathbf{R},\mathcal{P},(d-w)}(M, q_{L,\max}, Q_{L,\max}) \le M Q_{L,\max}/2^w. \tag{3}$$

Arbitrary-predicate leakage of this form also appears in OT-extension frameworks. For example, the $\mathcal{F}_{\Delta\text{-ROT}}$ functionalities[9] appearing in Guo et al. [GKWY20] and Chen and Tessaro [CT21] allow a corrupted receiver to specify a predicate P on the global correlation Δ and learn $P(\Delta)$, with the functionality aborting when $P(\Delta) = 0$. Guo et al. [GWYY25] further formulate predicate-based key leakage in the multi-user TCCR setting. Thus, the general predicate class considered in Proposition 3.2 is consistent with the leakage modeling used in these works; in particular, the proposition gives a concrete cmel bound for the leakage considered in [GWYY25] under its stated min-entropy assumption on the keys.

The bound of Proposition 3.2 is asymptotically tight in a natural parameter regime. In particular, without additional assumptions on the key distributions or the leakage predicates, the dependence on M, $Q_{L,\max}$, and 2^{-w} is unavoidable up to lower-order factors. We present details as Proposition F.1 in the appendix.

[7] This should not be interpreted as a direct implication from mTCCRL to Roy's TCR-like notion; rather, Roy's affine selective-abort model motivates the structured setting considered in this information-theoretic cmel calculation.

[8] ChatGPT was used to generate an initial draft of the proof of Propositions 3.2 and F.1 (Proposition F.1 shows the tightness of the bound in Proposition 3.2). The authors subsequently checked each step, corrected and completed the argument, and prepared the final proof. The authors take full responsibility for its correctness.

[9] Here, our reference to [GKWY20] concerns only the leakage interface of $\mathcal{F}_{\Delta\text{-ROT}}$. Subsequent works pointed out issues with the TCR formulation used in the malicious-security analysis of [GKWY20]; we discuss this point in Section 4.2. In addition, note that the observation here concerns only the modeling of leakage and does not by itself imply the corresponding protocol-security results. The application to the mTCCRL-based protocol analysis of [GWYY25] is established separately in Section 4.2.

3.2 Main Theorem

Let $H^E : \{0,1\}^n \times \{0,1\}^m \to \{0,1\}^n$ be a function defined by

$$H^E(x,t) := E(t, \varphi(x)) \oplus \sigma(x), \tag{4}$$

where φ and σ are permutations over $\{0,1\}^n$, and σ linear. This H captures MMO, $\widehat{\mathsf{MMO}}$, and EncFF as special cases. If both φ and σ are the identity permutation, then $H^E = \mathsf{MMO}^E$. If σ is an orthomorphism, then setting $\varphi(x) := x$ yields $H^E = \mathsf{EncFF}^E$, whereas setting $\varphi(x) := \sigma(x)$ yields $H^E = \widehat{\mathsf{MMO}}^E$.

We are now ready to state our main theorem, whose proof is deferred to Section 5.

Theorem 3.1. *Let E be an n-bit ideal cipher with m-bit keys, $\mathbf{R} = (\mathcal{R}_1, \ldots, \mathcal{R}_u)$ be any tuple of distributions over $\{0,1\}^n$, $\mathcal{P}$ be any set of oracle-free predicates over $\{0,1\}^n$, and H^E be the function defined in Eq. (4). Suppose that $q_C \le 2^{n-1}$, $0 < \rho \le n$, and $\mathsf{Adv}^{\mathrm{cmel}}_{\mathbf{R},\mathcal{P},\rho}(M, q_{L,\max}, Q_{L,\max}) < 1$. Let $\mathcal{A}$ be a $(q_E, q_C, u, M, q_{L,\max}, Q_{L,\max}, \mu)$-bounded adversary. Then, in the QICM, $\mathsf{Adv}^{\mathrm{mTCRL}}_{H^E,\mathbf{R},\mathcal{P}}(\mathcal{A}) \le \varepsilon$ (resp., $\mathsf{Adv}^{\mathrm{mTCCRL}}_{H^E,\mathbf{R},\mathcal{P}}(\mathcal{A}) \le \varepsilon$ if σ is an orthomorphism), where*

$$\begin{aligned}\varepsilon := {} & \frac{4\sqrt{\frac{(q_C)^2 q_E}{2^\rho}} + 2\sqrt{\frac{2\mu(q_E)^2}{2^\rho}} + q_C \cdot \mathsf{Adv}^{\mathrm{cmel}}_{\mathbf{R},\mathcal{P},\rho}(M, q_{L,\max}, Q_{L,\max})}{1 - \mathsf{Adv}^{\mathrm{cmel}}_{\mathbf{R},\mathcal{P},\rho}(M, q_{L,\max}, Q_{L,\max})} \\ & + \frac{13 q_C \mu}{2^\rho} + q_C(q_{L,\max} + 4) \cdot \mathsf{Adv}^{\mathrm{cmel}}_{\mathbf{R},\mathcal{P},\rho}(M, q_{L,\max}, Q_{L,\max}).\end{aligned} \tag{5}$$

A More Concrete Bound with Leakage. Suppose that, for each α, $\mathcal{R}_\alpha$ is uniform over an affine subspace of dimension at least d, and that $\mathcal{P}$ is a set of (oracle-free) linear predicates. Then, applying Proposition 3.1 with any integer w satisfying $M < 2^w < 2^d$, Eq. (5) becomes

$$\varepsilon := \frac{4\sqrt{\frac{(q_C)^2 q_E}{2^{d-w}}} + 2\sqrt{\frac{2\mu(q_E)^2}{2^{d-w}}} + (q_C M/2^w)}{1 - (M/2^w)} + \frac{13 q_C \mu}{2^{d-w}} + \frac{q_C(q_{L,\max} + 4)M}{2^w}$$

By applying Proposition 3.2 instead of Proposition 3.1, we also obtain a similar bound for arbitrary key distributions and predicates.

On the Security Bound without Leakage. Without leakage, the cmel advantage is zero by definition, yielding the following corollary for mTCR and mTCCR security.

Corollary 3.1. *Let E be an n-bit ideal cipher with m-bit keys, $\mathbf{R} = (\mathcal{R}_1, \ldots, \mathcal{R}_u)$ be a tuple of distributions over $\{0,1\}^n$, and H^E be the function defined in Eq. (4). Suppose that $q_C \le 2^{n-1}$ and the min-entropy of $\mathcal{R}_\alpha$ is at least ρ for all $\alpha = 1, \ldots, u$. Let $\mathcal{A}$ be a (q_E, q_C, u, μ)-bounded adversary. Then, in the QICM, $\mathsf{Adv}^{\mathrm{mTCR}}_{H^E,\mathbf{R}}(\mathcal{A}) \le \varepsilon$ (resp., $\mathsf{Adv}^{\mathrm{mTCCR}}_{H^E,\mathbf{R}}(\mathcal{A}) \le \varepsilon$ if σ is an orthomorphism), where $\varepsilon := 2\sqrt{2\mu q_E^2/2^\rho} + 13\mu q_C/2^\rho + 4\sqrt{q_E q_C^2/2^\rho}$.*

On Instantiations with Concrete Block Ciphers. Suppose $\rho = n-1$, as in the half-gates garbling schemes of [GKWY20, GKW+20]. If E is instantiated with AES-128 [Nat01] ($m = n = 128$), we obtain $\varepsilon \approx \sqrt{\mu q_E^2/2^{124}} + \mu q_C/2^{123.3} + \sqrt{q_E q_C^2/2^{123}}$, which shows that our bound guarantees security when $q_E, q_C \ll 2^{41}$.

Furthermore, if E is instantiated with Rijndael-256 [DR99, DR20] (which is a 256-bit block variant of AES, i.e., $m = n = 256$, and for which NIST has proposed for standardization [Nat24b, Nat26]), we obtain $\varepsilon \approx \sqrt{\mu q_E^2/2^{252}} + \mu q_C/2^{251.3} + \sqrt{q_E q_C^2/2^{251}}$. In this case, our bound guarantees security when $q_E, q_C \ll 2^{83.7}$.

From a practical perspective, our bound provides only a limited provable post-quantum security margin for AES-128 in these applications. When a larger provable margin is desired, Rijndael-256 provides a natural alternative: besides offering a substantially larger security margin under our bound, it can still efficiently exploit AES-NI instructions on modern x86-64 processors [DG22].

On Tightness of the Bound. Similar to the post-quantum security bound for the Even-Mansour cipher in [ABKM22], our bound is tight with respect to the query complexity, but not with respect to the distinguishing advantage. To see this, suppose that there is no leakage and that the key distributions are identical and uniform over an affine subspace. In realistic attack scenarios, an adversary can implement the quantum oracle of $E^{\pm}$ by itself, while obtaining a relatively small number of samples from the classical construction oracle; hence one typically has

$q_E \gg q_C\ (\geq \mu)$. In this regime, our bound reduces to $O\left(\sqrt{\mu q_E^2/2^\rho}\right)$. On the other hand, due to the similarity between our function H^E and the Even-Mansour cipher, BHT-based meet-in-the-middle attacks [BHT98, KM12, HS18] and the offline Simon attack [BHN+19] yield advantage $\Theta\left(\mu q_E^2/2^\rho\right)$. Both bounds reach constant order when $q_E \approx \sqrt{2^\rho/\mu}$, whereas they exhibit a quadratic gap for $q_E \ll \sqrt{2^\rho/\mu}$. Whether this quadratic gap can be closed remains open, as also remarked in [ABKM22, Footnote 1].

4 Applications

This section presents applications of our results to several representative protocols. By combining Theorem 3.1 and Corollary 3.1 with existing security analyses, we verify that when the CR-type hashes in these protocols are instantiated with the MMO-type constructions covered by the theorem, the resulting protocol remains secure in the QICM/QRPM. We also discuss the limitations on the applicability of our results. (See Table 2 for a brief summary.) When referring to theorems, propositions, etc. from prior work, we use the numbering in the version of each work cited in the references, unless otherwise noted.

4.1 Replacing CR-Type Hash Invocations in the QICM/QRPM

We first explain when our QICM/QRPM bounds can be plugged into existing classical protocol analyses based on CR-type assumptions. This is not a general composition theorem; it only identifies a common proof pattern in which the CR-type hash-replacement step remains justified when the hash is instantiated with the MMO-type constructions covered by our results.

We begin by explaining the basic correlation-robustness (CR) case in the QICM. Let H^E be a hash function constructed from an ideal cipher E. CR states that the real oracle $O^{\mathrm{CR}}_{H,k}(x) := H^E(x \oplus k)$ is indistinguishable from an independent random function f. Hence, in a protocol analysis, the outputs $H^E(x_1 \oplus k), \ldots, H^E(x_q \oplus k)$ can be replaced by $f(x_1), \ldots, f(x_q)$, where the inputs $x_1, \ldots, x_q$ may be chosen adaptively. After the replacement, each fresh input is answered with an independent uniform string, while repeated inputs are answered consistently.

Theorem 3.1 and Corollary 3.1 justify the same replacement in protocol security proofs when the hash function is instantiated with the covered MMO-type constructions, even if the adversary has quantum superposition access to the underlying ideal cipher E and E^{-1}: The random function f used for the replacement is independent of E. Consequently, if the remaining part of the protocol security proof does not rely on any further structure of the ideal cipher, it can be reused (with the CR-type hash assumption replaced by our QICM bound).

The same reasoning applies to stronger CR-type notions, such as CCR, TCR, TCCR, and their leakage or multi-user variants. It also applies in the QRPM when the hash construction is instantiated from a random permutation rather than an ideal cipher.

Limitations. There are two important limitations. First, if the protocol uses other cryptographic primitives, setup assumptions, or ideal functionalities, those components must themselves be post-quantum secure or must be available as ideal functionalities in the relevant hybrid model. Our result only addresses the hash values produced by the MMO-type constructions. Second, the original proof must use the ideal cipher or random permutation only through the CR-type hash invocations being replaced. Proofs that rely on additional RPM/ICM-specific arguments (e.g., those that extract information from a classical primitive-query transcript) are not automatically covered. Such steps require separate quantum analyses and are outside the scope of our results.

4.2 OT Extension

Oblivious transfer (OT), introduced by Rabin [Rab81], is one of the fundamental primitives for secure computation. OT extension, initiated by Beaver [Bea96] and made highly practical by Ishai et al. [IKNP03], expands a small number of base OTs into many OT instances using mostly symmetric-key primitives.

Applications Covered by Our Results.

Semi-Honest Security. Ishai et al. [IKNP03] showed that the semi-honest IKNP OT-extension protocol can be based on correlation robustness of the hash function H. Asharov et al. [ALSZ13] proposed optimized OT-extension protocols for correlated and random OT. Guo et al. [GKWY20] revisited these protocols, and showed that CR suffices in the semi-honest setting.

More concretely, Theorems 1, 2, and 7 of Guo et al. [GKWY20] analyze standard OT, correlated OT, and random OT, respectively, in the $\mathcal{F}_{\Delta\text{-ROT}}$-hybrid model. (Those theorems also discuss malicious security, but here we focus only on the results about semi-honest security.) They decompose OT extensions into two phases, and the $\mathcal{F}_{\Delta\text{-ROT}}$-hybrid model assumes that the first phase is ideally secure; Guo et al.'s discussion (and ours also) concerns the second, hash-dependent phase. The simulators in these proofs are straight-line and do not program the random permutation or ideal cipher, nor do they extract information from a primitive query transcript. The computational step is essentially the replacement of the real hash values by random strings.

Thus, the hash-replacement principle described in Section 4.1 is applicable. When the hash function is instantiated by the MMO-type constructions considered in this work, the CR hash-replacement steps in Theorems 1, 2, and 7 of [GKWY20] (in the semi-honest setting) are justified against quantum adversaries in the QICM/QRPM. If $\mathcal{F}_{\Delta\text{-ROT}}$ is further realized by a post-quantum secure protocol and the corresponding quantum composition theorem applies, then the corresponding instantiated OT-extension protocols obtain post-quantum security in the corresponding model.

Malicious Security. Guo et al. [GKWY20] also analyzed malicious security of OT extensions relying on TCR security, but subsequent works pointed out issues with their TCR formulation [CT21, Roy22]. To avoid the issue, Chen and Tessaro [CT21] introduced a notion which is named TCR* and reduced the security of their protocol for standard OT to the TCR* security of the underlying hash in the $\mathcal{F}_{\Delta\text{-ROT}}$-hybrid model [CT21, Theorem 1]. Guo et al. [GWYY25] subsequently gave a dedicated multi-user analysis of a random-IV variant of the random-OT-to-standard-OT transformation of Guo et al. [GKWY20], and their Theorem 2 directly reduces its malicious multi-user security to the mTCCRL security of the underlying hash function (with an additional statistical term).

The proof of [GWYY25, Theorem 2] again falls within the principle of Section 4.1; in particular, the mTCCRL term in [GWYY25, Theorem 2] can be replaced by our mTCCRL bound for $\widehat{\mathsf{MMO}}$ and EncFF in the QICM, while the remaining information-theoretic terms are unchanged. (Our result does not directly instantiate the TCR*-based result of [CT21, Theorem 1], because mTCCRL security does not imply TCR*; nevertheless, the standard-OT application is already covered by the mTCCRL-based reduction of [GWYY25, Theorem 2] discussed above, without invoking TCR*.)

A separate analysis is given by Roy [Roy22], who studies SoftSpokenOT through a subspace-VOLE framework. Theorem 5.1 of Roy [Roy22] reduces the malicious security of the hash-based subspace-VOLE-to-OT extension step to a TCR-like security notion for the underlying hash function.

Roy's TCR-like notion (Definition 2.5 of [Roy22]) is defined with a linear code $\mathcal{C}$ and involves hash inputs of the form $H(xG_{\mathcal{C}} \odot \Delta + y, \tau)$. Here, Δ is a secret value, and $x(\neq 0), y, \tau$ are the values that the adversary can choose, and $\odot$ is the component-wise multiplication. This is not, for arbitrary $\mathcal{C}$, covered by our TCRL notion. However, in the case of one-dimensional binary repetition-code[10], Roy's TCR oracle takes the form of $H(\Delta \oplus y, \tau)$ with selective-abort leakage on Δ, which is covered by our TCRL formulation. This binary repetition-code case is practically meaningful: it corresponds to the ordinary binary 1-out-of-2 OT-extension setting considered and implemented by Roy. For this repetition-code case, the principle of Section 4.1 is applicable, and the TCR-like hash assumption used in [Roy22, Theorem 5.1] can be instantiated by our MMO-type constructions in the QICM.

Results outside Our Scope. There are several OT extension protocols that fall outside the scope of our results, either because they rely on stronger security notions for hash functions or because their security proofs are carried out directly in the RPM or ICM.

- As mentioned above, our principle applies to Roy's result only for the repetition-code case, because Roy's TCR-like notion is not covered by our mTCRL for general codes. The full subspace-VOLE framework for arbitrary linear codes, and the resulting 1-out-of-N OT applications, would require a more general linear-code version of TCRL and are not directly covered by our consideration.
- Guo et al. [GYW+23] introduced the half-tree technique and applied it to obtain a single-point correlated OT extension protocol in the RPM. The security of their construction is shown in the UC framework with static corruption and semi-honest adversaries. Intuitively, their simulator must detect construction or primitive

[10] That is, $p = 2$, $k_{\mathcal{C}} = 1$, and $G_{\mathcal{C}} = [1, \ldots, 1]$ in Roy's notations.

queries related to a secret key k held by the environment. While this is achieved in the RPM using the H-coefficient technique, it is currently unclear how to extend this analysis to the QRPM and QICM.

– Motivated by the simulation requirements in UC security proofs, Cui et al. [CGG+25] defined a simulation-based notion of security for mTCCR. They demonstrated applications to UC-secure correlated OT extension in Ferret [YWL+20] and to authenticated AND triple generation [YWZ20]. Our notion does not imply such simulation-based security, so their applications are not covered, either.

4.3 Garbled Circuits

Garbled circuits [Yao86, BHR12] are a fundamental building block for secure two-party computation, and correlation robust hashing plays a central role in the design of efficient garbled circuits.

Applications Covered by Our Results. Among practical garbling optimizations, the half-gates construction of Zahur, Rosulek, and Evans [ZRE15] is particularly important because it is compatible with Free-XOR [KS08, CKKZ12] and reduces each AND gate to two ciphertexts. They proved that the construction is secure if the underlying hash function satisfies *circular correlation robustness for naturally derived keys* (CCRND), which is directly implied by TCCR (as observed by Guo et al. [GKWY20]). Finally, Guo et al. [GKW+20] showed the multi-instance security of the half-gates construction and its variants, assuming the underlying hash function is mTCCR. Our principle of Section 4.1 directly applies to the results of [ZRE15, Theorems 1-2] and [GKW+20, Theorems 1, 3-5] when the hash function is instantiated with $\widehat{\mathsf{MMO}}$ or EncFF. In particular, the same reduction bounds as in Theorems 1 and 3-5 of [GKW+20] apply to the corresponding QICM instantiations, with the mTCCR-advantage term replaced by our bound from Corollary 3.1[11].

Results outside Our Scope. The three-halves scheme by Rosulek and Roy [RR21] relies on *randomized tweakable circular correlation robustness* (RTCCR). It is stronger than our TCCR, and their security proof is outside our scope.

4.4 (Correlated) GGM Tree

GGM (or cGGM, respectively) trees [GGM84] instantiated with (T)CR (or (T)CCR, respectively) hash functions play an important role in building puncturable pseudorandom function (PPRF) constructions [BW13] for MPC-in-the-Head protocols and signature schemes [IKOS07].

Applications Covered by Our Results. Guo et al. [GKWY20, Section VI., B.] observed that if H is correlation robust (CR), then $G(x) := H(x \oplus 1) \| H(x \oplus 2)$ forms a doubling PRG, allowing the GGM tree to be instantiated using a fixed-key permutation. Furthermore, Guo et al. (Theorem 8 of [GYW+22], which is the full version of [GYW+23]) showed that their correlated GGM (cGGM) trees imply secure PPRFs, assuming that the underlying hash is CCR. The hash-dependent pseudorandomness parts of these results are directly covered by the principle of Section 4.1; by instantiating the underlying hash function with MMO (resp., $\widehat{\mathsf{MMO}}$ or EncFF), we obtain the corresponding QRPM/QICM analogue of the PRG statement of [GKWY20, Section VI., B.] (resp., the PPRF statement of [GYW+22, Theorem 8]).

Bui et al. [BCC+24] analyzed multi-instance PPRFs based on the doubling PRG $G_{t_0,t_1}(x) = (E(t_0, x) \oplus x) \| (E(t_1, x) \oplus x)$, where E denotes AES in the ideal cipher model and t_0, t_1 are random public salts (Theorems 7 and 9 of [BCC+24]). Their analysis is not formulated in terms of correlation robustness. However, since $G_{t_0,t_1}(x) = \mathsf{MMO}^E(x, t_0) \| \mathsf{MMO}^E(x, t_1)$ holds, our result on the TCR security of MMO^E (Corollary 3.1) implies the multi-instance PRG-security of G in the QICM, which can replace the corresponding statement of Theorem 9 of [BCC+24]. Plugging this into their generic PRG-to-PPRF reduction (Theorem 7 of [BCC+24]) yields the corresponding post-quantum multi-instance PPRF security statement in the QICM.

[11] In [GKW+20], the multi-instance security notion of TCCR is referred to as "miTCCR". Our mTCCR is slightly stronger than their miTCCR in that we allow the key distributions $\mathcal{R}_1, \ldots, \mathcal{R}_u$ to be distinct.

Limitations for Full Signature Security. Several other MPC-in-the-Head signature schemes also employ GGM or correlated-GGM trees [CLY+24, HJ24, KLS25, BCD25, WKK+25]. In particular, the schemes of [CLY+24, BCD25, WKK+25] explicitly use CCR-type hash functions inside their commitment mechanisms. For such schemes, our results can justify the CCR-dependent pseudorandomness or hiding components in the QICM/QRPM, whenever the corresponding hash is instantiated by one of the MMO-type constructions covered by Corollary 3.1. This includes direct applications of the principle of Section 4.1, as well as indirect applications in which an ideal-cipher-based GGM or PPRF statement is replaced by our QICM bound, as in the discussion of Bui et al. [BCC+24] above.

However, these observations do not yield full post-quantum security of the resulting signature schemes. Their proofs of the EUF-NMA security of signatures also rely on additional properties, such as the (extractable) binding of vector (semi-)commitments [BCD25, WKK+25, KLS25]. The existing proofs of such properties often use classical random-oracle, random-permutation, or ideal-cipher transcripts, which do not automatically lift to the QICM/QRPM from our results alone. Therefore, the full post-quantum unforgeability of the signature schemes is outside the scope of our results, unless these remaining commitment and other components are established separately in the appropriate post-quantum model.

4.5 Distributed Point and Comparison Functions (DPF/DCF)

A function $f_{\alpha,\beta}(x)$ is called a *point function* if $f_{\alpha,\beta}(x) = \beta$ when $x = \alpha$ and 0 otherwise. Similarly, a *comparison function* is defined as $f^{<}_{\alpha,\beta}(x) = \beta$ if $x < \alpha$ and 0 otherwise. Gilboa and Ishai [GI14] introduced *distributed point functions* (DPF), which split a secret point function into m additive shares $f_1, \ldots, f_m$ such that $\sum_i f_i(x) = f_{\alpha,\beta}(x)$. Boyle, Gilboa, and Ishai [BGI15] generalized DPF to *function secret sharing* (FSS), providing a unified framework for sharing a wide class of functions. Within this framework, they studied comparison and interval functions, and subsequent work [BCG+21] studied more efficient constructions of *distributed comparison functions* (DCF). These primitives have found numerous applications, including two-server PIR and private search, distributed ORAM, and preprocessing techniques for MPC [GI14, BGI15, BGI16, Ds17, BGI19, BCG+21].

Applications Covered by Our Results. Our principle of Section 4.1 applies to the CCR-hash-replacement steps in the half-tree framework of Guo et al. [GYW+23]. Their Theorems 3 and 4 construct DPF and DCF schemes from a CCR hash function, and Theorems 5 and 6 prove the security of the corresponding DPF/DCF protocols in the hybrid model. The hash-dependent parts of these proofs are again covered by our principle of Section 4.1. Hence, when the CCR hash is instantiated by $\widehat{\mathsf{MMO}}$ or EncFF, the CCR-dependent hash components in these DPF/DCF constructions are justified in the QICM/QRPM by Corollary 3.1.

5 Proof of Theorem 3.1

We prove the security of the hash function H^E of Eq. (4) by reducing it to multi-user security of an Even-Mansour-like tweakable block cipher construction.

5.1 Reduction to Tweakable Block-Cipher Security

Multi-User Tweakable Pseudorandom Permutation with Leakage in the QICM. Let $\mathsf{TBC}^E_k(x, t, b)$ be an n-bit tweakable block cipher with tweak space $\{0,1\}^m \times \{0,1\}$ (TBC^E_k depends on an ideal cipher E). That is, in addition to a key k and a plaintext x, TBC^E_k takes a pair $(t, b) \in \{0,1\}^m \times \{0,1\}$ as part of input, and $\mathsf{TBC}^E_k(\cdot, t, b)$ is a permutation for each (t, b).

Let $\mathbf{R} = (\mathcal{R}_1, \ldots, \mathcal{R}_u)$ be a tuple of distributions over $\{0,1\}^n$, and $\mathcal{P}$ be a set of oracle-free predicates defined over $\{0,1\}^n$. Suppose that u keys $\mathbf{k} = (k_1, \ldots, k_u)$ are sampled independently as $k_\alpha \leftarrow \mathcal{R}_\alpha$ for $\alpha = 1, \ldots, u$. For an adversary $\mathcal{A}$, we define its multi-user special tweakable pseudorandom permutation with leakage (mTPRPL*) advantage $\mathsf{Adv}^{\mathrm{mTPRPL}^*}_{\overline{\mathsf{TBC}},\mathbf{R},\mathcal{P}}(\mathcal{A})$ in the same way as the mTCCRL advantage, except that

- the oracle $\mathcal{O}^{\mathrm{mTCCRL}}_{H,\mathbf{k}}$ is replaced with $\overline{\mathsf{TBC}}^E_{\mathbf{k}}$ defined by $\overline{\mathsf{TBC}}^E_{\mathbf{k}}(\alpha; x, t, b) := \mathsf{TBC}^E_{k_\alpha}(x, t, b)$,
- the random function f is replaced with a tweakable random permutation[12], denoted by $\widetilde{\mathrm{RP}}(\alpha; x, t, b)$, such that $\widetilde{\mathrm{RP}}(\alpha; \cdot, t, b)$ is an independent n-bit random permutation for each (α, t, b).

[12] This is essentially just an ideal cipher, but we use a different name to distinguish it from the ideal cipher underlying hash functions and tweakable block ciphers.

We also define a $(q_E, q_C, u, M, q_{L,\max}, Q_{L,\max}, \mu)$-bounded adversary for mTPRPL* security in the same way as for mTCCRL security. (The same query restriction as in mTCCRL also applies to mTPRPL*.)

In addition, we define multi-user tweakable pseudo-random permutation with leakage (mTPRPL) advantage by fixing $b = 0$ in the definition of mTPRPL* advantage. A $(q_E, q_C, u, M, q_{L,\max}, Q_{L,\max}, \mu)$-bounded adversary is defined analogously. (The relationship between mTPRPL* and mTPRPL corresponds to that between mTCCRL and mTCRL.)

Reduction to the Security of TBC. The following proposition reduces the security of H^E defined in Eq. (4) to that of an Even-Mansour-like tweakable block cipher construction.

Proposition 5.1. *Let H^E be the function defined in Eq. (4). Define a tweakable block cipher TBC_k^E by*

$$\mathsf{TBC}_k^E(x,t,b) := \begin{cases} E(t, \varphi(x \oplus k)) \oplus \sigma(k) & \text{if } b = 0, \\ E(t, \varphi(x \oplus k)) \oplus \sigma(k) \oplus k & \text{if } b = 1, \end{cases} \tag{6}$$

where φ and σ are permutations over $\{0,1\}^n$, and σ is linear. If $\mathsf{Adv}^{\mathrm{mTPRPL}^*}_{\mathsf{TBC},\mathbf{R},\mathcal{P}}(\mathcal{B}) \le \varepsilon$ (resp., $\mathsf{Adv}^{\mathrm{mTPRPL}}_{\mathsf{TBC},\mathbf{R},\mathcal{P}}(\mathcal{B}) \le \varepsilon$) in the QICM for any $(q_E, q_C, u, M, q_{L,\max}, Q_{L,\max}, \mu)$-bounded adversary $\mathcal{B}$, then $\mathsf{Adv}^{\mathrm{mTCCRL}}_{H,\mathbf{R},\mathcal{P}}(\mathcal{A}) \le \varepsilon + (q_C\mu/2^n)$ (resp., $\mathsf{Adv}^{\mathrm{mTCRL}}_{H,\mathbf{R},\mathcal{P}}(\mathcal{A}) \le \varepsilon + (q_C\mu/2^n)$) in the QICM for any $(q_E, q_C, u, M, q_{L,\max}, Q_{L,\max}, \mu)$-bounded adversary $\mathcal{A}$.

The proof is straightforward; see Section B in the appendix for details. Note that the query restriction in mTPRPL* is essential for the TBC of Eq. (6): without it, the construction is trivially broken since $\mathsf{TBC}_k^E(x,t,0) \oplus \mathsf{TBC}_k^E(x,t,1) = k$.

Proving the Security of TBC. The main technical step in the proof of Theorem 3.1 is the following proposition. Together with Proposition 5.1, it immediately implies Theorem 3.1.

Proposition 5.2. *Suppose that $q_C \le 2^{n-1}$, $0 < \rho \le n$, and $\mathsf{Adv}^{\mathrm{cmel}}_{\mathbf{R},\mathcal{P},\rho}(M, q_{L,\max}, Q_{L,\max}) < 1$. Then, in the QICM, for the tweakable block cipher defined in Eq. (6) and any $(q_E, q_C, u, M, q_{L,\max}, Q_{L,\max}, \mu)$-bounded adversary $\mathcal{A}$, we have $\mathsf{Adv}^{\mathrm{mTPRPL}}_{\mathsf{TBC},\mathbf{R},\mathcal{P}}(\mathcal{A}) \le \varepsilon$ (resp., $\mathsf{Adv}^{\mathrm{mTPRPL}^*}_{\mathsf{TBC},\mathbf{R},\mathcal{P}}(\mathcal{A}) \le \varepsilon$ if σ is an orthomorphism), where*

$$\begin{aligned} \varepsilon := {} & \frac{4\sqrt{\frac{(q_C)^2 q_E}{2^\rho}} + 2\sqrt{\frac{2\mu(q_E)^2}{2^\rho}} + q_C \cdot \mathsf{Adv}^{\mathrm{cmel}}_{\mathbf{R},\mathcal{P},\rho}(M, q_{L,\max}, Q_{L,\max})}{1 - \mathsf{Adv}^{\mathrm{cmel}}_{\mathbf{R},\mathcal{P},\rho}(M, q_{L,\max}, Q_{L,\max})} \\ & + \frac{12 q_C \mu}{2^\rho} + q_C(q_{L,\max} + 4) \cdot \mathsf{Adv}^{\mathrm{cmel}}_{\mathbf{R},\mathcal{P},\rho}(M, q_{L,\max}, Q_{L,\max}). \end{aligned} \tag{7}$$

5.2 Proof of Proposition 5.2

In what follows, we will prove the claim for mTPRPL* security in Proposition 5.2, assuming that σ is an orthomorphism. The proof for the claim on mTPRPL can be obtained by fixing the bit b to be 0; in this case, the assumption that σ is a linear permutation suffices.

Our proof is based on the reprogramming-and-resampling proof paradigm introduced by [ABKM22]. The core technical tools are the arbitrary reprogramming lemma and a resampling lemma [ABKM22]. For the latter, we will use the variant from [ABMS26], which is suited to ideal ciphers and low-entropy keys. The proof technique has been developed in subsequent works [ABK+24, Hos25, SI26a, BBC+25, ABMS26]; see also Section 1.2. Proposition 5.2 does not follow by directly invoking an existing result as a black box: Our setting simultaneously involves an ideal cipher, possibly low-entropy and non-identically distributed keys, multi-user security, the special query restrictions inherited from TCCR security, and adaptive leakage[13].

The main additional technical difficulty comes from adaptive leakage. Leaking queries change the conditional distributions of the secret keys during the security experiment, and may occur between quantum primitive queries. The resulting entropy loss depends on the underlying key distributions and on the allowed class of

[13] If we focus only on the single-user leakage-free CR security of $\widehat{\mathrm{MMO}}^\pi$ in the QRPM, it can be derived from the analysis of [Hos25] for TEM with low-entropy keys, together with Proposition 5.1: when $m = 0$ and $b = 0$, the corresponding construction oracle is essentially the encryption oracle of a low-entropy Even-Mansour cipher. However, this restricted result is insufficient for many of the application settings discussed in Section 4.

predicates. To separate the leakage analysis from the ideal-cipher analysis, we use the cmel advantage defined in Section 3.1. Its role in the proof is to control the probability that the entropy condition needed for reprogramming and resampling fails at the points where we apply these techniques. The cmel advantage alone does not handle the interaction between adaptive leakage and key-dependent oracle replacements, so we introduce additional sub-hybrids that move the key-resampling and oracle-replacement steps across the intervening leaking queries one at a time. Further details are given below, after introducing the necessary notation and hybrids.

Setup and Notation. We next introduce the notation used in the proof.

Let $\psi_1(k,x) := \varphi(x \oplus k)$ and $\psi_2(k,y,b) := y \oplus \sigma(k) \oplus b \cdot k$. Then, we have $\overline{\mathsf{TBC}}^E_{\mathbf{k}}(\alpha;x,t,b) = \psi_2(k_\alpha, E(t,\psi_1(k_\alpha,x)),b)$. We will use this representation throughout the proof.

Recall that $\mathcal{A}$ is given classical oracle access to $\overline{\mathsf{TBC}}^E_{\mathbf{k}}$ and $\mathcal{L}_{\mathbf{k}}$ in the real world, or to $\widetilde{\mathrm{RP}}$ and $\mathcal{L}_{\mathbf{k}}$ in the ideal world. We refer to $\overline{\mathsf{TBC}}^E_{\mathbf{k}}/\widetilde{\mathrm{RP}}$ as the *construction oracle* and to $\mathcal{L}_{\mathbf{k}}$ as the *leaking oracle*. Queries to these oracles are called *construction queries* and *leaking queries*, respectively. We refer to the quantum oracle corresponding to $|E^{\pm}\rangle$ as the *primitive oracle*, and queries to the primitive oracle as *primitive queries*.

On Classical Queries and Transcripts. By T_i we denote the transcript of the first i construction queries. That is, $\mathsf{T}_i := (((\alpha_1;x_1,t_1,b_1),y_1),\dots,((\alpha_i;x_i,t_i,b_i),y_i))$, where $(\alpha_j;x_j,t_j,b_j)$ and y_j are the jth query and response, respectively.

In addition, by L_i we denote the transcript of all leaking queries made before the adversary submits the ith construction query. That is, if $\mathcal{A}$ makes q such leaking queries, then $\mathsf{L}_i := \left(\left(\left(\alpha'_1;P_1\right),\ell_1\right),\dots,\left(\left(\alpha'_q;P_q\right),\ell_q\right)\right)$.

Without loss of generality, we assume that $\mathcal{A}$ makes no redundant classical queries. That is, once the adversary queries a tuple to $\overline{\mathsf{TBC}}^E_{\mathbf{k}}$ or $\widetilde{\mathrm{RP}}$ (resp., $\mathcal{L}_{\mathbf{k}}$), it never queries the same tuple to the same oracle interface again. Together with the restriction that $\mathcal{A}$ does not query both $(\alpha;x,t,0)$ and $(\alpha;x,t,1)$ to $\overline{\mathsf{TBC}}^E_{\mathbf{k}}$ or $\widetilde{\mathrm{RP}}$ for any (α,x,t), this implies that $(\alpha,x,t) \neq (\alpha',x',t')$ for any two distinct construction queries $(\alpha;x,t,b)$ and $(\alpha';x',t',b')$.

On Reprogramming of Ideal Ciphers. In what follows, for arbitrary $t_* \in \{0,1\}^m$ and $s_0,s_1 \in \{0,1\}^n$, let $\mathsf{swap}^{t_*}_{s_0,s_1} : \{0,1\}^m \times \{0,1\}^n \to \{0,1\}^m \times \{0,1\}^n$ denote the map defined by

$$\mathsf{swap}^{t_*}_{s_0,s_1}(t,x) := \begin{cases} (t,s_1) & \text{if } t = t_* \text{ and } x = s_0, \\ (t,s_0) & \text{if } t = t_* \text{ and } x = s_1, \\ (t,x) & \text{otherwise.} \end{cases} \tag{8}$$

Let T_i be a transcript of construction queries defined above, and $\mathbf{k} := (k_1,\dots,k_u)$ be a list of keys. Define[14] $\mathsf{REP}_{\mathsf{T}_i,E,\mathbf{k}} : \{0,1\}^m \times \{0,1\}^n \to \{0,1\}^m \times \{0,1\}^n$ and $E^{\mathsf{T}_i,\mathbf{k}} \in \mathsf{BC}(m,n)$ by

$$\begin{aligned} \mathsf{REP}_{\mathsf{T}_i,E,\mathbf{k}} &:= \mathsf{swap}^{t_1}_{s_0^1,s_1^1} \circ \cdots \circ \mathsf{swap}^{t_i}_{s_0^i,s_1^i}, \\ E^{\mathsf{T}_i,\mathbf{k}}(t,x) &:= \text{the second component of } \mathsf{REP}_{\mathsf{T}_i,E,\mathbf{k}}(t,E(t,x)), \end{aligned}$$

where $s_0^j := E(t_j,\psi_1(k_{\alpha_j},x_j))$ and $s_1^j := \psi_2(k_{\alpha_j},y_j,b_j)$ for $1 \le j \le i$.

Hybrids. We prove the proposition using the following hybrids Hyb^j_i for $j = 1,\dots,5$ and $i = 0,\dots,q_C-1$.

Hyb^1_i: 1. The keys $\mathbf{k} = (k_1,\dots,k_u)$ are sampled as $k_\alpha \leftarrow \mathcal{R}_\alpha$ for each α at the beginning. Just as in the ideal world, the classical construction oracle $\widetilde{\mathrm{RP}}$, the leaking oracle $\mathcal{L}_{\mathbf{k}}$, and the quantum oracle of E (and its inverse) are given to the adversary $\mathcal{A}$; the adversary performs some computation, making some classical and quantum queries.

2. When $\mathcal{A}$ submits the $(i+1)$st construction query, resample the keys $\mathbf{k} = (k_1,\dots,k_u)$ as $k_\alpha \leftarrow \mathcal{R}_\alpha|_{\mathsf{L}_{i+1}}$ for each α. Thereafter, the construction oracle (resp., the quantum primitive oracle $|E^{\pm}\rangle$) is replaced with $\overline{\mathsf{TBC}}^{E^{\mathsf{T}_i,\mathbf{k}}}_{\mathbf{k}}$ (resp., the quantum oracle of $E^{\mathsf{T}_i,\mathbf{k}}$ and its inverse).

[14] REP is an abbreviation of reprogramming. $\mathsf{REP}_{\mathsf{T}_i,E,\mathbf{k}}$ is an operator to reprogram E, hoping that it becomes consistent with the first i classical construction queries in the hybrids defined later.

Hyb_i^2: 1. The keys $\mathbf{k} = (k_1, \ldots, k_u)$ are sampled as $k_\alpha \leftarrow \mathcal{R}_\alpha$ for each α at the beginning. The classical construction oracle $\widetilde{\mathrm{RP}}$, the leaking oracle $\mathcal{L}_{\mathbf{k}}$, and the quantum oracle of E (and its inverse) are given to the adversary $\mathcal{A}$.
2. When $\mathcal{A}$ submits the $(i+1)$st construction query (denoted by $(\alpha_{i+1}; x_{i+1}, t_{i+1}, b_{i+1})$), resample the keys $\mathbf{k} = (k_1, \ldots, k_u)$ as $k_\alpha \leftarrow \mathcal{R}_\alpha|_{\mathsf{L}_{i+1}}$ for each α, and set $E' := E \circ \mathsf{swap}^{t_{i+1}}_{s_0^*, s_1^*}$, where $s_0^* := \psi_1(k_{\alpha_{i+1}}, x_{i+1})$ and s_1^* is a uniform random n-bit string. Thereafter, the construction oracle (resp., the quantum primitive oracle $|E^{\pm}\rangle$) is replaced with $\overline{\mathsf{TBC}}_{\mathbf{k}}^{(E')^{\mathsf{T}_i,\mathbf{k}}}$ (resp., the quantum oracle of $(E')^{\mathsf{T}_i,\mathbf{k}}$ and its inverse).

Hyb_i^3: This hybrid is defined by modifying Hyb_i^2 in such a way that the $(i+1)$th construction query is answered with $\widetilde{y} := \psi_2\left(k_{\alpha_{i+1}}, E^{\mathsf{T}_i,\mathbf{k}}(t_{i+1}, s_1^*), b_{i+1}\right)$.

Hyb_i^4: This hybrid is also defined by modifying Hyb_i^2 in such a way that the $(i+1)$th construction query is answered using $\widetilde{\mathrm{RP}}$, with the response denoted by y_{i+1}, and that s_1^* is set as $s_1^* := \left(E^{\mathsf{T}_i,\mathbf{k}}\right)^{-1}\left(t_{i+1}, \psi_2\left(k_{\alpha_{i+1}}, y_{i+1}, b_{i+1}\right)\right)$.

Hyb_i^5: 1. The keys $\mathbf{k} = (k_1, \ldots, k_u)$ are sampled as $k_\alpha \leftarrow \mathcal{R}_\alpha$ for each α at the beginning. The classical construction oracle $\widetilde{\mathrm{RP}}$, the leaking oracle $\mathcal{L}_{\mathbf{k}}$, and the quantum oracle of E (and its inverse) are given to the adversary $\mathcal{A}$.
2. When $\mathcal{A}$ receives the response to the $(i+1)$st construction query, resample the keys $\mathbf{k} = (k_1, \ldots, k_u)$ as $k_\alpha \leftarrow \mathcal{R}_\alpha|_{\mathsf{L}_{i+1}}$ for each α. Thereafter, the construction oracle (resp., the quantum primitive oracle $|E^{\pm}\rangle$) is replaced with $\overline{\mathsf{TBC}}_{\mathbf{k}}^{E^{\mathsf{T}_{i+1},\mathbf{k}}}$ (resp., the quantum oracle of $E^{\mathsf{T}_{i+1},\mathbf{k}}$ and its inverse).

We also define $\mathsf{Hyb}_{q_C}^1$ to be the ideal world. Note that the output distribution of $\mathcal{A}$ in Hyb_0^1 is equal to that in the real world. (See also Table 7 in the appendix for comparison of the hybrids.)

Sketch of the Proof. What we have to bound is the difference between the real and ideal worlds, i.e., Hyb_0^1 and $\mathsf{Hyb}_{q_C}^1$. We do this by bounding the difference between the successive hybrids $\mathsf{Hyb}_i^1, \mathsf{Hyb}_i^2, \ldots, \mathsf{Hyb}_i^5, \mathsf{Hyb}_{i+1}^1$ for $i = 0, \ldots, q_C - 1$, and then summing the resulting bounds.

Applying cmel before Key-Dependent Oracle Replacement. Before key resampling in each hybrid, construction and primitive queries are answered using $\widetilde{\mathrm{RP}}$ and the unmodified E, both independently sampled from the initial keys. Thus, a cmel adversary can simulate these oracles itself and forward only the leaking queries. Stopping the simulation just before key resampling shows that the probability that a relevant conditional key distribution has min-entropy below ρ is bounded by the cmel advantage.

From Hyb_i^1 to Hyb_i^5. The transition from Hyb_i^1 to Hyb_i^2 is bounded using the resampling lemma of [ABMS26]. Outside the entropy bad event, the point $s_0^* = \psi_1(k_{\alpha_{i+1}}, x_{i+1})$ has min-entropy at least ρ, since φ is a permutation, while s_1^* is independently uniform. These are the entropy conditions used to apply the lemma; the cmel advantage accounts for the bad event.

The transitions from Hyb_i^j to Hyb_i^{j+1} for $j = 2, 3, 4$ are bounded by classical collision and permutation arguments. The TCCR-specific query restriction is used to rule out the relevant same-user input collisions, and collisions involving distinct users are bounded using the independence and conditional min-entropy of the resampled keys, with the entropy-bad event controlled by cmel.

Moving the replacement past leaking queries. The transition from Hyb_i^5 to Hyb_{i+1}^1 is more delicate because leaking queries may occur between two consecutive construction queries and change the conditional key distributions. To handle this, we introduce additional sub-hybrids that move the key-resampling and oracle-replacement point across the intervening leaking queries one at a time. At each step, the cmel advantage controls the probability that the relevant conditional key distributions have min-entropy below ρ, while the arbitrary reprogramming lemma of [ABKM22] bounds the effect of changing the primitive oracle between consecutive classical queries.

Summing these interval bounds and the bounds for the other transitions yields Proposition 5.2. See Section C for the complete details.

References

AAB+25. Anasuya Acharya, Karen Azari, Mirza Ahad Baig, Dennis Hofheinz, and Chethan Kamath. Securely instantiating 'half gates' garbling in the standard model. In Tibor Jager and Jiaxin Pan, editors, *PKC 2025, Part IV*, volume 15677 of *LNCS*, pages 37–75. Springer, Cham, May 2025. 5

ABC+26. Gorjan Alagic, Maxime Bros, Pierre Ciadoux, Quynh Dang, Thinh Dang, John Kelsey, Jacob Lichtinger, Yi-Kai Liu, Carl Miller, Dustin Moody, Rene Peralta, Ray Perlner, Angela Robinson, Hamilton Silberg, Daniel Smith-Tone, and Noah Waller. Status report on the second round of the additional digital signature schemes for the NIST post-quantum cryptography standardization process. Technical Report NIST IR 8610, National Institute of Standards and Technology, 2026. 9

ABK+24. Gorjan Alagic, Chen Bai, Jonathan Katz, Christian Majenz, and Patrick Struck. Post-quantum security of tweakable Even-Mansour, and applications. In Marc Joye and Gregor Leander, editors, *EUROCRYPT 2024, Part I*, volume 14651 of *LNCS*, pages 310–338. Springer, Cham, May 2024. 5, 6, 16

ABKM22. Gorjan Alagic, Chen Bai, Jonathan Katz, and Christian Majenz. Post-quantum security of the Even-Mansour cipher. In Orr Dunkelman and Stefan Dziembowski, editors, *EUROCRYPT 2022, Part III*, volume 13277 of *LNCS*, pages 458–487. Springer, Cham, May / June 2022. 5, 6, 11, 12, 16, 18, 32

ABMS26. Gorjan Alagic, Chen Bai, Christian Majenz, and Kaiyan Shi. Post-quantum security of block cipher constructions, 2026. To appear at SAC 2026. 5, 6, 16, 18, 25

ABP19. Michel Abdalla, Fabrice Benhamouda, and Alain Passelègue. Algebraic XOR-RKA-secure pseudorandom functions from post-zeroizing multilinear maps. In Steven D. Galbraith and Shiho Moriai, editors, *ASIACRYPT 2019, Part II*, volume 11922 of *LNCS*, pages 386–412. Springer, Cham, December 2019. 5

ADW09. Joël Alwen, Yevgeniy Dodis, and Daniel Wichs. Leakage-resilient public-key cryptography in the bounded-retrieval model. In Shai Halevi, editor, *CRYPTO 2009*, volume 5677 of *LNCS*, pages 36–54. Springer, Berlin, Heidelberg, August 2009. 5

AHU19. Andris Ambainis, Mike Hamburg, and Dominique Unruh. Quantum security proofs using semi-classical oracles. In Alexandra Boldyreva and Daniele Micciancio, editors, *CRYPTO 2019, Part II*, volume 11693 of *LNCS*, pages 269–295. Springer, Cham, August 2019. 6

ALSZ13. Gilad Asharov, Yehuda Lindell, Thomas Schneider, and Michael Zohner. More efficient oblivious transfer and extensions for faster secure computation. In Ahmad-Reza Sadeghi, Virgil D. Gligor, and Moti Yung, editors, *ACM CCS 2013*, pages 535–548. ACM Press, November 2013. 13

BBC+25. Jyotirmoy Basak, Ritam Bhaumik, Amit Kumar Chauhan, Ravindra Jejurikar, Ashwin Jha, Anandarup Roy, André Schrottenloher, and Suprita Talnikar. Post-quantum security of key-alternating feistel ciphers. In Goichiro Hanaoka and Bo-Yin Yang, editors, *ASIACRYPT 2025, Part I*, volume 16245 of *LNCS*, pages 446–478. Springer, Singapore, December 2025. 6, 16

BBFR24. Ryad Benadjila, Charles Bouillaguet, Thibauld Feneuil, and Matthieu Rivain. MQOM — MQ on my mind. Technical report, National Institute of Standards and Technology, 2024. Submission to the NIST PQC Standardization Process: Additional Digital Signature Schemes, Round 2. 3, 9

BCC+24. Dung Bui, Eliana Carozza, Geoffroy Couteau, Dahmun Goudarzi, and Antoine Joux. Faster signatures from MPC-in-the-head. In Kai-Min Chung and Yu Sasaki, editors, *ASIACRYPT 2024, Part I*, volume 15484 of *LNCS*, pages 396–428. Springer, Singapore, December 2024. 4, 14, 15

BCD25. Dung Bui, Kelong Cong, and Cyprien Delpech de Saint Guilhem. Faster VOLEitH signatures from all-but-one vector commitment and half-tree. In Willy Susilo and Josef Pieprzyk, editors, *ACISP 25, Part I*, volume 15658 of *LNCS*, pages 205–223. Springer, Singapore, July 2025. 4, 15

BCG+19. Elette Boyle, Geoffroy Couteau, Niv Gilboa, Yuval Ishai, Lisa Kohl, and Peter Scholl. Efficient pseudorandom correlation generators: Silent OT extension and more. In Alexandra Boldyreva and Daniele Micciancio, editors, *CRYPTO 2019, Part III*, volume 11694 of *LNCS*, pages 489–518. Springer, Cham, August 2019. 2, 5

BCG+21. Elette Boyle, Nishanth Chandran, Niv Gilboa, Divya Gupta, Yuval Ishai, Nishant Kumar, and Mayank Rathee. Function secret sharing for mixed-mode and fixed-point secure computation. In Anne Canteaut and François-Xavier Standaert, editors, *EUROCRYPT 2021, Part II*, volume 12697 of *LNCS*, pages 871–900. Springer, Cham, October 2021. 4, 15

BDF+11. Dan Boneh, Özgür Dagdelen, Marc Fischlin, Anja Lehmann, Christian Schaffner, and Mark Zhandry. Random oracles in a quantum world. In Dong Hoon Lee and Xiaoyun Wang, editors, *ASIACRYPT 2011*, volume 7073 of *LNCS*, pages 41–69. Springer, Berlin, Heidelberg, December 2011. 6

BDK+20. Niklas Büscher, Daniel Demmler, Nikolaos P. Karvelas, Stefan Katzenbeisser, Juliane Krämer, Deevashwer Rathee, Thomas Schneider, and Patrick Struck. Secure two-party computation in a quantum world. In Mauro Conti, Jianying Zhou, Emiliano Casalicchio, and Angelo Spognardi, editors, *ACNS 2020, Part I*, volume 12146 of *LNCS*, pages 461–480. Springer, Cham, October 2020. 5, 6

BDPV08. Guido Bertoni, Joan Daemen, Michaël Peeters, and Gilles Van Assche. On the indifferentiability of the Sponge construction. In Nigel P. Smart, editor, *EUROCRYPT 2008*, volume 4965 of *LNCS*, pages 181–197. Springer, Berlin, Heidelberg, April 2008. 5

Bea96. Donald Beaver. Correlated pseudorandomness and the complexity of private computations. In *28th ACM STOC*, pages 479–488. ACM Press, May 1996. 2, 12

BGI15. Elette Boyle, Niv Gilboa, and Yuval Ishai. Function secret sharing. In Elisabeth Oswald and Marc Fischlin, editors, *EUROCRYPT 2015, Part II*, volume 9057 of *LNCS*, pages 337–367. Springer, Berlin, Heidelberg, April 2015. 4, 15

BGI16. Elette Boyle, Niv Gilboa, and Yuval Ishai. Function secret sharing: Improvements and extensions. In Edgar R. Weippl, Stefan Katzenbeisser, Christopher Kruegel, Andrew C. Myers, and Shai Halevi, editors, *ACM CCS 2016*, pages 1292–1303. ACM Press, October 2016. 15

BGI19. Elette Boyle, Niv Gilboa, and Yuval Ishai. Secure computation with preprocessing via function secret sharing. In Dennis Hofheinz and Alon Rosen, editors, *TCC 2019, Part I*, volume 11891 of *LNCS*, pages 341–371. Springer, Cham, December 2019. 15

BHN+19. Xavier Bonnetain, Akinori Hosoyamada, María Naya-Plasencia, Yu Sasaki, and André Schrottenloher. Quantum attacks without superposition queries: The offline Simon's algorithm. In Steven D. Galbraith and Shiho Moriai, editors, *ASIACRYPT 2019, Part I*, volume 11921 of *LNCS*, pages 552–583. Springer, Cham, December 2019. 12

BHR12. Mihir Bellare, Viet Tung Hoang, and Phillip Rogaway. Foundations of garbled circuits. In Ting Yu, George Danezis, and Virgil D. Gligor, editors, *ACM CCS 2012*, pages 784–796. ACM Press, October 2012. 14

BHT98. Gilles Brassard, Peter Høyer, and Alain Tapp. Quantum cryptanalysis of hash and claw-free functions. In *LATIN'98: Theoretical Informatics*, volume 1380 of *Lecture Notes in Computer Science*, pages 163–169. Springer, 1998. 12

BK03. Mihir Bellare and Tadayoshi Kohno. A theoretical treatment of related-key attacks: RKA-PRPs, RKA-PRFs, and applications. In Eli Biham, editor, *EUROCRYPT 2003*, volume 2656 of *LNCS*, pages 491–506. Springer, Berlin, Heidelberg, May 2003. 5

BK12. Zvika Brakerski and Yael Tauman Kalai. A parallel repetition theorem for leakage resilience. In Ronald Cramer, editor, *TCC 2012*, volume 7194 of *LNCS*, pages 248–265. Springer, Berlin, Heidelberg, March 2012. 5

BN18. Srimanta Bhattacharya and Mridul Nandi. Full indifferentiable security of the xor of two or more random permutations using the χ^2 method. In Jesper Buus Nielsen and Vincent Rijmen, editors, *EUROCRYPT 2018, Part I*, volume 10820 of *LNCS*, pages 387–412. Springer, Cham, April / May 2018. 5

BW13. Dan Boneh and Brent Waters. Constrained pseudorandom functions and their applications. In Kazue Sako and Palash Sarkar, editors, *ASIACRYPT 2013, Part II*, volume 8270 of *LNCS*, pages 280–300. Springer, Berlin, Heidelberg, December 2013. 14

CDMP05. Jean-Sébastien Coron, Yevgeniy Dodis, Cécile Malinaud, and Prashant Puniya. Merkle-Damgård revisited: How to construct a hash function. In Victor Shoup, editor, *CRYPTO 2005*, volume 3621 of *LNCS*, pages 430–448. Springer, Berlin, Heidelberg, August 2005. 5

CGG+25. Hongrui Cui, Chun Guo, Xiaojie Guo, Xiao Wang, Kang Yang, and Yu Yu. Simulation-based security notion of correlation robust hashing with applications to MPC. Cryptology ePrint Archive, Report 2025/1818, 2025. 4, 5, 14

CGW+24. Hongrui Cui, Chun Guo, Xiao Wang, Chenkai Weng, Kang Yang, and Yu Yu. AES-based CCR hash with high security and its application to zero-knowledge proofs. Cryptology ePrint Archive, Report 2024/1271, 2024. 3

CKKZ12. Seung Geol Choi, Jonathan Katz, Ranjit Kumaresan, and Hong-Sheng Zhou. On the security of the "free-XOR" technique. In Ronald Cramer, editor, *TCC 2012*, volume 7194 of *LNCS*, pages 39–53. Springer, Berlin, Heidelberg, March 2012. 2, 5, 14

CLL19. Wonseok Choi, ByeongHak Lee, and Jooyoung Lee. Indifferentiability of truncated random permutations. In Steven D. Galbraith and Shiho Moriai, editors, *ASIACRYPT 2019, Part I*, volume 11921 of *LNCS*, pages 175–195. Springer, Cham, December 2019. 5

CLS15. Benoit Cogliati, Rodolphe Lampe, and Yannick Seurin. Tweaking Even-Mansour ciphers. In Rosario Gennaro and Matthew J. B. Robshaw, editors, *CRYPTO 2015, Part I*, volume 9215 of *LNCS*, pages 189–208. Springer, Berlin, Heidelberg, August 2015. 6

CLY+24. Hongrui Cui, Hanlin Liu, Di Yan, Kang Yang, Yu Yu, and Kaiyi Zhang. ReSolveD: Shorter signatures from regular syndrome decoding and VOLE-in-the-head. In Qiang Tang and Vanessa Teague, editors, *PKC 2024, Part I*, volume 14601 of *LNCS*, pages 229–258. Springer, Cham, April 2024. 15

CT21. Yu Long Chen and Stefano Tessaro. Better security-efficiency trade-offs in permutation-based two-party computation. In Mehdi Tibouchi and Huaxiong Wang, editors, *ASIACRYPT 2021, Part II*, volume 13091 of *LNCS*, pages 275–304. Springer, Cham, December 2021. 2, 3, 5, 10, 13

DG22. Nir Drucker and Shay Gueron. Software optimization of rijndael for modern x86-64 platforms. In *ITNG 2022 19th International Conference on Information Technology—New Generations*, volume 1421 of *Advances in Intelligent Systems and Computing*, pages 147–153. Springer, 2022. 11

DORS08. Yevgeniy Dodis, Rafail Ostrovsky, Leonid Reyzin, and Adam Smith. Fuzzy extractors: How to generate strong keys from biometrics and other noisy data. *SIAM Journal on Computing*, 38(1):97–139, 2008. 5

DP08. Stefan Dziembowski and Krzysztof Pietrzak. Leakage-resilient cryptography. In *49th FOCS*, pages 293–302. IEEE Computer Society Press, October 2008. 5

DR99. Joan Daemen and Vincent Rijmen. AES Proposal: Rijndael. Technical report, National Institute of Standards and Technology, September 1999. Document version 2, September 3, 1999. 11

DR20. Joan Daemen and Vincent Rijmen. *The Design of Rijndael: The Advanced Encryption Standard (AES).* Information Security and Cryptography. Springer, 2 edition, 2020. 11

Ds17. Jack Doerner and abhi shelat. Scaling ORAM for secure computation. In Bhavani M. Thuraisingham, David Evans, Tal Malkin, and Dongyan Xu, editors, *ACM CCS 2017*, pages 523–535. ACM Press, October / November 2017. 15

EM93. Shimon Even and Yishay Mansour. A construction of a cipher from a single pseudorandom permutation. In Hideki Imai, Ronald L. Rivest, and Tsutomu Matsumoto, editors, *ASIACRYPT'91*, volume 739 of *LNCS*, pages 210–224. Springer, Berlin, Heidelberg, November 1993. 6

GBJ+23. Aldo Gunsing, Ritam Bhaumik, Ashwin Jha, Bart Mennink, and Yaobin Shen. Revisiting the indifferentiability of the sum of permutations. In Helena Handschuh and Anna Lysyanskaya, editors, *CRYPTO 2023, Part III*, volume 14083 of *LNCS*, pages 628–660. Springer, Cham, August 2023. 5

GBK16. Praveen Gauravaram, Nasour Bagheri, and Lars R. Knudsen. Building indifferentiable compression functions from the PGV compression functions. *DCC*, 78(2):547–581, 2016. 5

GGM84. Oded Goldreich, Shafi Goldwasser, and Silvio Micali. How to construct random functions (extended abstract). In *25th FOCS*, pages 464–479. IEEE Computer Society Press, October 1984. 14

GI14. Niv Gilboa and Yuval Ishai. Distributed point functions and their applications. In Phong Q. Nguyen and Elisabeth Oswald, editors, *EUROCRYPT 2014*, volume 8441 of *LNCS*, pages 640–658. Springer, Berlin, Heidelberg, May 2014. 4, 15

GKW+20. Chun Guo, Jonathan Katz, Xiao Wang, Chenkai Weng, and Yu Yu. Better concrete security for half-gates garbling (in the multi-instance setting). In Daniele Micciancio and Thomas Ristenpart, editors, *CRYPTO 2020, Part II*, volume 12171 of *LNCS*, pages 793–822. Springer, Cham, August 2020. 2, 3, 4, 5, 7, 8, 11, 14

GKWY20. Chun Guo, Jonathan Katz, Xiao Wang, and Yu Yu. Efficient and secure multiparty computation from fixed-key block ciphers. In *2020 IEEE Symposium on Security and Privacy*, pages 825–841. IEEE Computer Society Press, May 2020. 2, 3, 4, 5, 8, 10, 11, 13, 14

GOR11. Vipul Goyal, Adam O'Neill, and Vanishree Rao. Correlated-input secure hash functions. In Yuval Ishai, editor, *TCC 2011*, volume 6597 of *LNCS*, pages 182–200. Springer, Berlin, Heidelberg, March 2011. 5

GWYY25. Chun Guo, Xiao Wang, Kang Yang, and Yu Yu. On tweakable correlation robust hashing against key leakages. *DCC*, 93(8):3375–3412, 2025. 2, 3, 4, 5, 7, 8, 9, 10, 13

GYW+22. Xiaojie Guo, Kang Yang, Xiao Wang, Wenhao Zhang, Xiang Xie, Jiang Zhang, and Zheli Liu. Half-tree: Halving the cost of tree expansion in COT and DPF. Cryptology ePrint Archive, Paper 2022/1431, 2022. Version 20231221:062849. 14

GYW+23. Xiaojie Guo, Kang Yang, Xiao Wang, Wenhao Zhang, Xiang Xie, Jiang Zhang, and Zheli Liu. Half-tree: Halving the cost of tree expansion in COT and DPF. In Carmit Hazay and Martijn Stam, editors, *EUROCRYPT 2023, Part I*, volume 14004 of *LNCS*, pages 330–362. Springer, Cham, April 2023. 4, 13, 14, 15

HJ24. Janik Huth and Antoine Joux. MPC in the head using the subfield bilinear collision problem. In Leonid Reyzin and Douglas Stebila, editors, *CRYPTO 2024, Part I*, volume 14920 of *LNCS*, pages 39–70. Springer, Cham, August 2024. 15

Hos25. Akinori Hosoyamada. Post-quantum security of keyed sponge-based constructions through a modular approach. In Goichiro Hanaoka and Bo-Yin Yang, editors, *ASIACRYPT 2025, Part I*, volume 16245 of *LNCS*, pages 411–445. Springer, Singapore, December 2025. 5, 6, 16

HS18. Akinori Hosoyamada and Yu Sasaki. Quantum Demiric-Selçuk meet-in-the-middle attacks: Applications to 6-round generic Feistel constructions. In Dario Catalano and Roberto De Prisco, editors, *SCN 18*, volume 11035 of *LNCS*, pages 386–403. Springer, Cham, September 2018. 12

HY18. Akinori Hosoyamada and Kan Yasuda. Building quantum-one-way functions from block ciphers: Davies-Meyer and Merkle-Damgård constructions. In Thomas Peyrin and Steven Galbraith, editors, *ASIACRYPT 2018, Part I*, volume 11272 of *LNCS*, pages 275–304. Springer, Cham, December 2018. 3

IKNP03. Yuval Ishai, Joe Kilian, Kobbi Nissim, and Erez Petrank. Extending oblivious transfers efficiently. In Dan Boneh, editor, *CRYPTO 2003*, volume 2729 of *LNCS*, pages 145–161. Springer, Berlin, Heidelberg, August 2003. 2, 5, 6, 12, 13

IKOS07. Yuval Ishai, Eyal Kushilevitz, Rafail Ostrovsky, and Amit Sahai. Zero-knowledge from secure multiparty computation. In David S. Johnson and Uriel Feige, editors, *39th ACM STOC*, pages 21–30. ACM Press, June 2007. 14

Int19. International Organization for Standardization. *ISO/IEC 29192-6:2019: Information technology – Lightweight cryptography – Part 6: Message authentication codes (MACs)*, September 2019. Includes Chaskey-12. 6

JST21. Joseph Jaeger, Fang Song, and Stefano Tessaro. Quantum key-length extension. In Kobbi Nissim and Brent Waters, editors, *TCC 2021, Part I*, volume 13042 of *LNCS*, pages 209–239. Springer, Cham, November 2021. 6

KCP16. John Kelsey, Shu-jen Chang, and Ray Perlner. SHA-3 derived functions: cSHAKE, KMAC, TupleHash, and ParallelHash. NIST Special Publication 800-185, National Institute of Standards and Technology, December 2016. 6

KKKS15. Carmen Kempka, Ryo Kikuchi, Susumu Kiyoshima, and Koutarou Suzuki. Garbling scheme for formulas with constant size of garbled gates. In Tetsu Iwata and Jung Hee Cheon, editors, *ASIACRYPT 2015, Part I*, volume 9452 of *LNCS*, pages 758–782. Springer, Berlin, Heidelberg, November / December 2015. 5

KLLN16. Marc Kaplan, Gaëtan Leurent, Anthony Leverrier, and María Naya-Plasencia. Quantum differential and linear cryptanalysis. *IACR Trans. Symm. Cryptol.*, 2016(1):71–94, 2016. 3

KLS25. Seongkwang Kim, Byeonghak Lee, and Mincheol Son. Relaxed vector commitment for shorter signatures. In Serge Fehr and Pierre-Alain Fouque, editors, *EUROCRYPT 2025, Part IV*, volume 15604 of *LNCS*, pages 427–455. Springer, Cham, May 2025. 15

KM12. Hidenori Kuwakado and Masakatu Morii. Security on the quantum-type even-mansour cipher. In *2012 International Symposium on Information Theory and its Applications*, pages 312–316. IEEE, 2012. 12

KR96. Joe Kilian and Phillip Rogaway. How to protect DES against exhaustive key search. In Neal Koblitz, editor, *CRYPTO'96*, volume 1109 of *LNCS*, pages 252–267. Springer, Berlin, Heidelberg, August 1996. 6

KS08. Vladimir Kolesnikov and Thomas Schneider. Improved garbled circuit: Free XOR gates and applications. In Luca Aceto, Ivan Damgård, Leslie Ann Goldberg, Magnús M. Halldórsson, Anna Ingólfsdóttir, and Igor Walukiewicz, editors, *ICALP 2008, Part II*, volume 5126 of *LNCS*, pages 486–498. Springer, Berlin, Heidelberg, July 2008. 2, 14

KX26. Haruhisa Kosuge and Keita Xagawa. Towards formal security proofs of mqom. Cryptology ePrint Archive, Paper 2026/629, 2026. 3, 9

MMV+14. Nicky Mouha, Bart Mennink, Anthony Van Herrewege, Dai Watanabe, Bart Preneel, and Ingrid Verbauwhede. Chaskey: An efficient MAC algorithm for 32-bit microcontrollers. In Antoine Joux and Amr M. Youssef, editors, *SAC 2014*, volume 8781 of *LNCS*, pages 306–323. Springer, Cham, August 2014. 6

MRH04. Ueli M. Maurer, Renato Renner, and Clemens Holenstein. Indifferentiability, impossibility results on reductions, and applications to the random oracle methodology. In Moni Naor, editor, *TCC 2004*, volume 2951 of *LNCS*, pages 21–39. Springer, Berlin, Heidelberg, February 2004. 5

Nat01. National Institute of Standards and Technology. Advanced Encryption Standard (AES). Federal Information Processing Standards Publication NIST FIPS 197-upd1, U.S. Department of Commerce, Washington, DC, 2001. Updated May 9, 2023. 2, 11

Nat24a. National Institute of Standards and Technology. NIST Releases First 3 Finalized Post-Quantum Encryption Standards. https://www.nist.gov/news-events/news/2024/08/nist-releases-first-3-finalized-post-quantum-encryption-standards, August 2024. Accessed: 2026-05-15. 2

Nat24b. National Institute of Standards and Technology. PRE-DRAFT Call for Comments: NIST Proposes to Standardize a Wider Variant of AES. NIST Special Publication 800-197, Initial Preliminary Draft, December 2024. Date published: December 23, 2024; public comment period closed June 25, 2025; accessed May 8, 2026. 11

Nat26. National Institute of Standards and Technology. Second pre-draft call for comments: GCM and GMAC block cipher modes of operation. NIST Special Publication 800-38D Rev. 1, June 2026. Published June 1, 2026. 11

NK26. Kohei Nakamura and Akinori Kawachi. Improving the post-quantum security reduction of oblvious transfer extension. SCIS 2026, 2026 Symposium on Cryptography and Information Security, 2026. 6

NS12. Moni Naor and Gil Segev. Public-key cryptosystems resilient to key leakage. *SIAM Journal on Computing*, 41(4):772–814, 2012. 5

Pat09. Jacques Patarin. The "coefficients H" technique (invited talk). In Roberto Maria Avanzi, Liam Keliher, and Francesco Sica, editors, *SAC 2008*, volume 5381 of *LNCS*, pages 328–345. Springer, Berlin, Heidelberg, August 2009. 3

Pie09. Krzysztof Pietrzak. A leakage-resilient mode of operation. In Antoine Joux, editor, *EUROCRYPT 2009*, volume 5479 of *LNCS*, pages 462–482. Springer, Berlin, Heidelberg, April 2009. 5

Rab81. Michael O. Rabin. How to exchange secrets by oblivious transfer. Technical Report TR-81, Harvard University, 1981. 2, 12

Roy22. Lawrence Roy. SoftSpokenOT: Quieter OT extension from small-field silent VOLE in the minicrypt model. In Yevgeniy Dodis and Thomas Shrimpton, editors, *CRYPTO 2022, Part I*, volume 13507 of *LNCS*, pages 657–687. Springer, Cham, August 2022. 2, 3, 4, 5, 7, 8, 10, 13

RR21. Mike Rosulek and Lawrence Roy. Three halves make a whole? Beating the half-gates lower bound for garbled circuits. In Tal Malkin and Chris Peikert, editors, *CRYPTO 2021, Part I*, volume 12825 of *LNCS*, pages 94–124, Virtual Event, August 2021. Springer, Cham. 4, 14

RSS11. Thomas Ristenpart, Hovav Shacham, and Thomas Shrimpton. Careful with composition: Limitations of the indifferentiability framework. In Kenneth G. Paterson, editor, *EUROCRYPT 2011*, volume 6632 of *LNCS*, pages 487–506. Springer, Berlin, Heidelberg, May 2011. 5

SI26a. Rentaro Shiba and Tetsu Iwata. Multi-key security in the quantum world: Revisiting tweakable even-mansour and fx. IACR Cryptology ePrint Archive 2026/382, 2026. 5, 6, 16

SI26b. Rentaro Shiba and Tetsu Iwata. Post-quantum security of tweakable key-alternating feistel ciphers in the multi-key setting. In *SCN 2026*, 2026. To appear. 6

TMK+25. Meltem Sönmez Turan, Kerry McKay, Jinkeon Kang, John Kelsey, and Donghoon Chang. Ascon-based lightweight cryptography standards for constrained devices: Authenticated encryption, hash, and extendable output functions. NIST Special Publication 800-232, National Institute of Standards and Technology, August 2025. 6

Unr15. Dominique Unruh. Revocable quantum timed-release encryption. *J. ACM*, 62(6):49:1–49:76, 2015. 6

WKK+25. Yalan Wang, Bryan Kumara, Harsh Kasyap, Liqun Chen, Sumanta Sarkar, Christopher J. P. Newton, Carsten Maple, and Ugur-Ilker Atmaca. An improved vector commitment construction with applications to signatures. In *TrustCom 2025*, pages 1634–1639. IEEE, 2025. See also https://eprint.iacr.org/2025/1411. 4, 15

Yao86. Andrew Chi-Chih Yao. How to generate and exchange secrets (extended abstract). In *27th FOCS*, pages 162–167. IEEE Computer Society Press, October 1986. 2, 6, 14

YWL+20. Kang Yang, Chenkai Weng, Xiao Lan, Jiang Zhang, and Xiao Wang. Ferret: Fast extension for correlated OT with small communication. In Jay Ligatti, Xinming Ou, Jonathan Katz, and Giovanni Vigna, editors, *ACM CCS 2020*, pages 1607–1626. ACM Press, November 2020. 14

YWZ20. Kang Yang, Xiao Wang, and Jiang Zhang. More efficient MPC from improved triple generation and authenticated garbling. In Jay Ligatti, Xinming Ou, Jonathan Katz, and Giovanni Vigna, editors, *ACM CCS 2020*, pages 1627–1646. ACM Press, November 2020. 14

ZRE15. Samee Zahur, Mike Rosulek, and David Evans. Two halves make a whole - reducing data transfer in garbled circuits using half gates. In Elisabeth Oswald and Marc Fischlin, editors, *EUROCRYPT 2015, Part II*, volume 9057 of *LNCS*, pages 220–250. Springer, Berlin, Heidelberg, April 2015. 2, 4, 14

Supporting Materials

AI Tool Use Disclosure

The authors used generative AI tools (OpenAI GPT-5.6 Sol and GPT-6 Astra) for editorial assistance, making figures/tables, identifying possible issues in the exposition and technical arguments, and assisting in drafting the proof of Propositions 3.2 and F.1. Except for the proof of Propositions 3.2 and F.1, all mathematical content, including definitions, theorem statements, and proofs, was written by the authors. The proof of Propositions 3.2 and F.1, including the AI-assisted portions, was carefully reviewed and verified by the authors. The authors take full responsibility for the correctness, originality, and integrity of the manuscript.

A Parameter Index

Table 4. Main parameters used throughout the paper.

Parameter	Meaning	Parameter	Meaning
n	Block/output length and length of each secret key	m	Tweak length (and ideal-cipher key length)
u	Number of users (secret keys)	q_E	Total number of quantum primitive queries
q_C	Total number of classical construction queries	μ	Maximum number of construction queries having the same tweak
$q_{L,\max}$	Maximum total number of leaking queries	$Q_{L,\max}$	Maximum number of leaking queries for each user
M	Maximum number of distinct users on which leaking queries are made	ρ	(Conditional) min-entropy threshold used in the security bound
d	Lower bound on the initial key min-entropy in Propositions 3.1 and 3.2	w	Entropy-loss parameter in the cmel bounds of Propositions 3.1 and 3.2

B Proof of Proposition 5.1

Proof (of Proposition 5.1). Below, we show the claim for mTPRPL* and mTCCRL. The proof for mTPRPL and mTCRL is similar.

First, consider modifying the definition of the mTCCRL advantage of Eq. (1): For each (classical) query $(\alpha; x, t, b)$ to the $\mathcal{O}_{H,\mathbf{k}}^{\mathrm{mTCCRL}}$ oracle in the real world and the f oracle in the ideal world ($1 \le \alpha \le u$), modify the response by taking the XOR with $\sigma(x)$.

Regarding the oracles in the real world, we observe that

$$\begin{aligned}\mathcal{O}_{H,\mathbf{k}}^{\mathrm{mTCCRL}}(\alpha; x, t, b) \oplus \sigma(x) &= H^E(x \oplus k_\alpha, t) \oplus b \cdot k_\alpha \oplus \sigma(x) \\ &= \overline{\mathsf{TBC}}_{\mathbf{k}}^E(\alpha; x, t, b)\end{aligned} \tag{9}$$

holds by the definitions of the hash function H^E in Eq. (4) and TBC_k^E in Eq. (6) (here we used the assumption that σ is linear). Moreover, the output distribution of f is not changed by taking the XOR with $\sigma(x)$ because f is a random function.

Hence, for each adversary $\mathcal{A}$ on the mTCCRL security of H^E, there exists an adversary $\mathcal{B}$ against the mTPRPL* security of TBC such that $\mathcal{B}$ makes exactly the same number of queries as $\mathcal{A}$ and

$$\begin{aligned}&\mathsf{Adv}_{H,\mathbf{R},\mathcal{P}}^{\mathrm{mTCCRL}}(\mathcal{A}) \\ &= \left|\Pr_{\mathbf{k},E}\left[1 \leftarrow \mathcal{B}^{\overline{\mathsf{TBC}}_{\mathbf{k}}^E, \mathcal{L}_{\mathbf{k}}, |E^{\pm}\rangle}\right] - \Pr_{f,\mathbf{k},E}\left[1 \leftarrow \mathcal{B}^{f, \mathcal{L}_{\mathbf{k}}, |E^{\pm}\rangle}\right]\right| \\ &\le \mathsf{Adv}_{\mathsf{TBC},\mathbf{R},\mathcal{P}}^{\mathrm{mTPRPL}^*}(\mathcal{B}) + \left|\Pr_{f,\mathbf{k},E}\left[1 \leftarrow \mathcal{B}^{f, \mathcal{L}_{\mathbf{k}}, |E^{\pm}\rangle}\right] - \Pr_{\widetilde{\mathrm{RP}},\mathbf{k},E}\left[1 \leftarrow \mathcal{B}^{\widetilde{\mathrm{RP}}, \mathcal{L}_{\mathbf{k}}, |E^{\pm}\rangle}\right]\right| \\ &\le \mathsf{Adv}_{\mathsf{TBC},\mathbf{R},\mathcal{P}}^{\mathrm{mTPRPL}^*}(\mathcal{B}) + \frac{q_C \mu}{2^n}\end{aligned} \tag{10}$$

holds, where the first inequality follows from the definition of mTPRPL* security, and the second inequality holds because $\mathcal{A}$ makes at most μ queries of the form $(*; *, t, *)$ to the first oracle for each t, and so does $\mathcal{B}$ (f is distinguished from $\widetilde{\mathrm{RP}}$ only when a collision is found in the outputs of $f(\alpha; \cdot, t, b)$ for some (α, t, b)). Therefore, the claim of the proposition follows. □

C The Remaining Part of the Proof for Proposition 5.2

This section provides the remaining part of the proof for Proposition 5.2.

Bounding the Distance between Hyb_i^1 and Hyb_i^2.

Lemma C.1. *Suppose* $\mathsf{Adv}_{\mathbf{R},\mathcal{P},\rho}^{\mathrm{cmel}}(M, q_{L,\max}, Q_{L,\max}) < 1$ *and* $\rho \le n$. *Then, it holds that*

$$\begin{aligned}&\left|\Pr[1 \leftarrow \mathcal{A} \text{ in } \mathsf{Hyb}_i^1] - \Pr[1 \leftarrow \mathcal{A} \text{ in } \mathsf{Hyb}_i^2]\right| \\ &\quad\le \frac{4\sqrt{q_E/2^\rho} + \mathsf{Adv}_{\mathbf{R},\mathcal{P},\rho}^{\mathrm{cmel}}(M, q_{L,\max}, Q_{L,\max})}{1 - \mathsf{Adv}_{\mathbf{R},\mathcal{P},\rho}^{\mathrm{cmel}}(M, q_{L,\max}, Q_{L,\max})} + \mathsf{Adv}_{\mathbf{R},\mathcal{P},\rho}^{\mathrm{cmel}}(M, q_{L,\max}, Q_{L,\max}).\end{aligned} \tag{11}$$

To prove this lemma, we need the resampling lemma below.

Lemma C.2 (Resampling lemma of [ABMS26]). *Suppose that a (quantum) distinguisher* $\mathcal{D}$ *runs in the following experiment.*

1. *Let* E_0 *be an* n*-bit ideal cipher with* m*-bit keys.* $\mathcal{D}$ *is given quantum oracle access to* $|E_0^{\pm}\rangle$. $\mathcal{D}$ *performs some computation, making quantum queries to* $|E_0^{\pm}\rangle$, *and outputs the description of a distribution* D *over* $\{0,1\}^m \times \{0,1\}^n \times \{0,1\}^n$.
2. *A tuple* (t_*, s_0, s_1) *is sampled from* D, *a cipher* $E_1 \in \mathsf{BC}(m,n)$ *is defined by* $E_1 := E_0 \circ \mathsf{swap}_{s_0,s_1}^{t_*}$ (swap *is defined as in Eq. (8)), and a uniform random bit* r *is chosen. Then,* $\mathcal{D}$ *is given quantum oracle access to* $|E_r^{\pm}\rangle$ *and the tuple* (t_*, s_0, s_1). *After performing some computation and making quantum queries to* $|E_r^{\pm}\rangle$, *the distinguisher* $\mathcal{D}$ *returns a bit* r' *as an output.*

If the min-entropy of the distribution D is always at least ρ, then we have

$$|\Pr[1 \leftarrow \mathcal{D} \mid r = 1] - \Pr[1 \leftarrow \mathcal{D} \mid r = 0]| \leq 4\sqrt{q/2^{\rho-n}}, \tag{12}$$

where q is the number of queries made by $\mathcal{D}$ to $|E_0^{\pm}\rangle$ in the first stage.

The proof of Lemma C.1 is as follows.

Proof (of Lemma C.1). Let Bad be the event that the min-entropy of $\mathcal{R}_{\alpha_{i+1}}|_{\mathsf{L}_{i+1}}$ is less than ρ, when $\mathcal{A}$ submits the $(i+1)$st construction query. In addition, let $\mathcal{C}$ be an adversary against the cmel experiment of p. 9 that runs $\mathcal{A}$ until $\mathcal{A}$ submits the $(i+1)$th construction query, simulating the oracles of $\widetilde{\mathrm{RP}}$ and $|E^{\pm}\rangle$ by sampling these oracles by itself, and forwarding leaking queries to the oracle given to $\mathcal{C}$ itself. Then,

$$\Pr[\mathsf{Bad}] \leq \mathsf{Adv}^{\mathrm{cmel}}_{\mathbf{R},\mathcal{P},\rho}(\mathcal{C}) \leq \mathsf{Adv}^{\mathrm{cmel}}_{\mathbf{R},\mathcal{P},\rho}(M, q_{L,\max}, Q_{L,\max}) \tag{13}$$

holds in both Hyb_i^1 and Hyb_i^2.

Let $\mathcal{D}$ be a distinguisher for the resampling lemma (Lemma C.2) defined as follows.

1-1. As an initial set-up, $\mathcal{D}$ is given quantum oracle access to $|E_0^{\pm}\rangle$. $\mathcal{D}$ samples a tweakable random permutation $\widetilde{\mathrm{RP}}$ and keys $\mathbf{k} = (k_1, \ldots, k_u)$ as $k_\alpha \leftarrow \mathcal{R}_\alpha$ by itself.

1-2. $\mathcal{D}$ runs the adversary $\mathcal{A}$. The construction oracle, the leaking oracle, and the primitive oracle for $\mathcal{A}$ are simulated with $\widetilde{\mathrm{RP}}$, $\mathcal{L}_{\mathbf{k}}$, and the $|E_0^{\pm}\rangle$ oracle given to $\mathcal{D}$, respectively.

1-3. When $\mathcal{A}$ submits the $(i+1)$st construction query $(\alpha_{i+1}; x_{i+1}, t_{i+1}, b_{i+1})$, the distinguisher $\mathcal{D}$ computes the min-entropy of $\mathcal{R}_{\alpha_{i+1}}|_{\mathsf{L}_{i+1}}$ to check if Bad has happened or not. If Bad has not happened, then $\mathcal{D}$ outputs the description of the distribution D such that

$$D(t_*, s_0, s_1) = \begin{cases} \Pr_{k_{\alpha_{i+1}} \leftarrow \mathcal{R}_{\alpha_{i+1}}|_{\mathsf{L}_{i+1}}}\left[\psi_1(k_{\alpha_{i+1}}, x_{i+1}) = s_0\right] \cdot \frac{1}{2^n} & \text{if } t_* = t_{i+1}, \\ 0 & \text{if } t_* \neq t_{i+1}. \end{cases}$$

If Bad has happened, then $\mathcal{D}$ outputs the uniform distribution over $\{0,1\}^{m+2n}$ as D.

2-1. A tuple (t_*, s_0, s_1) is sampled from the distribution D and given to the distinguisher $\mathcal{D}$. Then, $\mathcal{D}$ is given quantum oracle access to $|E_r^{\pm}\rangle$ (r is a random bit and $E_1 := E_0 \circ \mathsf{swap}^{t_*}_{s_0,s_1}$).

2-2. $\mathcal{D}$ sets $k_{\alpha_{i+1}} \leftarrow \varphi^{-1}(s_0) \oplus x_{i+1}$, and resamples $k_\alpha \leftarrow \mathcal{R}_\alpha|_{\mathsf{L}_{i+1}}$ for $1 \leq \alpha \neq \alpha_{i+1} \leq u$.

2-3. $\mathcal{D}$ continues running $\mathcal{A}$. The construction oracle, the leaking oracle, and the primitive oracle for $\mathcal{A}$ are simulated with $\overline{\mathsf{TBC}}_{\mathbf{k}}^{(E_r)^{\mathsf{T}_i,\mathbf{k}}}$, $\mathcal{L}_{\mathbf{k}}$, and the quantum oracle of $(E_r)^{\mathsf{T}_i,\mathbf{k}}$ (and its inverse), respectively.

2-4. When $\mathcal{A}$ returns an output bit, $\mathcal{D}$ returns it as its own output.

Then,

$$\Pr[1 \leftarrow \mathcal{D} \mid r = 0, \neg\mathsf{Bad}] = \Pr\left[1 \leftarrow \mathcal{A} \text{ in } \mathsf{Hyb}_i^1 \mid \neg\mathsf{Bad}\right],$$
$$\Pr[1 \leftarrow \mathcal{D} \mid r = 1, \neg\mathsf{Bad}] = \Pr\left[1 \leftarrow \mathcal{A} \text{ in } \mathsf{Hyb}_i^2 \mid \neg\mathsf{Bad}\right]$$

follow from the definition of $\mathcal{D}$. In addition,

$$|\Pr[1 \leftarrow \mathcal{D} \mid r = 0] - \Pr[1 \leftarrow \mathcal{D} \mid r = 1]| \leq 4\sqrt{q_E/2^{\rho}} \tag{14}$$

follows from Lemma C.2 (the value ρ in Lemma C.2 corresponds to $(\rho + n)$ here because the min-entropy of D is at least $(\rho + n)$ if Bad does not happen, and $(m + 2n)$ if Bad happens). Moreover,

$$\Pr[\mathsf{Bad} \mid r = 0] = \Pr[\mathsf{Bad} \mid r = 1] = \Pr[\mathsf{Bad}] \leq \mathsf{Adv}^{\mathrm{cmel}}_{\mathbf{R},\mathcal{P},\rho}(M, q_{L,\max}, Q_{L,\max}) \tag{15}$$

holds.

Therefore, we have

$$
\begin{aligned}
&\left|\Pr\left[1 \leftarrow \mathcal{A} \text{ in } \mathsf{Hyb}_i^1\right] - \Pr\left[1 \leftarrow \mathcal{A} \text{ in } \mathsf{Hyb}_i^2\right]\right| \\
&\leq \left|\Pr\left[1 \leftarrow \mathcal{A} \text{ in } \mathsf{Hyb}_i^1 \mid \neg\mathsf{Bad}\right] - \Pr\left[1 \leftarrow \mathcal{A} \text{ in } \mathsf{Hyb}_i^2 \mid \neg\mathsf{Bad}\right]\right| + \Pr\left[\mathsf{Bad}\right] \\
&= \left|\Pr\left[1 \leftarrow \mathcal{D} \mid r = 0, \neg\mathsf{Bad}\right] - \Pr\left[1 \leftarrow \mathcal{D} \mid r = 1, \neg\mathsf{Bad}\right]\right| + \Pr\left[\mathsf{Bad}\right] \\
&= \left|\Pr\left[1 \leftarrow \mathcal{D} \wedge \neg\mathsf{Bad} \mid r = 0\right] - \Pr\left[1 \leftarrow \mathcal{D} \wedge \neg\mathsf{Bad} \mid r = 1\right]\right| \cdot \frac{1}{\Pr[\neg\mathsf{Bad}]} + \Pr\left[\mathsf{Bad}\right] \\
&\leq \underbrace{\left|\Pr\left[1 \leftarrow \mathcal{D} \mid r = 0\right] - \Pr\left[1 \leftarrow \mathcal{D} \mid r = 1\right]\right|}_{\leq 4\sqrt{q_E/2^\rho} \text{ by Eq. (14)}} \cdot \frac{1}{\Pr[\neg\mathsf{Bad}]} \\
&\quad + \underbrace{\left|\Pr\left[1 \leftarrow \mathcal{D} \wedge \mathsf{Bad} \mid r = 0\right] - \Pr\left[1 \leftarrow \mathcal{D} \wedge \mathsf{Bad} \mid r = 1\right]\right|}_{\leq \Pr[\mathsf{Bad}] \text{ by Eq. (15)}} \cdot \frac{1}{\Pr[\neg\mathsf{Bad}]} \\
&\quad + \Pr\left[\mathsf{Bad}\right] \\
&\leq \frac{4\sqrt{q_E/2^\rho}}{1 - \Pr[\mathsf{Bad}]} + \frac{\Pr\left[\mathsf{Bad}\right]}{1 - \Pr\left[\mathsf{Bad}\right]} + \Pr\left[\mathsf{Bad}\right] \\
&\leq \frac{4\sqrt{q_E/2^\rho} + \mathsf{Adv}^{\mathrm{cmel}}_{\mathbf{R},\mathcal{P},\rho}(M, q_{L,\max}, Q_{L,\max})}{1 - \mathsf{Adv}^{\mathrm{cmel}}_{\mathbf{R},\mathcal{P},\rho}(M, q_{L,\max}, Q_{L,\max})} + \mathsf{Adv}^{\mathrm{cmel}}_{\mathbf{R},\mathcal{P},\rho}(M, q_{L,\max}, Q_{L,\max}). \qquad (16)
\end{aligned}
$$

Therefore, the claim of the lemma follows. □

Bounding the Distance between Hyb_i^2 and Hyb_i^3.

Lemma C.3. *Suppose* $\rho \leq n$. *Then,* $\left|\Pr[1 \leftarrow \mathcal{A} \text{ in } \mathsf{Hyb}_i^2] - \Pr[1 \leftarrow \mathcal{A} \text{ in } \mathsf{Hyb}_i^3]\right| \leq 2\mu/2^\rho + \mathsf{Adv}^{\mathrm{cmel}}_{\mathbf{R},\mathcal{P},\rho}(M, q_{L,\max}, Q_{L,\max})$.

Proof. The difference between Hyb_i^2 and Hyb_i^3 lies in how the response to the $(i+1)$st query (i.e., the value y_{i+1}) is sampled. The experiment Hyb_i^2 samples y_{i+1} as

$$
\begin{aligned}
y_{i+1} &\leftarrow \overline{\mathsf{TBC}}^{(E')^{\mathsf{T}_i,\mathbf{k}}}_{k_{\alpha_{i+1}}}(x_{i+1}, t_{i+1}, b_{i+1}) \\
&= \psi_2\left(k_{\alpha_{i+1}}, (E')^{\mathsf{T}_i,\mathbf{k}}(t_{i+1}, s_0^*), b_{i+1}\right),
\end{aligned}
$$

and Hyb_i^3 samples y_{i+1} as

$$
y_{i+1} \leftarrow \widetilde{y} := \psi_2\left(k_{\alpha_{i+1}}, E^{\mathsf{T}_i,\mathbf{k}}(t_{i+1}, s_1^*), b_{i+1}\right)
$$

where s_1^* is a uniform random value in both Hyb_i^2 and Hyb_i^3. Hence, the two experiments are identical if

$$
(E')^{\mathsf{T}_i,\mathbf{k}}(t_{i+1}, s_0^*) = E^{\mathsf{T}_i,\mathbf{k}}(t_{i+1}, s_1^*). \qquad (17)
$$

Let

$$
\begin{aligned}
I_{i+1} &:= \left\{ j \mid 1 \leq j \leq i,\ t_j = t_{i+1} \right\}, \\
K_{\mathsf{bad}} &:= \left\{ \psi_1(k_{\alpha_j}, x_j) \mid j \in I_{i+1} \text{ and } \alpha_j \neq \alpha_{i+1} \right\}, \\
X_{\mathsf{bad}} &:= \left\{ \psi_1(k_{\alpha_j}, x_j) \mid j \in I_{i+1} \right\},
\end{aligned}
$$

and define $\mathsf{Bad}_{\mathrm{coll}}$ as the event that $\psi_1(k_{\alpha_{i+1}}, x_{i+1}) \in K_{\mathsf{bad}}$ or $s_1^* \in X_{\mathsf{bad}}$. Then, the following claim holds.

Claim. Assume $\mathsf{Bad}_{\mathrm{coll}}$ does not hold in Hyb_i^2, i.e., $\psi_1(k_{\alpha_{i+1}}, x_{i+1}) \notin K_{\mathsf{bad}}$ and $s_1^* \notin X_{\mathsf{bad}}$. Then, $(E')^{\mathsf{T}_i,\mathbf{k}}(t_{i+1}, s_0^*) = E^{\mathsf{T}_i,\mathbf{k}}(t_{i+1}, s_1^*)$ holds (and thus Hyb_i^2 and Hyb_i^3 are identical).

Proof. Let j be an index in I_{i+1} (which implies $t_j = t_{i+1}$). Since $(\alpha_j, x_j, t_j) \neq (\alpha_{i+1}, x_{i+1}, t_{i+1})$ is assumed (see p. 17)[15], we have that $\alpha_{i+1} = \alpha_j \wedge x_{i+1} \neq x_j$ or $\alpha_{i+1} \neq \alpha_j$. In both cases,

$$s_0^* = \psi_1(k_{\alpha_{i+1}}, x_{i+1}) \neq \psi_1(k_{\alpha_j}, x_j).$$

holds. (In the former case, the non-equation is trivial because $k_{\alpha_{i+1}} = k_{\alpha_j}$ and $x_{i+1} \neq x_j$. In the latter case, the non-equation follows from the assumption of the claim that $\psi_1(k_{\alpha_{i+1}}, x_{i+1}) \notin K_{\mathsf{bad}}$.)

In addition, $s_1^* \neq \psi_1(k_{\alpha_j}, x_j)$ follows from the assumption $s_1^* \notin X_{\mathsf{bad}}$. Therefore, we have

$$E'(t_j, \psi_1(k_{\alpha_j}, x_j)) = E(t_j, \psi_1(k_{\alpha_j}, x_j)) \text{ for } 1 \leq j \leq i,$$

which implies

$$\begin{aligned}
\mathsf{REP}_{\mathsf{T}_i,E',\mathbf{k}} &= \mathsf{swap}^{t_1}_{E'(t_1,\psi_1(k_{\alpha_1},x_1)),\psi_2(k_{\alpha_1},y_1,b_1)} \circ \cdots \\
&\qquad \cdots \circ \mathsf{swap}^{t_i}_{E'(t_i,\psi_1(k_{\alpha_i},x_i)),\psi_2(k_{\alpha_i},y_i,b_i)} \\
&= \mathsf{swap}^{t_1}_{E(t_1,\psi_1(k_{\alpha_1},x_1)),\psi_2(k_{\alpha_1},y_1,b_1)} \circ \cdots \\
&\qquad \cdots \circ \mathsf{swap}^{t_i}_{E(t_i,\psi_1(k_{\alpha_i},x_i)),\psi_2(k_{\alpha_i},y_i,b_i)} \\
&= \mathsf{REP}_{\mathsf{T}_i,E,\mathbf{k}}. \qquad (18)
\end{aligned}$$

Therefore,

$$\begin{aligned}
(E')^{\mathsf{T}_i,\mathbf{k}}(t_{i+1}, s_0^*) &= \mathsf{proj}_2\left(\mathsf{REP}_{\mathsf{T}_i,E',\mathbf{k}}(t_{i+1}, E'(t_{i+1}, s_0^*))\right) \\
&= \mathsf{proj}_2\left(\mathsf{REP}_{\mathsf{T}_i,E,\mathbf{k}}\left(t_{i+1}, E(t_{i+1}, s_1^*)\right)\right) \\
&= E^{\mathsf{T}_i,\mathbf{k}}(t_{i+1}, s_1^*) \qquad (19)
\end{aligned}$$

follows, where $\mathsf{proj}_2 : \{0,1\}^m \times \{0,1\}^n \to \{0,1\}^n$ is the projection to the second component defined by $\mathsf{proj}_2(t, y) = y$. □

Define $\mathsf{Bad}_{\mathsf{ent}}$ to be the event that the min-entropy of $\mathcal{R}_\alpha|_{\mathsf{L}_{i+1}}$ is less than ρ for some α, when $\mathcal{A}$ submits the $(i+1)$st construction query. Then, we can show that

$$\Pr[\mathsf{Bad}_{\mathsf{ent}}] \leq \mathsf{Adv}^{\mathsf{cmel}}_{\mathbf{R},\mathcal{P},\rho}(M, q_{L,\max}, Q_{L,\max})$$

holds in both Hyb_i^2 and Hyb_i^3 in the same way as Eq. (13).

In addition, let $\mathsf{Bad} := \mathsf{Bad}_{\mathsf{coll}} \vee \mathsf{Bad}_{\mathsf{ent}}$. The probability of the event Bad is the same between Hyb_i^2 and Hyb_i^3 because the two games are identical until before y_{i+1} is sampled, and whether Bad happens is determined before that.

In both Hyb_i^2 and Hyb_i^3, since s_1^* is chosen uniformly at random, we have

$$\Pr[s_1^* \in X_{\mathsf{bad}}] \leq \frac{|X_{\mathsf{bad}}|}{2^n} \leq \frac{\mu}{2^n}.$$

In addition, by the definition of ψ_1 that $\psi_1(k, x) = \varphi(x \oplus k)$ (where φ is a permutation),

$$\begin{aligned}
&\Pr[\psi_1(k_{\alpha_{i+1}}, x_{i+1}) \in K_{\mathsf{bad}} \mid \neg\, \mathsf{Bad}_{\mathsf{ent}}] \\
&\leq \sum_{z \in K_{\mathsf{bad}}} \Pr_{k_{\alpha_{i+1}} \leftarrow \mathcal{R}_{\alpha_{i+1}}|_{\mathsf{L}_{i+1}}}[k_{\alpha_{i+1}} = x_{i+1} \oplus \varphi^{-1}(z) \mid \neg\, \mathsf{Bad}_{\mathsf{ent}}] \\
&\leq \sum_{z \in K_{\mathsf{bad}}} \frac{1}{2^\rho} \leq \frac{|K_{\mathsf{bad}}|}{2^\rho} \leq \frac{\mu}{2^\rho} \qquad (20)
\end{aligned}$$

also holds. Therefore, we have

$$\begin{aligned}
&\left|\Pr[1 \leftarrow \mathcal{A} \text{ in } \mathsf{Hyb}_i^2] - \Pr[1 \leftarrow \mathcal{A} \text{ in } \mathsf{Hyb}_i^3]\right| \\
&\leq \underbrace{\left|\Pr[1 \leftarrow \mathcal{A} \text{ in } \mathsf{Hyb}_i^2 \mid \neg\mathsf{Bad}] - \Pr[1 \leftarrow \mathcal{A} \text{ in } \mathsf{Hyb}_i^3 \mid \neg\mathsf{Bad}]\right|}_{=0 \text{ by the above claim}} + \Pr[\mathsf{Bad}] \\
&\leq \Pr[s_1^* \in X_{\mathsf{bad}}] + \Pr[\psi_1(k_{\alpha_{i+1}}, x_{i+1}) \in K_{\mathsf{bad}} \mid \neg\, \mathsf{Bad}_{\mathsf{ent}}] + \Pr[\mathsf{Bad}_{\mathsf{ent}}] \\
&\leq \frac{\mu}{2^n} + \frac{\mu}{2^\rho} + \mathsf{Adv}^{\mathsf{cmel}}_{\mathbf{R},\mathcal{P},\rho}(M, q_{L,\max}, Q_{L,\max}) \leq \frac{2\mu}{2^\rho} + \mathsf{Adv}^{\mathsf{cmel}}_{\mathbf{R},\mathcal{P},\rho}(M, q_{L,\max}, Q_{L,\max}).
\end{aligned}$$

[15] Here, we use the query restriction rule inherited from the definition of TCCR security.

□

Bounding the Distance between Hyb_i^3 and Hyb_i^4.

Lemma C.4. $\left|\Pr[1 \leftarrow \mathcal{A} \text{ in } \mathsf{Hyb}_i^3] - \Pr[1 \leftarrow \mathcal{A} \text{ in } \mathsf{Hyb}_i^4]\right| \leq \mu/2^n$.

Proof. The difference between Hyb_i^3 and Hyb_i^4 lies in how the response to the $(i+1)$st query (i.e., the value y_{i+1}) and s_1^* are sampled.

The experiment Hyb_i^3 samples y_{i+1} and s_1^* as follows.

1. Choose s_1^* from $\{0,1\}^n$ uniformly at random.
2. Set y_{i+1} as

$$y_{i+1} \leftarrow \widetilde{y} := \psi_2\left(k_{\alpha_{i+1}}, E^{\mathsf{T}_i,\mathbf{k}}(t_{i+1}, s_1^*), b_{i+1}\right).$$

When considered on its own, the distribution of y_{i+1} is also uniform (because $\psi_2(k_{\alpha_{i+1}}, \cdot, b_{i+1})$ is a permutation). Hence, the distribution of the pair (s_1^*, y_{i+1}) remains unchanged even if we modify the sampling as follows.

1. Choose y_{i+1} from $\{0,1\}^n$ uniformly at random.
2. Set s_1^* as

$$s_1^* \leftarrow \left(E^{\mathsf{T}_i,\mathbf{k}}\right)^{-1}\left(t_{i+1}, \psi_2\left(k_{\alpha_{i+1}}, y_{i+1}, b_{i+1}\right)\right).$$

By contrast, Hyb_i^4 samples y_{i+1} and s_1^* as below.

1. Sample y_{i+1} by

$$y_{i+1} \leftarrow \widetilde{\mathrm{RP}}(\alpha_{i+1}; x_{i+1}, t_{i+1}, b_{i+1}),$$

2. Set s_1^* as

$$s_1^* \leftarrow \left(E^{\mathsf{T}_i,\mathbf{k}}\right)^{-1}\left(t_{i+1}, \psi_2\left(k_{\alpha_{i+1}}, y_{i+1}, b_{i+1}\right)\right).$$

After (equivalently) modifying the sampling in Hyb_i^3, how to set s_1^* is common between the two experiments. Hence, the distinguishing advantage between Hyb_i^3 and Hyb_i^4 is upper bounded by the advantage between the two distributions of y_{i+1}, which is the uniform distribution in Hyb_i^3 while the output distribution of $\widetilde{\mathrm{RP}}(\alpha_{i+1}; \cdot, \cdot, \cdot)$ in Hyb_i^4. Since queries of the form $(*; *, t, *)$ are made to $\widetilde{\mathrm{RP}}$ at most μ times in Hyb_i^4, we have

$$\left|\Pr[1 \leftarrow \mathcal{A} \text{ in } \mathsf{Hyb}_i^3] - \Pr[1 \leftarrow \mathcal{A} \text{ in } \mathsf{Hyb}_i^4]\right| \leq \frac{\mu}{2^n}.$$

□

Bounding the Distance between Hyb_i^4 and Hyb_i^5.

Lemma C.5. *For any $\rho \leq n$, $\left|\Pr[1 \leftarrow \mathcal{A} \text{ in } \mathsf{Hyb}_i^4] - \Pr[1 \leftarrow \mathcal{A} \text{ in } \mathsf{Hyb}_i^5]\right| \leq 9\mu/2^\rho + \mathsf{Adv}^{\mathrm{cmel}}_{\mathbf{R},\mathcal{P},\rho}(M, q_{L,\max}, Q_{L,\max})$ holds.*

Proof. Recall that s_0^* and s_1^* in Hyb_i^4 are defined as

$$s_0^* := \psi_1(k_{\alpha_{i+1}}, x_{i+1}), \quad s_1^* := (E^{\mathsf{T}_i,\mathbf{k}})^{-1}(t_{i+1}, \psi_2(k_{\alpha_{i+1}}, y_{i+1}, b_{i+1})), \tag{21}$$

where $y_{i+1} := \widetilde{\mathrm{RP}}(\alpha_{i+1}; x_{i+1}, t_{i+1}, b_{i+1})$. Let

$$I_{i+1} := \left\{ j \mid 1 \leq j \leq i,\ t_j = t_{i+1} \right\}, \quad K_{\mathsf{bad}} := \left\{ \psi_1(k_{\alpha_j}, x_j) \mid j \in I_{i+1} \text{ and } \alpha_j \neq \alpha_{i+1} \right\},$$
$$X_{\mathsf{bad}} := \left\{ \psi_1(k_{\alpha_j}, x_j) \mid j \in I_{i+1} \right\}, \quad Y_{\mathsf{bad}} := \bigcup_{j \in I_{i+1}} \{E(t_j, \psi_1(k_{\alpha_j}, x_j)), \psi_2(k_{\alpha_j}, y_j, b_j)\}.$$

and define $\mathsf{Bad}_{\mathrm{coll}}$ as the event that $\psi_1(k_{\alpha_{i+1}}, x_{i+1}) \in K_{\mathsf{bad}}$, $s_1^* \in X_{\mathsf{bad}}$, or $\psi_2(k_{\alpha_{i+1}}, y_{i+1}, b_{i+1}) \in Y_{\mathsf{bad}}$. Then, the following claim holds.

Claim. Assume $\mathsf{Bad}_{\mathrm{coll}}$ does not hold in Hyb_i^4, i.e., $\psi_1(k_{\alpha_{i+1}}, x_{i+1}) \notin K_{\mathsf{bad}}$, $s_1^* \notin X_{\mathsf{bad}}$, and $\psi_2(k_{\alpha_{i+1}}, y_{i+1}, b_{i+1}) \notin Y_{\mathsf{bad}}$ in Hyb_i^4. Then, $(E')^{\mathsf{T}_i,\mathbf{k}} = E^{\mathsf{T}_{i+1},\mathbf{k}}$ holds. In particular, Hyb_i^4 is perfectly indistinguishable from Hyb_i^5.

Proof. By using the assumption that $\psi_1(k_{\alpha_{i+1}}, x_{i+1}) \notin K_{\mathsf{bad}}$ and $s_1^* \notin X_{\mathsf{bad}}$, we can prove

$$\mathsf{REP}_{\mathsf{T}_i,E',\mathbf{k}} = \mathsf{REP}_{\mathsf{T}_i,E,\mathbf{k}} \tag{22}$$

in the same way as we showed Eq. (18). In addition, from the assumption $\psi_2(k_{\alpha_{i+1}}, y_{i+1}, b_{i+1}) \notin Y_{\mathsf{bad}}$, we have

$$\begin{aligned} s_1^* &= \left(E^{\mathsf{T}_i,\mathbf{k}}\right)^{-1}\left(t_{i+1}, \psi_2(k_{\alpha_{i+1}}, y_{i+1}, b_{i+1})\right) \\ &= E^{-1}\left(t_{i+1}, \psi_2(k_{\alpha_{i+1}}, y_{i+1}, b_{i+1})\right), \end{aligned}$$

which implies that

$$\mathsf{swap}^{t_{i+1}}_{s_0^*,s_1^*} = \mathsf{swap}^{t_{i+1}}_{\psi_1(k_{\alpha_{i+1}},x_{i+1}),E^{-1}\left(t_{i+1},\psi_2(k_{\alpha_{i+1}},y_{i+1},b_{i+1})\right)}. \tag{23}$$

Hence, for any (t, x),

$$\begin{aligned} (E')^{\mathsf{T}_i,\mathbf{k}}(t,x) &= \mathsf{proj}_2\left(\mathsf{REP}_{\mathsf{T}_i,E',\mathbf{k}}\left(t, E'(t,x)\right)\right) \\ &\overset{(a)}{=} \mathsf{proj}_2\left(\mathsf{REP}_{\mathsf{T}_i,E',\mathbf{k}}\left(t, E\left(\mathsf{swap}^{t_{i+1}}_{s_0^*,s_1^*}(t,x)\right)\right)\right) \\ &\overset{(b)}{=} \mathsf{proj}_2\left(\mathsf{REP}_{\mathsf{T}_i,E,\mathbf{k}}\left(t, E\left(\mathsf{swap}^{t_{i+1}}_{s_0^*,s_1^*}(t,x)\right)\right)\right) \\ &\overset{(c)}{=} \mathsf{proj}_2\left(\mathsf{REP}_{\mathsf{T}_i,E,\mathbf{k}}\left(\mathsf{swap}^{t_{i+1}}_{s_0^{i+1},s_1^{i+1}}(t, E(t,x))\right)\right) \\ &= \mathsf{proj}_2\left(\mathsf{REP}_{\mathsf{T}_{i+1},E,\mathbf{k}}(t, E(t,x))\right) \\ &= E^{\mathsf{T}_{i+1},\mathbf{k}}(t,x) \end{aligned}$$

holds, where $\mathsf{proj}_2 : \{0,1\}^m \times \{0,1\}^n \to \{0,1\}^n$ is the projection to the second component defined by $\mathsf{proj}_2(t,y) = y$, (a) follows from the definition of E', (b) follows from Eq. (22), and (c) follows from Eq. (23). Therefore, $(E')^{\mathsf{T}_i,\mathbf{k}} = E^{\mathsf{T}_{i+1},\mathbf{k}}$ follows. □

Define $\mathsf{Bad}_{\mathsf{ent}}$ to be the event that the min-entropy of $\mathcal{R}_\alpha|_{\mathsf{L}_{i+1}}$ is less than ρ for some α, when $\mathcal{A}$ receives the response to the $(i+1)$st construction query. Then, we can show that

$$\Pr\left[\mathsf{Bad}_{\mathsf{ent}}\right] \le \mathsf{Adv}^{\mathrm{cmel}}_{\mathbf{R},\mathcal{P},\rho}(M, q_{L,\max}, Q_{L,\max})$$

holds in both Hyb_i^4 and Hyb_i^5 in the same way as Eq. (13).

The variables s_0^* and s_1^* do not appear in the definition of Hyb_i^5, but we can define them as in Eq. (21) and can argue whether $\mathsf{Bad}_{\mathsf{coll}}$ happens in Hyb_i^5 or not. Let $\mathsf{Bad} := \mathsf{Bad}_{\mathsf{ent}} \vee \mathsf{Bad}_{\mathsf{coll}}$. Then, the probability of Bad is the same between Hyb_i^4 and Hyb_i^5 because the two games are identical until the processing of the $(i+1)$-th classical query has been completed, and whether Bad happens is determined once y_{i+1} is sampled and returned to the adversary.

y_{i+1} is set as $y_{i+1} := \widetilde{\mathrm{RP}}(\alpha_{i+1}; x_{i+1}, t_{i+1}, b_{i+1})$ in both Hyb_i^4 and Hyb_i^5, where $\widetilde{\mathrm{RP}}$ is an ideally random tweakable permutation. Since the number of queries of the form $(\alpha; *, t, *)$ made to the classical oracle of $\widetilde{\mathrm{RP}}$ is at most μ for each t, we have

$$\begin{aligned} \Pr\left[\psi_2(k_{\alpha_{i+1}}, y_{i+1}, b_{i+1}) \in Y_{\mathsf{bad}}\right] &= \Pr\left[y_{i+1} \oplus \sigma(k_{\alpha_{i+1}}) \oplus b_{i+1} \cdot k_{\alpha_{i+1}} \in Y_{\mathsf{bad}}\right] \\ &\le \frac{|Y_{\mathsf{bad}}|}{(2^n - \mu)} \le \frac{4\mu}{2^n} \end{aligned} \tag{24}$$

(because $(\mu \le)\ q_C \le 2^{n-1}$ is assumed), and similarly

$$\Pr\left[s_1^* \in X_{\mathsf{bad}}\right] \le \frac{4\mu}{2^n}$$

holds (because s_1^* is defined from y_{i+1} via Eq. (21) and $(E^{\mathsf{T}_i,\mathbf{k}})^{-1}(t_{i+1}, \cdot)$ is a permutation). Furthermore, we can show

$$\Pr\left[\psi_1(k_{\alpha_{i+1}}, x_{i+1}) \in K_{\mathsf{bad}} \,\middle|\, \neg\, \mathsf{Bad}_{\mathsf{ent}}\right] \le \frac{\mu}{2^\rho}$$

in the same way as Eq. (20). Therefore,

$$\begin{aligned}
&\left|\Pr[1 \leftarrow \mathcal{A} \text{ in } \mathsf{Hyb}_i^4] - \Pr[1 \leftarrow \mathcal{A} \text{ in } \mathsf{Hyb}_i^5]\right| \\
&\leq \underbrace{\left|\Pr\left[1 \leftarrow \mathcal{A} \text{ in } \mathsf{Hyb}_i^4 \mid \neg\mathsf{Bad}\right] - \Pr\left[1 \leftarrow \mathcal{A} \text{ in } \mathsf{Hyb}_i^5 \mid \neg\mathsf{Bad}\right]\right|}_{=0 \text{ by the above claim}} + \Pr\left[\mathsf{Bad}\right] \\
&\leq \Pr\left[s_1^* \in X_{\mathsf{bad}}\right] + \Pr\left[\psi_2(k_{\alpha_{i+1}}, y_{i+1}, b_{i+1}) \in Y_{\mathsf{bad}}\right] \\
&\quad + \Pr\left[\psi_1(k_{\alpha_{i+1}}, x_{i+1}) \in K_{\mathsf{bad}} \,\middle|\, \neg\, \mathsf{Bad}_{\mathsf{ent}}\right] + \Pr\left[\mathsf{Bad}_{\mathsf{ent}}\right] \\
&\leq \frac{4\mu}{2^n} + \frac{4\mu}{2^n} + \frac{\mu}{2^\rho} + \mathsf{Adv}^{\mathrm{cmel}}_{\mathbf{R},\mathcal{P},\rho}(M, q_{L,\max}, Q_{L,\max}) \leq \frac{9\mu}{2^\rho} + \mathsf{Adv}^{\mathrm{cmel}}_{\mathbf{R},\mathcal{P},\rho}(M, q_{L,\max}, Q_{L,\max})
\end{aligned}$$

follows. □

Bounding the Distance between Hyb_i^5 and Hyb_{i+1}^1.

Lemma C.6. *For $i \leq q_C - 2$ (resp., $i = q_C - 1$), let $q_{E,i+1}$ be the expected number of the quantum primitive queries $\mathcal{A}$ makes between the $(i+1)$st and $(i+2)$nd construction queries (resp., $\mathcal{A}$ makes after the last construction query) in the ideal world. If $\mathsf{Adv}^{\mathrm{cmel}}_{\mathbf{R},\mathcal{P},\rho}(M, q_{L,\max}, Q_{L,\max}) < 1$, then we have*

$$\begin{aligned}
&\left|\Pr[1 \leftarrow \mathcal{A} \text{ in } \mathsf{Hyb}_i^5] - \Pr[1 \leftarrow \mathcal{A} \text{ in } \mathsf{Hyb}_{i+1}^1]\right| \\
&\qquad \leq \frac{2q_{E,i+1}\sqrt{2\mu/2^\rho}}{1 - \mathsf{Adv}^{\mathrm{cmel}}_{\mathbf{R},\mathcal{P},\rho}(M, q_{L,\max}, Q_{L,\max})} + (q_{L,\max}+1)\mathsf{Adv}^{\mathrm{cmel}}_{\mathbf{R},\mathcal{P},\rho}(M, q_{L,\max}, Q_{L,\max})
\end{aligned} \tag{25}$$

To prove this lemma, we introduce additional hybrids $\mathsf{Hyb}_i^{5\text{-}j}$ for $j = 0, 1, 2, \ldots, q_{L,\max}+1$ between Hyb_i^5 and Hyb_{i+1}^1. The difference of $\mathsf{Hyb}_i^{5\text{-}j}$ from Hyb_i^5 is when to resample the keys and change the primitive oracle. For $i \leq q_C - 2$ (resp., $i = q_C - 1$), let $\mathsf{NumL}_{i+1}^{<j}$ be the event that $\mathcal{A}$ does not make j or more leaking queries between the $(i+1)$st and $(i+2)$nd construction queries (resp., after the last construction query). (This event $\mathsf{NumL}_{i+1}^{<j}$ will be used solely to make explanation on case-by-case analysis simpler.)

$\underline{\mathsf{Hyb}_i^{5\text{-}j}}$:

1. The keys $\mathbf{k} = (k_1, \ldots, k_u)$ are sampled as $k_\alpha \leftarrow \mathcal{R}_\alpha$ for each α at the beginning. The classical oracle of $\widetilde{\mathrm{RP}}$, the leaking oracle $\mathcal{L}_{\mathbf{k}}$, and the quantum oracle of E (and its inverse) are given to the adversary $\mathcal{A}$.
2. If $j = 0$: When $\mathcal{A}$ receives the response to the $(i+1)$st construction query, let $\mathsf{L}_{i+1}^{(j)}$ be the transcript of the leaking queries made so far (and the responses), and then resample the keys $\mathbf{k} = (k_1, \ldots, k_u)$ as $k_\alpha \leftarrow \mathcal{R}_\alpha|_{\mathsf{L}_{i+1}^{(j)}}$ for each α. Go to Step 3.
   If $j > 0$ and $\mathsf{NumL}_{i+1}^{<j}$ does not happen: When $\mathcal{A}$ receives the response to the jth leaking query, counting from after receiving the $(i+1)$st construction query, define $\mathsf{L}_{i+1}^{(j)}$ and resample the keys $\mathbf{k}$ similarly, and then go to Step 3.
   If $j > 0$ and $\mathsf{NumL}_{i+1}^{<j}$ happens, and $i \leq q_C - 2$: When $\mathcal{A}$ submits the $(i+2)$nd construction query, define $\mathsf{L}_{i+1}^{(j)}$ and resample $\mathbf{k}$ similarly, and then go to Step 3.
   If $j > 0$ and $\mathsf{NumL}_{i+1}^{<j}$ happens, and $i = q_C - 1$: In this case, just keep running $\mathcal{A}$ until it returns the final output (without changing the oracles).
3. Replace the construction oracle (resp., the quantum primitive oracle $|E^{\pm}\rangle$) with $\overline{\mathsf{TBC}}_{\mathbf{k}}^{E^{\mathsf{T}_{i+1},\mathbf{k}}}$ (resp., the quantum oracle of $E^{\mathsf{T}_{i+1},\mathbf{k}}$ and its inverse). Continue running $\mathcal{A}$ until it returns the final output.

The following properties follow by definition:

- $\mathsf{Hyb}_i^{5\text{-}0} = \mathsf{Hyb}_i^5$.
- If $\mathsf{NumL}_{i+1}^{<j}$ happens, then the responses from the oracles completely match between $\mathsf{Hyb}_i^{5\text{-}j}$ and Hyb_{i+1}^1.
- If $j = q_{L,\max}+1$, then $\mathsf{NumL}_{i+1}^{<j}$ always happens (and thus $\mathsf{Hyb}_i^{5\text{-}(q_{L,\max}+1)} = \mathsf{Hyb}_{i+1}^1$).

Moreover, for $i \leq q_C - 2$ (resp., $i = q_C - 1$), define $q_{E,i+1}^{(j+1)}$ as the expected number of the quantum primitive queries that $\mathcal{A}$ makes in the ideal world after the jth leaking query (counting from after receiving the response

to the $(i+1)$st construction query)[16] and before the next leaking or construction queries (resp., before the next leaking query or $\mathcal{A}$ returns the final output).

Then, the following lemma holds.

Lemma C.7. *Let $i < q_C$ and $j \le q_{L,\max}$. If $\mathsf{Adv}^{\mathrm{cmel}}_{\mathbf{R},\mathcal{P},\rho}(M, q_{L,\max}, Q_{L,\max}) < 1$, then we have*

$$\left|\Pr[1 \leftarrow \mathcal{A} \text{ in } \mathsf{Hyb}_i^{5\text{-}j}] - \Pr[1 \leftarrow \mathcal{A} \text{ in } \mathsf{Hyb}_i^{5\text{-}(j+1)}]\right| \le \frac{2q_{E,i+1}^{(j+1)}\sqrt{2\mu/2^\rho}}{1 - \mathsf{Adv}^{\mathrm{cmel}}_{\mathbf{R},\mathcal{P},\rho}(M, q_{L,\max}, Q_{L,\max})} + \mathsf{Adv}^{\mathrm{cmel}}_{\mathbf{R},\mathcal{P},\rho}(M, q_{L,\max}, Q_{L,\max}) \tag{26}$$

To prove this lemma, we need the arbitrary reprogramming lemma below.

Lemma C.8 (Arbitrary reprogramming lemma [ABKM22]). *Let $\mathcal{X}$ and $\mathcal{Y}$ be some non-empty finite sets. Consider the following three-stage experiment with a (quantum) distinguisher $\mathcal{D}$.*

1. *After performing some computation, $\mathcal{D}$ outputs the descriptions of a function $F_0 : \mathcal{X} \to \mathcal{Y}$ and a randomized algorithm $\mathcal{B}$. Here, $\mathcal{B}$ samples and outputs a function $F_1 : \mathcal{X} \to \mathcal{Y}$.*
2. *Run $\mathcal{B}$ to sample F_1. Choose a bit r uniformly at random. Resume running $\mathcal{D}$, giving quantum oracle access to F_r.*
3. *$\mathcal{D}$ is given the randomness used in $\mathcal{B}$ to sample F_1, performs some additional computation* without *oracle access to F_r, and finally outputs a bit r'.*

If the expected number of queries by $\mathcal{D}$ in the second stage is at most q when $r = 0$, then

$$|\Pr[1 \leftarrow \mathcal{D} \mid r = 1] - \Pr[1 \leftarrow \mathcal{D} \mid r = 0]| \le 2q \cdot \sqrt{\varepsilon}$$

holds, where $\varepsilon := \sup_{(F_0,\mathcal{B})} \max_{x \in \mathcal{X}} \Pr[F_0(x) \neq F_1(x)]$. Here, the supremum is taken over all the possible pairs $(F_0, \mathcal{B})$ output by $\mathcal{D}$ in the first stage, and the probability is taken over the randomness used in $\mathcal{B}$ to sample F_1 for each fixed $(F_0, \mathcal{B})$.

The proof of Lemma C.7 is as follows.

Proof (of Lemma C.7). Below, we prove the claim when $i \le q_C - 2$ and $j > 0$. (See also Tables 5 and 6 for comparison between $\mathsf{Hyb}_i^{5\text{-}j}$ and $\mathsf{Hyb}_i^{5\text{-}(j+1)}$ in this case.) The proof for $i = q_C - 1$ or $j = 0$ is analogous.

Let Bad be the event that the min-entropy of $\mathcal{R}_\alpha|_{\mathsf{L}_{i+1}^{(j)}}$ is less than ρ for some α, when $\mathcal{A}$ receives the response to jth leaking query, counting from after receiving the $(i+1)$st construction query, or $\mathcal{A}$ submits the $(i+2)$nd construction query. Then, we can show

$$\Pr[\mathsf{Bad}] \le \mathsf{Adv}^{\mathrm{cmel}}_{\mathbf{R},\mathcal{P},\rho}(M, q_{L,\max}, Q_{L,\max})$$

holds in both $\mathsf{Hyb}_i^{5\text{-}j}$ and $\mathsf{Hyb}_i^{5\text{-}(j+1)}$, in the same way as Eq. (13).

Let $\mathcal{D}$ be a distinguisher for the arbitrary reprogramming lemma (Lemma C.8) defined as follows.

1-1. $\mathcal{D}$ samples an ideal cipher E, a tweakable random permutation $\widetilde{\mathrm{RP}}$, and keys $\mathbf{k} = (k_1, \dots, k_u)$ as $k_\alpha \leftarrow \mathcal{R}_\alpha$ by itself.

1-2. $\mathcal{D}$ runs the adversary $\mathcal{A}$. The construction oracle, the leaking oracle, and the primitive oracle for $\mathcal{A}$ are simulated with $\widetilde{\mathrm{RP}}$, $\mathcal{L}_{\mathbf{k}}$, and $|E^{\pm}\rangle$, respectively.

1-3. If $\mathsf{NumL}_{i+1}^{<j}$ does not happen (resp., if $\mathsf{NumL}_{i+1}^{<j}$ happens), when $\mathcal{A}$ receives the response to the jth leaking query, counting from after receiving the $(i+1)$st construction query (resp., when $\mathcal{A}$ submits the $(i+2)$nd construction query), the distinguisher $\mathcal{D}$ computes the conditional distributions $\mathcal{R}_\alpha|_{\mathsf{L}_{i+1}^{(j)}}$ for $1 \le \alpha \le u$ to check if Bad has happened or not, and then outputs the description of a function F_0 and an algorithm $\mathcal{B}$ (to sample another function F_1) as follows.
   - Define $F_0 : \{0,1\} \times \{0,1\}^m \times \{0,1\}^n \to \{0,1\}^n$ by $F_0(a,t,x) := E^{1-2a}(t,x)$.
   - $\mathcal{B}$ first resamples the keys $\mathbf{k} = (k_1, \dots, k_u)$ as $k_\alpha \leftarrow \mathcal{R}_\alpha|_{\mathsf{L}_{i+1}^{(j)}}$ for $\alpha = 1, \dots, u$. If neither Bad nor $\mathsf{NumL}_{i+1}^{<j}$ has happened, then $\mathcal{B}$ defines F_1 as $F_1(a,t,x) := \left(E^{\mathsf{T}_{i+1},\mathbf{k}}\right)^{1-2a}(t,x)$. If Bad or $\mathsf{NumL}_{i+1}^{<j}$ has happened, then $\mathcal{B}$ sets F_1 just as $F_1 := F_0$.

[16] When $j = 0$, we regard that it is the $(i+1)$st construction query, by abuse of notation. In addition, when the aforementioned jth leaking query is not defined, then the number of quantum primitive queries is regarded as 0.

2-1. $\mathcal{B}$ is run to sample F_1, and a random bit r is chosen. $\mathcal{D}$ is given quantum oracle access to F_r.
2-2. $\mathcal{D}$ continues running $\mathcal{A}$, simulating the primitive oracle by using F_r (for example, if $r = 0$, $F_0(0,\cdot,\cdot)$ and $F_0(1,\cdot,\cdot)$ are used to simulate E and E^{-1}, respectively). When $\mathcal{A}$ makes the next classical query (either leaking query or construction query), go to the next step.
3-1. $\mathcal{D}$ receives the randomness to sample F_1, i.e., the keys $\mathbf{k} = (k_1, \dots, k_u)$.
3-2. $\mathcal{D}$ updates the key used by the leaking oracle to the resampled $\mathbf{k}$, and replaces the construction oracle and the primitive oracle for $\mathcal{A}$ with $\overline{\mathsf{TBC}}_{\mathbf{k}}^{E^{\mathsf{T}_{i+1},\mathbf{k}}}$ and $E^{\mathsf{T}_{i+1},\mathbf{k}}$ (and its inverse), respectively, and continues running $\mathcal{A}$. Thereafter, the queries (including the one $\mathcal{A}$ made at the end of Step 2-2) are responded with the updated oracles.
3-3. When $\mathcal{A}$ returns an output bit, $\mathcal{D}$ returns it as its own output.

Then, the following claim holds.

Claim. It holds that

$$\Pr\left[1 \leftarrow \mathcal{D} \mid r = 0, \neg\mathsf{Bad}\right] = \Pr\left[1 \leftarrow \mathcal{A} \text{ in } \mathsf{Hyb}_i^{5\text{-}(j+1)} \mid \neg\mathsf{Bad}\right], \tag{27}$$

$$\Pr\left[1 \leftarrow \mathcal{D} \mid r = 1, \neg\mathsf{Bad}\right] = \Pr\left[1 \leftarrow \mathcal{A} \text{ in } \mathsf{Hyb}_i^{5\text{-}j} \mid \neg\mathsf{Bad}\right], \tag{28}$$

$$\Pr\left[1 \leftarrow \mathcal{D} \wedge \mathsf{Bad} \mid r = 0\right] = \Pr\left[1 \leftarrow \mathcal{D} \wedge \mathsf{Bad} \mid r = 1\right]. \tag{29}$$

Proof. Eq. (28) is obvious because the experiment $\mathsf{Hyb}_i^{5\text{-}j}$ is perfectly simulated for $\mathcal{A}$ when $r = 1$ as long as Bad does not happen. In addition, Eq. (29) holds because $F_1 = F_0$ if Bad happens, and thus the oracles simulated for $\mathcal{A}$ are the same between $r = 0$ and $r = 1$.

In what follows, we show Eq. (27). (See also Tables 5 and 6.) Note that, up to the point where $\mathcal{D}$ outputs F_0 and $\mathcal{B}$, the simulation is identical for $r = 0$ and $r = 1$, and whether the event Bad occurs has already been determined.

When $r = 0$, the primitive oracle available to $\mathcal{A}$ until the next classical query is implemented by using F_0, and hence is simply $|E^{\pm}\rangle$. This agrees with $\mathsf{Hyb}_i^{5\text{-}(j+1)}$. There is, however, a difference in the timing of key resampling. In the experiment simulated by $\mathcal{D}$, $\mathcal{B}$ re-samples the keys as $k_\alpha \leftarrow \mathcal{R}_\alpha|_{\mathsf{L}_{i+1}^{(j)}}$ for $\alpha = 1, \dots, u$ before the next classical query, whereas in $\mathsf{Hyb}_i^{5\text{-}(j+1)}$, if the next classical query is a leakage query, the keys are resampled after its response is returned. Nevertheless, this difference does not affect the joint distribution seen by $\mathcal{A}$: Indeed, suppose that the next classical query is a leakage query $(\alpha; P)$, let ℓ denote its response, and let

$$\mathsf{L}_{i+1}^{(j+1)} := \mathsf{L}_{i+1}^{(j)} \cup \{((\alpha; P), \ell)\}$$

be the resulting leakage transcript. By the definition of the conditional key distributions, conditioned additionally on the response ℓ, a key sampled from $\mathcal{R}_\alpha|_{\mathsf{L}_{i+1}^{(j)}}$ is distributed exactly according to $\mathcal{R}_\alpha|_{\mathsf{L}_{i+1}^{(j+1)}}$. Moreover, before this next classical query, the resampled keys are not used by any oracle visible to $\mathcal{A}$ when $r = 0$, since the primitive oracle is $|E^{\pm}\rangle$. Thus, whether the keys are sampled immediately before or immediately after the leakage response is unobservable to $\mathcal{A}$. Therefore, Eq. (27) holds. (If the next classical query is not a leakage query but the $(i+2)$-nd construction query, the claim is immediate: the keys sampled by $\mathcal{B}$ have distribution $\mathcal{R}_\alpha|_{\mathsf{L}_{i+1}^{(j)}}$ and have not been used while the primitive oracle is E, so sampling them earlier is equivalent to the resampling performed by $\mathsf{Hyb}_i^{5\text{-}(j+1)}$ when that construction query is submitted.) □

Moreover, the value of $F_1(a,t,x) = \left(E^{\mathsf{T}_{i+1},\mathbf{k}}\right)^{1-2a}(t,x)$ is different from $F_0(a,t,x) = E^{1-2a}(t,x)$ only if Bad does not happen and $(t,x) = \left(t_h, z_h^{a,0}\right)$ or $(t,x) = \left(t_h, z_h^{a,1}\right)$ for some $h \le i+1$, where $z_h^{0,0} := \psi_1(k_{\alpha_h}, x_h)$, $z_h^{0,1} := E^{-1}\left(t_h, \psi_2(k_{\alpha_h}, y_h, b_h)\right)$, $z_h^{1,0} := E\left(t_h, \psi_1(k_{\alpha_h}, x_h)\right)$, $z_h^{1,1} := \psi_2\left(k_{\alpha_h}, y_h, b_h\right)$. Thus, for any (a,t,x), we have

$$\begin{aligned}
&\Pr\left[F_0(a,t,x) \neq F_1(a,t,x)\right] \le \Pr\left[F_0(a,t,x) \neq F_1(a,t,x) \mid \neg\mathsf{Bad}\right] \\
&\le \sum_{h=1,\dots,i+1} \left(\Pr\left[(t,x) = \left(t_h, z_h^{a,0}\right) \middle| \neg\,\mathsf{Bad}\right] + \Pr\left[(t,x) = \left(t_h, z_h^{a,1}\right) \middle| \neg\,\mathsf{Bad}\right]\right) \\
&\le \sum_{\substack{h=1,\dots,i+1 \\ \text{s.t. } t=t_h}} \left(\Pr\left[x = z_h^{a,0} \middle| \neg\,\mathsf{Bad}\right] + \Pr\left[x = z_h^{a,1} \middle| \neg\,\mathsf{Bad}\right]\right) \\
&\le \sum_{\substack{h=1,\dots,i+1 \\ \text{s.t. } t=t_h}} \left(\frac{1}{2^\rho} + \frac{1}{2^\rho}\right) \le \frac{2\mu}{2^\rho},
\end{aligned} \tag{30}$$

where the second last inequality holds because the min-entropy of k_α is at least ρ for each α when resampled if Bad does not happen (in addition that φ is a permutation and σ is a linear orthomorphism), and the last inequality holds because $\mathcal{A}$ makes at most μ queries of the form $(*, t, *)$ for any t across the classical queries.

Therefore, it follows that

$$
\begin{aligned}
&\left|\Pr\left[1 \leftarrow \mathcal{A} \text{ in } \mathsf{Hyb}_i^{5\text{-}(j+1)}\right] - \Pr\left[1 \leftarrow \mathcal{A} \text{ in } \mathsf{Hyb}_i^{5\text{-}j}\right]\right| \\
&\leq \Bigg|\underbrace{\Pr\left[1 \leftarrow \mathcal{A} \text{ in } \mathsf{Hyb}_i^{5\text{-}(j+1)} \mid \neg\mathsf{Bad}\right]}_{=\Pr[1\leftarrow\mathcal{D}|r=0,\neg\mathsf{Bad}] \text{ by Eq. (27)}} - \underbrace{\Pr\left[1 \leftarrow \mathcal{A} \text{ in } \mathsf{Hyb}_i^{5\text{-}j} \mid \neg\mathsf{Bad}\right]}_{=\Pr[1\leftarrow\mathcal{D}|r=1,\neg\mathsf{Bad}] \text{ by Eq. (28)}}\Bigg| + \Pr\left[\mathsf{Bad}\right] \\
&= \left|\Pr\left[1 \leftarrow \mathcal{D} \mid r = 0, \neg\mathsf{Bad}\right] - \Pr\left[1 \leftarrow \mathcal{D} \mid r = 1, \neg\mathsf{Bad}\right]\right| + \Pr\left[\mathsf{Bad}\right] \\
&= \left|\Pr\left[1 \leftarrow \mathcal{D} \wedge \neg\mathsf{Bad} \mid r = 0\right] - \Pr\left[1 \leftarrow \mathcal{D} \wedge \neg\mathsf{Bad} \mid r = 1\right]\right| \cdot \frac{1}{\Pr[\neg\mathsf{Bad}]} + \Pr\left[\mathsf{Bad}\right] \\
&\leq \underbrace{\left|\Pr\left[1 \leftarrow \mathcal{D} \mid r = 0\right] - \Pr\left[1 \leftarrow \mathcal{D} \mid r = 1\right]\right|}_{\leq 2q_{E,i+1}^{(j+1)}\sqrt{2\mu/2^\rho} \text{ by Lemma C.8 and Eq. (30)}} \cdot \frac{1}{\Pr[\neg\mathsf{Bad}]} \\
&\quad + \underbrace{\left|\Pr\left[1 \leftarrow \mathcal{D} \wedge \mathsf{Bad} \mid r = 0\right] - \Pr\left[1 \leftarrow \mathcal{D} \wedge \mathsf{Bad} \mid r = 1\right]\right|}_{=0 \text{ by Eq. (29)}} \cdot \frac{1}{\Pr[\neg\mathsf{Bad}]} + \Pr\left[\mathsf{Bad}\right] \\
&\leq \frac{2q_{E,i+1}^{(j+1)}\sqrt{2\mu/2^\rho}}{1 - \Pr[\mathsf{Bad}]} + \Pr[\mathsf{Bad}] \\
&\leq \frac{2q_{E,i+1}^{(j+1)}\sqrt{2\mu/2^\rho}}{1 - \mathsf{Adv}^{\mathrm{cmel}}_{\mathbf{R},\mathcal{P},\rho}(M, q_{L,\max}, Q_{L,\max})} + \mathsf{Adv}^{\mathrm{cmel}}_{\mathbf{R},\mathcal{P},\rho}(M, q_{L,\max}, Q_{L,\max}).
\end{aligned}
$$

□

Finally, Lemma C.6 is proven as follows.

Proof (of Lemma C.6). Since $\mathsf{Hyb}_i^{5\text{-}0} = \mathsf{Hyb}_i^5$ and $\mathsf{Hyb}_i^{5\text{-}(q_{L,\max}+1)} = \mathsf{Hyb}_{i+1}^1$, we have

$$
\begin{aligned}
&\left|\Pr[1 \leftarrow \mathcal{A} \text{ in } \mathsf{Hyb}_i^5] - \Pr[1 \leftarrow \mathcal{A} \text{ in } \mathsf{Hyb}_{i+1}^1]\right| \\
&\leq \sum_{0\leq j\leq q_{L,\max}} \left|\Pr[1 \leftarrow \mathcal{A} \text{ in } \mathsf{Hyb}_i^{5\text{-}j}] - \Pr[1 \leftarrow \mathcal{A} \text{ in } \mathsf{Hyb}_i^{5\text{-}(j+1)}]\right| \\
&\leq \sum_{0\leq j\leq q_{L,\max}} \left(\frac{2q_{E,i+1}^{(j+1)}\sqrt{2\mu/2^\rho}}{1 - \mathsf{Adv}^{\mathrm{cmel}}_{\mathbf{R},\mathcal{P},\rho}(M, q_{L,\max}, Q_{L,\max})} + \mathsf{Adv}^{\mathrm{cmel}}_{\mathbf{R},\mathcal{P},\rho}(M, q_{L,\max}, Q_{L,\max})\right) \\
&\leq \frac{2q_{E,i+1}\sqrt{2\mu/2^\rho}}{1 - \mathsf{Adv}^{\mathrm{cmel}}_{\mathbf{R},\mathcal{P},\rho}(M, q_{L,\max}, Q_{L,\max})} + \left(q_{L,\max}+1\right)\mathsf{Adv}^{\mathrm{cmel}}_{\mathbf{R},\mathcal{P},\rho}(M, q_{L,\max}, Q_{L,\max}) \qquad (31)
\end{aligned}
$$

follows from Lemma C.7 (note that $q_{E,i+1}^{(1)} + \cdots + q_{E,i+1}^{(q_{L,\max}+1)} = q_{E,i+1}$). □

Finishing the Proof of Proposition 5.2. The claim of Proposition 5.2 immediately follows from Lemmas C.1–C.6 because

$$\begin{aligned}
&\mathsf{Adv}^{\mathrm{mTPRPL}^*}_{\mathbf{TBC},\mathbf{R},\mathcal{P}}(\mathcal{A}) \\
&= \left|\Pr[1 \leftarrow \mathcal{A} \text{ in } \mathsf{Hyb}^1_0] - \Pr[1 \leftarrow \mathcal{A} \text{ in } \mathsf{Hyb}^1_{q_C}]\right| \\
&\le \sum_{0\le i\le q_C-1, 1\le j\le 4} \left|\Pr[1 \leftarrow \mathcal{A} \text{ in } \mathsf{Hyb}^j_i] - \Pr[1 \leftarrow \mathcal{A} \text{ in } \mathsf{Hyb}^{j+1}_i]\right| \\
&\quad + \sum_{0\le i\le q_C-1} \left|\Pr[1 \leftarrow \mathcal{A} \text{ in } \mathsf{Hyb}^5_i] - \Pr[1 \leftarrow \mathcal{A} \text{ in } \mathsf{Hyb}^1_{i+1}]\right| \\
&\le \frac{1}{1-\mathsf{Adv}^{\mathrm{cmel}}_{\mathbf{R},\mathcal{P},\rho}(M, q_{L,\max}, Q_{L,\max})} \cdot \left(4\sqrt{\frac{(q_C)^2 q_E}{2^\rho}} + 2\sqrt{\frac{2\mu(q_E)^2}{2^\rho}}\right) \\
&\quad + \frac{q_C \cdot \mathsf{Adv}^{\mathrm{cmel}}_{\mathbf{R},\mathcal{P},\rho}(M, q_{L,\max}, Q_{L,\max})}{1-\mathsf{Adv}^{\mathrm{cmel}}_{\mathbf{R},\mathcal{P},\rho}(M, q_{L,\max}, Q_{L,\max})} \\
&\quad + \frac{12 q_C \mu}{2^\rho} + q_C(q_{L,\max}+4) \cdot \mathsf{Adv}^{\mathrm{cmel}}_{\mathbf{R},\mathcal{P},\rho}(M, q_{L,\max}, Q_{L,\max}). \qquad (32)
\end{aligned}$$

□

D Proof of Proposition 3.1

To prove the proposition, we use the following lemma.

Lemma D.1. *Let S be a d-dimensional affine subspace of $\{0,1\}^n$ and P be a linear predicate. That is, there is a d-dimensional $\mathbb{F}_2$-vector space $V \subset \{0,1\}^n$ and $c \in \{0,1\}^n$ such that $S = V \oplus c$, and there exists some $a \in \{0,1\}^n$ and $b \in \{0,1\}$ such that $P(k) = 1$ iff $\langle a, k\rangle = b$. Then, the following properties hold.*

1. *If $\langle a, v\rangle = 0$ for all $v \in V$, then $P(k) = 0$ for all $k \in S$ or $P(k) = 1$ for all $k \in S$.*
2. *If $\langle a, v_*\rangle = 1$ for some $v_* \in V$, then both $S_0 := \{k \in S : P(k) = 0\}$ and $S_1 := \{k \in S : P(k) = 1\}$ are $(d-1)$-dimensional affine subspaces.*

The lemma can be shown by straightforward and basic discussions of linear algebra, but we give a proof below for completeness.

Proof (of Lemma D.1). If $\langle a, v\rangle = 0$ for all $v \in V$, then $\langle a, k\rangle = \langle a, c\rangle$ for all $k \in S = V \oplus c$. Hence, the first claim holds.

Next, we prove the second claim. Let $V_0 := \{v \in V : \langle a, v\rangle = 0\}$ and $V_1 := \{v \in V : \langle a, v\rangle = 1\}$. Then, $V = V_0 \cup V_1$ and $V_0 \cap V_1 = \emptyset$ hold, and V_0 is a $\mathbb{F}_2$ linear subspace of V. Define an $\mathbb{F}_2$-linear map $F : V \to \mathbb{F}_2$ by $F(v) := \langle a, v\rangle$. Then, $\mathrm{Im}(F) = \mathbb{F}_2$ because there is a $v_* \in V$ such that $\langle a, v_*\rangle = 1$ by assumption. Hence, the dimension of V_0 is equal to $\dim(\mathrm{Ker}(F)) = \dim(V) - \dim(\mathrm{Im}(F)) = d - 1$. Moreover, $V_0 \oplus v_* \subset V_1$ holds because $\langle a, v \oplus v_*\rangle = 1$ for any $v \in V_0$. This implies that $2^{d-1} = |V_0 \oplus v_*| \le |V_1| = |V| - |V_0| = 2^{d-1}$, and thus we have $V_1 = V_0 \oplus v_*$. Furthermore, we have $S_0 = V_{\langle a,c\rangle \oplus b \oplus 1} \oplus c$ and $S_1 = V_{\langle a,c\rangle \oplus b} \oplus c$ by definition of S_0 and S_1. Therefore, the second claim also follows. □

Next, we prove Proposition 3.1.

Proof (of Proposition 3.1). Since we care only about the number of queries as the computational resource of adversaries (time complexity, space complexity, etc., are ignored), we can assume that adversaries are classical deterministic algorithms, without loss of generality.

First, we prove the proposition when $u = 1$. Since $\mathcal{R}_1$ is the uniform distribution over an affine space of dimension $\ge d$, the min-entropy of $\mathcal{R}_1$ is at least d. Due to Lemma D.1, each query to the leaking oracle $\mathcal{L}_{\mathbf{k}}$ either (1) does not provide any new information about the key, or (2) decreases the entropy of the key exactly by 1. Without loss of generality, we can assume that adversaries make queries of only the type of (2), and the entropy of k_1 decreases by 1 at each query to $\mathcal{L}_{\mathbf{k}}$.

Then, in the experiment $\mathsf{Expt}^{\mathrm{cmel}}_{\mathbf{R},\mathcal{P},\rho,\mathcal{A}}$, an adversary $\mathcal{A}$ can decrease the entropy of k_1 by more than w only if $\mathcal{A}$ succeeds in making $(w+1)$ leaking queries. This is because the leaking query is prohibited after "0" is returned,

and the entropy decreases by 1 per each query, regardless of the response. Due to Lemma D.1 and the assumption of the proposition, this happens with probability at most $1/2^w$. Therefore, we have

$$\mathsf{Adv}^{\mathrm{cmel}}_{\mathbf{R},\mathcal{P},(d-w)}(\mathcal{A}) = \Pr\left[\rho_1 < d - w \text{ in } \mathsf{Expt}^{\mathrm{cmel}}_{\mathbf{R},\mathcal{P},(d-w),\mathcal{A}}\right] \le \frac{1}{2^w}$$

when $u = 1$.

When $u > 1$, by taking the union bound regarding α on which $\mathcal{A}$ makes queries of the form $(\alpha; *)$ to $\mathcal{L}_{\mathbf{k}}$, we obtain

$$\mathsf{Adv}^{\mathrm{cmel}}_{\mathbf{R},\mathcal{P},(d-w)}(\mathcal{A}) \le \frac{M}{2^w}. \tag{33}$$

□

E Proof of Proposition 3.2

Proof. First, we prove the claim when $u = 1$ and a single key K_1 is sampled according to a distribution $\mathcal{R}_1$. By the assumptions of the proposition, $\mathcal{A}$ makes at most $Q_{L,\max}$ leakage queries, and no further query if a leakage query returns 0. Without loss of generality, we assume that $\mathcal{A}$ is deterministic. Since no further leakage query is allowed after a response 0, an execution of $\mathcal{A}$ consists of a sequence of responses 1, possibly followed by a single response 0; $\mathcal{A}$ may also stop voluntarily after a sequence of 1's. Hence, a deterministic adversary making at most $Q_{L,\max}$ leakage queries has at most $Q_{L,\max} + 1$ possible terminal transcripts. Each such transcript L corresponds to a subset $S_{\mathsf{L}} \subseteq \{0,1\}^n$ of values of K_1 compatible with L, which satisfies $\mathcal{R}_1|_{\mathsf{L}}(k) = \mathcal{R}_1(k)/\mathcal{R}_1(S_{\mathsf{L}})$ for $k \in S_{\mathsf{L}}$. These subsets S_{L} are pairwise disjoint and partition the support of $\mathcal{R}_1$.

Suppose that a transcript L is bad, namely,

$$\rho_* = \rho_1 < d - w.$$

This implies

$$\max_{k \in S_{\mathsf{L}}} \frac{\mathcal{R}_1(k)}{\mathcal{R}_1(S_{\mathsf{L}})} > 2^{-(d-w)}.$$

On the other hand, since the min-entropy of $\mathcal{R}_1$ is at least d, we have

$$\max_k \mathcal{R}_1(k) \le 2^{-d}.$$

Hence, it follows that

$$\mathcal{R}_1(S_{\mathsf{L}}) < 2^{-w}.$$

If $Q_{L,\max} \ge 2^w$, then $Q_{L,\max}/2^w \ge 1$, and the desired bound is trivial. Hence, suppose that $Q_{L,\max} < 2^w$. If all terminal transcripts were bad, then, since there are at most $Q_{L,\max} + 1$ of them,

$$1 = \sum_{\mathsf{L}} \mathcal{R}_1(S_{\mathsf{L}}) < \frac{Q_{L,\max} + 1}{2^w} \le 1,$$

which is impossible. Therefore, at least one terminal transcript is not bad. Consequently, there are at most $Q_{L,\max}$ bad terminal transcripts, and hence

$$\mathsf{Adv}^{\mathrm{cmel}}_{\mathbf{R},\mathcal{P},(d-w)}(\mathcal{A}) = \Pr[\rho_* < d - w] \le \frac{Q_{L,\max}}{2^w}.$$

Next, we prove the claim when $u > 1$. Again, we may assume without loss of generality that $\mathcal{A}$ is deterministic. Let $\mathtt{I} \subseteq \{1, \dots, u\}$ be the set of indices on which $\mathcal{A}$ makes leakage queries ($\mathtt{I}$ is a random variable; it may change with each execution of $\mathcal{A}$). By assumption, $|\mathtt{I}| \le M$ in every execution.

Fix an index $\alpha \in \{1, \dots, u\}$, and fix arbitrary values of all keys other than K_α (in their support). Whether $\mathcal{A}$ ever queries index α is then determined: until its first query at index α, all oracle responses depend only on the fixed keys. Moreover, by independence, conditioning on these other keys leaves K_α distributed according to $\mathcal{R}_\alpha$.

If index α is never queried, then $\rho_\alpha =$ (min-entropy of $\mathcal{R}_\alpha$) $\ge d$, so the event $\rho_\alpha < d - w$ cannot occur. If index α is queried, consider a single-user adversary $\mathcal{B}$ that simulates $\mathcal{A}$, answering queries at other indices using

the fixed keys and forwarding queries at index α to its own leakage oracle. The adversary $\mathcal{B}$ makes at most $Q_{L,\max}$ oracle queries and makes no further oracle query after receiving a response 0. Its leakage transcript imposes exactly the same conditions on K_α as the original experiment does. Hence, the single-user bound shows that, for these fixed other keys, the probability of $\rho_\alpha < d - w$ is at most $Q_{L,\max}/2^w$.

Averaging over the other keys, and noting that this probability is zero whenever index α is not queried, gives

$$\Pr[\rho_\alpha < d - w] \le \frac{Q_{L,\max}}{2^w} \Pr[\alpha \in \mathrm{I}]. \tag{34}$$

Therefore, by the union bound and $\sum_{\alpha=1}^{u} \Pr[\alpha \in \mathrm{I}] = \mathbb{E}[|\mathrm{I}|] \le M$, we obtain

$$\begin{aligned}\mathsf{Adv}^{\mathrm{cmel}}_{\mathbf{R},\mathcal{P},(d-w)}(\mathcal{A}) &= \Pr[\rho_* < d - w] \le \sum_{\alpha=1}^{u} \Pr[\rho_\alpha < d - w] \le \frac{Q_{L,\max}}{2^w} \sum_{\alpha=1}^{u} \Pr[\alpha \in \mathrm{I}] \\ &= \frac{Q_{L,\max}}{2^w} \mathbb{E}[|\mathrm{I}|] \le \frac{M Q_{L,\max}}{2^w}.\end{aligned}$$

□

F Tightness of the Bound in Proposition 3.2

In this section, we show that the general upper bound in Proposition 3.2 is asymptotically tight. More precisely, we exhibit a setting in which the cmel advantage matches the upper bound up to lower-order factors, provided that the total leakage-query budget is sufficiently large and that the resulting advantage is small.

Proposition F.1. *Let d and M be positive integers, let w be a non-negative integer such that $w < d$, and suppose that $1 \le Q_{L,\max} < 2^w$. Suppose moreover that $q_{L,\max} \ge M Q_{L,\max}$. Then, for any $n \ge d$, there exist $u = M$ key distributions $\mathbf{R} = (\mathcal{R}_1, \dots, \mathcal{R}_M)$ over $\{0,1\}^n$, each having min-entropy d, a set $\mathcal{P}$ of (oracle-free) predicates, and an admissible adversary[17] $\mathcal{A}$ such that*

$$\mathsf{Adv}^{\mathrm{cmel}}_{\mathbf{R},\mathcal{P},(d-w)}(\mathcal{A}) = 1 - \left(1 - \frac{Q_{L,\max}(2^{d-w} - 1)}{2^d}\right)^M.$$

In particular, in any asymptotic parameter regime[18] satisfying $d - w \to \infty$ and $M Q_{L,\max}/2^w = o(1)$, we have

$$\mathsf{Adv}^{\mathrm{cmel}}_{\mathbf{R},\mathcal{P},(d-w)}(M, q_{L,\max}, Q_{L,\max}) \ge (1 - o(1)) \frac{M Q_{L,\max}}{2^w}.$$

Proof. Let $Q := Q_{L,\max}$. Fix an arbitrary set $S \subseteq \{0,1\}^n$ of size 2^d, and let $\mathcal{R}_\alpha$ be the uniform distribution over S for every $\alpha \in \{1, \dots, M\}$. Thus, each $\mathcal{R}_\alpha$ has min-entropy exactly d.

We first partition a subset of S into Q pairwise disjoint sets $B_1, \dots, B_Q$ such that $|B_j| = 2^{d-w} - 1$ for every $j \in \{1, \dots, Q\}$. Such sets exist because $Q < 2^w$ and hence $Q(2^{d-w} - 1) < 2^d$. Let $G := S \setminus \bigcup_{j=1}^{Q} B_j$. For each $j \in \{1, \dots, Q\}$, define a predicate $P_j : \{0,1\}^n \to \{0,1\}$ by

$$P_j(k) := \begin{cases} 0 & \text{if } k \in B_j, \\ 1 & \text{otherwise.} \end{cases}$$

These predicates are oracle-free. We let

$$\mathcal{P} := \{P_1, \dots, P_Q\}.$$

Consider the following adversary $\mathcal{A}$. For each index $\alpha = 1, \dots, M$, in order, it queries $(\alpha; P_1), \dots, (\alpha; P_Q)$, stopping the queries for index α as soon as a response 0 is received. It then proceeds to the next index. The adversary therefore involves exactly M indices, makes at most Q leakage queries for each index, and makes at most $MQ \le q_{L,\max}$ leakage queries in total. Hence, it is admissible for the parameter bounds under consideration.

Fix an index α. Suppose that the first response 0 for this index occurs on the jth query. Since the sets $B_1, \dots, B_Q$ are pairwise disjoint, the leakage transcript for index α then implies $K_\alpha \in B_j$. Accordingly, the conditional

[17] That is, $\mathcal{A}$ makes at most $q_{L,\max}$ queries to $\mathcal{L}_{\mathbf{k}}$ in total, involving at most M distinct values of α, and at most $Q_{L,\max}$ queries for any fixed α.

[18] For example, assume $d = n$, $w = n/2$, $M Q_{L,\max} \ll 2^{n/2}$, and $n \to \infty$.

distribution of K_α is uniform over B_j, and its conditional min-entropy is $\log_2 |B_j| = \log_2(2^{d-w} - 1) < d - w$. Thus, every transcript ending in a response 0 is bad.

On the other hand, suppose that all Q responses for index α are 1. Then the leakage transcript implies that $K_\alpha \in G$, and the conditional distribution of K_α is uniform over G. We have $|G| = 2^d - Q(2^{d-w} - 1) = (2^w - Q)2^{d-w} + Q > 2^{d-w}$, where the last inequality follows from $Q < 2^w$. Hence, in this case the conditional min-entropy is strictly larger than $d - w$, and the transcript is not bad.

It follows that, for every fixed α, the event $\rho_\alpha < d - w$ occurs exactly when $K_\alpha \in \bigcup_{j=1}^{Q} B_j$. Therefore,

$$\Pr[\rho_\alpha < d - w] = \frac{Q(2^{d-w} - 1)}{2^d} =: p.$$

Since the keys $K_1, \ldots, K_M$ are mutually independent, the events $\rho_\alpha < d - w$ are also mutually independent. Recalling that $\rho_* = \min_{\alpha \in \{1,\ldots,M\}} \rho_\alpha$, we obtain

$$\begin{aligned}
\mathsf{Adv}^{\mathrm{cmel}}_{\mathbf{R},\mathcal{P},(d-w)}(\mathcal{A}) &= \Pr[\rho_* < d - w] \\
&= 1 - \Pr[\rho_\alpha \geq d - w \text{ for all } \alpha \in \{1, \ldots, M\}] \\
&= 1 - (1 - p)^M \\
&= 1 - \left(1 - \frac{Q(2^{d-w} - 1)}{2^d}\right)^M.
\end{aligned}$$

This proves the first two claims.

It remains to consider the stated asymptotic regime. We have $p = \frac{Q}{2^w}\left(1 - 2^{-(d-w)}\right)$. Thus, if $d - w \to \infty$, $p = (1 - o(1))\frac{Q}{2^w}$. Moreover, the assumption $MQ/2^w = o(1)$ implies $Mp = o(1)$. Consequently,

$$1 - (1 - p)^M = (1 - o(1))Mp = (1 - o(1))\frac{MQ}{2^w}.$$

Substituting $Q = Q_{L,\max}$ gives

$$\mathsf{Adv}^{\mathrm{cmel}}_{\mathbf{R},\mathcal{P},(d-w)}(M, q_{L,\max}, Q_{L,\max}) \geq (1 - o(1))\frac{MQ_{L,\max}}{2^w},$$

as claimed. □

G Comparison of Hybrids

Table 5. Brief comparison of $\mathsf{Hyb}_i^{5\text{-}j}$ and $\mathsf{Hyb}_i^{5\text{-}(j+1)}$ between the $(i+1)$st and $(i+2)$nd construction queries, when $\mathsf{NumL}_{i+1}^{<j}$ and $\mathsf{NumL}_{i+1}^{<(j+1)}$ do not happen, $i \leq q_C - 2$, and $j > 0$. See also Table 6.

	$(i+1)$st const. query	prim. query	1st leak. query	prim. query	2nd leak. query	...
$\mathsf{Hyb}_i^{5\text{-}j}$	$\widetilde{\mathrm{RP}}$	$\lvert E^{\pm}\rangle$	$\mathcal{L}_{\mathbf{k}}$	$\lvert E^{\pm}\rangle$	$\mathcal{L}_{\mathbf{k}}$	⋯
$\mathsf{Hyb}_i^{5\text{-}(j+1)}$	$\widetilde{\mathrm{RP}}$	$\lvert E^{\pm}\rangle$	$\mathcal{L}_{\mathbf{k}}$	$\lvert E^{\pm}\rangle$	$\mathcal{L}_{\mathbf{k}}$	⋯

...	jth leak. query		prim. query	$(j+1)$st leak. query	...	prim. query	$(i+2)$nd const. query
⋯	$\mathcal{L}_{\mathbf{k}}$	resample $\mathbf{k}$	$\lvert\left(E^{\mathsf{T}_{i+1},\mathbf{k}}\right)^{\pm}\rangle$	$\mathcal{L}_{\mathbf{k}}$	⋯	$\lvert\left(E^{\mathsf{T}_{i+1},\mathbf{k}}\right)^{\pm}\rangle$	$\overline{\mathsf{TBC}}_{\mathbf{k}}^{E^{\mathsf{T}_{i+1},\mathbf{k}}}$
⋯	$\mathcal{L}_{\mathbf{k}}$		$\lvert E^{\pm}\rangle$	$\mathcal{L}_{\mathbf{k}}$	resample $\mathbf{k}$ ⋯	$\lvert\left(E^{\mathsf{T}_{i+1},\mathbf{k}}\right)^{\pm}\rangle$	$\overline{\mathsf{TBC}}_{\mathbf{k}}^{E^{\mathsf{T}_{i+1},\mathbf{k}}}$

Table 6. Invariants in the transition from $\mathsf{Hyb}_i^{5\text{-}j}$ to $\mathsf{Hyb}_i^{5\text{-}(j+1)}$. We consider the case $i \le q_C - 2$, $j > 0$, where both the j-th and $(j+1)$-st leaking queries occur between the $(i+1)$-st and $(i+2)$-nd construction queries. Here, $\mathsf{L}_{i+1}^{(j)}$ denotes the leaking transcript up to and including the j-th such leaking query, and $\mathbf{R}\,|_{\mathsf{L}_{i+1}^{(j)}}$ denotes the tuple of conditional key distributions determined by the leakage transcript $\mathsf{L}_{i+1}^{(j)}$. The boundary cases are analogous. See also Table 5.

	$\mathsf{Hyb}_i^{5\text{-}j}$	$\mathsf{Hyb}_i^{5\text{-}(j+1)}$	Invariant / difference
Up to the response to the j-th leaking query	Uses $\widetilde{\mathrm{RP}}$, $\mathcal{L}_{\mathbf{k}}$, and $\lvert E^{\pm}\rangle$.	Uses $\widetilde{\mathrm{RP}}$, $\mathcal{L}_{\mathbf{k}}$, and $\lvert E^{\pm}\rangle$.	The two hybrids are identical up to this point. In particular, T_{i+1} and $\mathsf{L}_{i+1}^{(j)}$ have the same distribution.
Immediately after the j-th leaking response	Resample $\mathbf{k} \leftarrow \mathbf{R}\|_{\mathsf{L}_{i+1}^{(j)}}$ and switch the primitive oracle to $\lvert(E^{\mathsf{T}_{i+1},\mathbf{k}})^{\pm}\rangle$.	No resampling or reprogramming is performed yet.	Resampling and reprogramming are performed only in $\mathsf{Hyb}^{5\text{-}j}$.
Between the j-th and $(j+1)$-st leaking queries	Primitive queries are answered using $\lvert(E^{\mathsf{T}_{i+1},\mathbf{k}})^{\pm}\rangle$.	Primitive queries are answered using $\lvert E^{\pm}\rangle$.	This is the only interval in which the responses from the primitive oracles may differ. The arbitrary reprogramming lemma is applied to bound the difference in this interval.
At the $(j+1)$-st leaking query	The response is computed using the resampled key $\mathbf{k}$.	The response is computed using the key $\mathbf{k}$ sampled at the beginning of the game.	$\mathsf{Hyb}_i^{5\text{-}j}$ uses the resampled key, while $\mathsf{Hyb}_i^{5\text{-}(j+1)}$ uses the key sampled at the beginning. (This difference causes no additional security loss, which is essentially shown in the proof of Eq. (27).)
Immediately after the $(j+1)$-st leaking response	Keep the current keys and oracles.	Resample $\mathbf{k} \leftarrow \mathbf{R}\|_{\mathsf{L}_{i+1}^{(j+1)}}$ and switch the primitive oracle to $\lvert(E^{\mathsf{T}_{i+1},\mathbf{k}})^{\pm}\rangle$.	The two hybrids again have the same form. (The construction oracles are also replaced with $\overline{\mathsf{TBC}}_{\mathbf{k}}^{E^{\mathsf{T}_{i+1},\mathbf{k}}}$ in both hybrids.)

Table 7. Overview of the hybrids around the $(i+1)$-st construction query. Here $Q_{i+1} = (\alpha_{i+1}; x_{i+1}, t_{i+1}, b_{i+1})$ denotes the $(i+1)$-st construction query, and $\mathbf{R}\,|_{\mathsf{L}_{i+1}}$ denotes the tuple of conditional key distributions determined by the leakage transcript L_{i+1}.

	Hyb_i^1	Hyb_i^2	Hyb_i^3	Hyb_i^4	Hyb_i^5
Before the $(i+1)$-st construction query	$\mathbf{k} \leftarrow \mathbf{R}$ at the beginning; the construction oracle is $\widetilde{\mathrm{RP}}$, the leakage oracle is $\mathcal{L}_{\mathbf{k}}$, and the primitive oracle is $\lvert E^{\pm}\rangle$.				
When $Q_{i+1} = (\alpha_{i+1}; x_{i+1}, t_{i+1}, b_{i+1})$ is submitted	Resample $\mathbf{k} \leftarrow \mathbf{R}\,\lvert_{\mathsf{L}_{i+1}}$.	Resample $\mathbf{k} \leftarrow \mathbf{R}\,\lvert_{\mathsf{L}_{i+1}}$; $s_0^* := \psi_1(k_{\alpha_{i+1}}, x_{i+1})$, $s_1^* \xleftarrow{\$} \{0,1\}^n$, $E' := E \circ \mathsf{swap}^{t_{i+1}}_{s_0^*, s_1^*}$.	Same as Hyb_i^2.	Resample $\mathbf{k} \leftarrow \mathbf{R}\,\lvert_{\mathsf{L}_{i+1}}$; $s_0^* := \psi_1(k_{\alpha_{i+1}}, x_{i+1})$.	No resampling yet.
Response to Q_{i+1}	Answer using $\overline{\mathsf{TBC}}_{\mathbf{k}}^{E^{\mathsf{T}_i,\mathbf{k}}}$.	Answer using $\overline{\mathsf{TBC}}_{\mathbf{k}}^{(E')^{\mathsf{T}_i,\mathbf{k}}}$.	Return $\widetilde{y} := \psi_2\big(k_{\alpha_{i+1}}, E^{\mathsf{T}_i,\mathbf{k}}(t_{i+1}, s_1^*), b_{i+1}\big)$.	Return $y_{i+1} := \widetilde{\mathrm{RP}}(Q_{i+1})$, then set $s_1^* := (E^{\mathsf{T}_i,\mathbf{k}})^{-1}(t_{i+1}, \psi_2(k_{\alpha_{i+1}}, y_{i+1}, b_{i+1}))$, $E' := E \circ \mathsf{swap}^{t_{i+1}}_{s_0^*, s_1^*}$.	Return $y_{i+1} := \widetilde{\mathrm{RP}}(Q_{i+1})$.
After the response to Q_{i+1}	Use $\overline{\mathsf{TBC}}_{\mathbf{k}}^{E^{\mathsf{T}_i,\mathbf{k}}}$ and $\lvert(E^{\mathsf{T}_i,\mathbf{k}})^{\pm}\rangle$.	Use $\overline{\mathsf{TBC}}_{\mathbf{k}}^{(E')^{\mathsf{T}_i,\mathbf{k}}}$ and $\lvert((E')^{\mathsf{T}_i,\mathbf{k}})^{\pm}\rangle$.	Use $\overline{\mathsf{TBC}}_{\mathbf{k}}^{(E')^{\mathsf{T}_i,\mathbf{k}}}$ and $\lvert((E')^{\mathsf{T}_i,\mathbf{k}})^{\pm}\rangle$.	Use $\overline{\mathsf{TBC}}_{\mathbf{k}}^{(E')^{\mathsf{T}_i,\mathbf{k}}}$ and $\lvert((E')^{\mathsf{T}_i,\mathbf{k}})^{\pm}\rangle$.	Resample $\mathbf{k} \leftarrow \mathbf{R}\,\lvert_{\mathsf{L}_{i+1}}$; use $\overline{\mathsf{TBC}}_{\mathbf{k}}^{E^{\mathsf{T}_{i+1},\mathbf{k}}}$ and $\lvert(E^{\mathsf{T}_{i+1},\mathbf{k}})^{\pm}\rangle$.